

Intel[®] RAID Software User Guide

- Intel[®] Embedded Server RAID Technology 2
- Intel[®] IT/IR RAID
- Intel[®] Integrated Server RAID
- Intel[®] RAID Software Stack 3

Revision 23.0

July, 2014

Intel Order Number: D29305-023

DISCLAIMER

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.

A "Mission Critical Application" is any application in which failure of the Intel Product could result, directly or indirectly, in personal injury or death. SHOULD YOU PURCHASE OR USE INTEL'S PRODUCTS FOR ANY SUCH MISSION CRITICAL APPLICATION, YOU SHALL INDEMNIFY AND HOLD INTEL AND ITS SUBSIDIARIES, SUBCONTRACTORS AND AFFILIATES, AND THE DIRECTORS, OFFICERS, AND EMPLOYEES OF EACH, HARMLESS AGAINST ALL CLAIMS COSTS, DAMAGES, AND EXPENSES AND REASONABLE ATTORNEYS' FEES ARISING OUT OF, DIRECTLY OR INDIRECTLY, ANY CLAIM OF PRODUCT LIABILITY, PERSONAL INJURY, OR DEATH ARISING IN ANY WAY OUT OF SUCH MISSION CRITICAL APPLICATION, WHETHER OR NOT INTEL OR ITS SUBCONTRACTOR WAS NEGLIGENT IN THE DESIGN, MANUFACTURE, OR WARNING OF THE INTEL PRODUCT OR ANY OF ITS PARTS.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. The information here is subject to change without notice. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained by calling 1-800-548-4725, or go to:
<http://www.intel.com/design/literature.htm>.

Intel warrants that this product will perform to its published specifications. However, all computer systems are inherently subject to unpredictable system behavior under various environmental and other conditions.

This product is not intended to be the sole source for any critical data and the user must maintain a verified backup. Failure to do so or to comply with other user notices in the product user guide and specification documents may result in loss of or access to data.

Table of Contents

1	Overview	1
1.1	Supported Hardware	1
1.2	Software	4
1.3	RAID Terminology	5
1.3.1	Fault Tolerance	5
1.3.2	Enclosure Management	7
1.3.3	Performance	7
2	RAID Levels	9
2.1	Summary of RAID Levels	9
2.2	Selecting a RAID Level	10
2.2.1	RAID 0 - Data Striping	10
2.2.2	RAID 1 - Disk Mirroring/Disk Duplexing	11
2.2.3	RAID 5 - Data Striping with Striped Parity	11
2.2.4	RAID 6 - Distributed Parity and Disk Striping	12
2.2.5	RAID IME	14
2.2.6	RAID 10 - Combination of RAID 1 and RAID 0	14
2.2.7	RAID 50 - Combination of RAID 5 and RAID 0	15
2.2.8	RAID 60 - Combination of RAID 0 and RAID 6	16
2.3	RAID Configuration Strategies	19
2.3.1	Maximizing Fault Tolerance	19
2.3.2	Maximizing Performance	20
2.3.3	Maximizing Storage Capacity	22
2.4	RAID Availability	23
2.4.1	RAID Availability Concept	23
2.4.2	Spare Drives	23
2.4.3	Rebuilding	23
2.4.4	Drive in Foreign State	24
2.4.5	Copyback	24
2.4.6	Configuration Planning	25
2.4.7	Dimmer Switch Feature	25
2.4.8	Serial ATA III Features	25
2.4.9	UEFI 2.0 Support	26
2.4.10	Number of Physical Disks	26
2.4.11	MegaRAID Fast Path	26
2.4.12	4K Sector Drive Support	26
2.4.13	Larger than 2TB Drive Support	27
2.4.14	Power Save settings	27
2.4.15	Shield State	27
2.4.16	Array Purpose	28

3	RAID Utilities	29
3.1	Intel® Embedded Server RAID Technology II BIOS Configuration Utility	29
3.2	LSI MPT* SAS BIOS Configuration Utility	29
3.3	Intel® RAID BIOS Console 2 Configuration Utility for Intelligent RAID	30
3.4	Intel® RAID Web Console 2 Configuration and Monitoring Utility	31
3.4.1	Drive Hierarchy within the RAID Firmware	32
3.5	Intel® Intelligent RAID Controller Features	32
3.5.1	Enterprise Features	32
3.5.2	Fault Tolerant Features	33
3.5.3	Cache Options and Settings	34
3.5.4	Background Tasks	34
3.5.5	Error Handling	35
3.5.6	Audible Alarm	35
4	Intel® RAID Drivers	37
4.1	RAID Driver Installation for Microsoft Windows*	37
4.1.1	Installation in a New Microsoft Windows* Operating System	37
4.1.2	Installation in an Existing Microsoft Windows* Operating System	38
4.2	RAID Driver Installation for Red Hat* Enterprise Linux	38
4.3	RAID Driver Installation for SuSE* Linux	39
4.4	RAID Driver Installation for Novell NetWare*	39
4.4.1	Installation in a New Novell Netware* System	39
4.4.2	Installation in an Existing Novell Netware* System	40
4.5	RAID Driver Installation for Solaris* 10	40
4.5.1	Installation in a New Solaris* System	40
4.5.2	Installation in an Existing Solaris* System	41
5	Intel® Embedded Server RAID BIOS Configuration Utility	42
5.1	Creating, Adding or Modifying a Virtual Drive Configuration	43
5.2	Setting the Write Cache and Read Ahead Policies	45
5.3	Working with a Global Hot-spare Drive	46
5.3.1	Adding a Hot-spare Drive	46
5.3.2	Removing a Hot-spare Drive	46
5.4	Rebuilding a Drive	46
5.4.1	Auto Rebuild and Auto Resume	47
5.5	Checking Data Consistency	47
5.6	Viewing and Changing Device Properties	47
5.7	Forcing a Drive Online or Offline	48
5.8	Configuring a Bootable Virtual Drive	48
5.9	Deleting (Clearing) a Storage Configuration	49
6	Intel® IT/IR RAID Controller Configuration Utility	50
6.1	Overview	50
6.1.1	Integrated Mirroring and Integrated Mirroring Enhanced Features	50
6.1.2	Integrated Striping Features	51

6.2	Creating Mirrored Volumes	51
6.2.1	Mirrored Volume Configuration Overview	52
6.2.2	Creating Mirrored Volumes	52
6.2.3	Managing Hot Spare Disks	56
6.2.4	Other Configuration Tasks	58
6.3	Creating Integrated Striping Volumes	61
6.3.1	Integrated Striping Configuration Overview	62
6.3.2	Creating Integrated Striping Volumes	62
6.3.3	Other Configuration Tasks	64
7	12Gb/s Intel® RAID Controller Configuration Utility	68
7.1	Overview	68
7.2	Starting the Ctrl-R Utility	68
7.3	Exiting the Ctrl-R Utility	69
7.4	Ctrl-R Utility Keystrokes	69
7.5	Ctrl-R Utility Menus	70
7.5.1	VD Mgmt Menu	70
7.5.2	PD Mgmt Menu	71
7.5.3	Ctrl Mgmt Menu	72
7.5.4	Properties Menu	73
7.5.5	Foreign View Menu	74
7.6	Controller Advanced Software Options	75
7.6.1	Advanced Software Options	75
7.6.2	RAID Premium Feature Key	77
7.7	Creating a Storage Configuration	78
7.7.1	Selecting Additional Virtual Drive Properties	80
7.7.2	Creating a CacheCade Virtual Drive	81
7.7.3	Modifying a CacheCade Virtual Drive	82
7.7.4	Creating a CacheCade Pro 2.0 Virtual Drive	83
7.7.5	Modifying a CacheCade Pro 2.0 Virtual Drive	84
7.7.6	Enabling SSD Caching on a Virtual Drive	86
7.7.7	Disabling SSD Caching on a Virtual Drive	86
7.7.8	Enabling or Disabling SSD Caching on Multiple Virtual Drives	87
7.7.9	Deleting a Virtual Drive with SSD Caching Enabled	88
7.8	Clearing the Configuration	88
7.9	Intel SafeStore Encryption Services	89
7.9.1	Enabling Drive Security	89
7.9.2	Changing Security Settings	91
7.9.3	Disabling Drive Security	92
7.9.4	Importing or Clearing a Foreign Configuration	92
7.10	Discarding Preserved Cache	95
7.11	Converting JBOD Drives to Unconfigured Good Drives	96
7.12	Converting Unconfigured Good Drives to JBOD Drives	97
7.13	Viewing and Changing Device Properties	99

7.13.1	Viewing Controller Properties	99
7.13.2	Modifying Controller Properties	99
7.13.3	Viewing and Changing Virtual Drive Properties	101
7.13.4	Deleting a Virtual Drive	103
7.13.5	Deleting a Virtual Drive Group	103
7.13.6	Expanding a Virtual Drive	104
7.13.7	Erasing a Virtual Drive	105
7.13.8	Managing Link Speed	106
7.13.9	Managing Power Save Settings for the Controller	107
7.13.10	Managing Power Save Settings for the Drive Group	108
7.13.11	Managing BBU Information	108
7.13.12	Managing Dedicated Hot Spares	110
7.13.13	Securing a Drive Group	111
7.13.14	Setting LED Blinking	111
7.13.15	Performing a Break Mirror Operation	111
7.13.16	Performing a Join Mirror Operation	112
7.14	Managing Storage Configurations	114
7.14.1	Initializing a Virtual Drive	114
7.14.2	Running a Consistency Check	114
7.14.3	Rebuilding a Physical Drive	115
7.14.4	Performing a Copyback Operation	116
7.14.5	Removing a Physical Drive	116
7.14.6	Creating Global Hot Spares	117
7.14.7	Removing a Hot Spare Drive	117
7.14.8	Making a Drive Offline	117
7.14.9	Making a Drive Online	118
7.14.10	Instant Secure Erase	118
7.14.11	Erasing a Physical Drive	118
8	6Gb/s Intel® RAID Controller Configuration Utility	120
8.1	Quick Configuration Steps	120
8.2	Detailed Configuration Steps using the Intel® RAID BIOS Console 2	121
8.2.1	Start the Intel® RAID BIOS Console 2 Utility	121
8.2.2	Screen and Option Descriptions	121
8.3	Setting Up a RAID Array Using the Configuration Wizard	128
8.4	Creating RAID 0, 1, 5, or 6 using Intel® RAID BIOS Console 2 (detailed)	131
8.5	Creating RAID 10, RAID 50, and RAID 60 using Intel® RAID BIOS Console 2	135
8.6	Setting Drive Parameters	141
8.7	Creating a Hot Spare	142
8.8	Viewing Event Details	145
9	Intel® RAID Web Console 2	148
9.1	Configuration Functions	148
9.2	Monitoring Functions	148
9.3	Maintenance Functions	149

9.4	Hardware and Software Requirements	149
9.5	Installing the Intel® RAID Web Console 2 on a Microsoft Windows* Operating System 149	
9.6	Installing the Intel® RAID Web Console 2 on Linux or SuSE* Linux Enterprise Server . 151	
9.7	Intel® RAID Web Console 2 Support and Installation on VMWare	152
9.7.1	Installing Intel® RAID Web Console 2 for VMWare Classic	152
9.7.2	Uninstalling Intel® RAID Web Console 2 for VMWare	152
9.7.3	Installing Intel® RAID Web Console 2 Support on the VMWare ESX	152
9.8	Starting the Intel® RAID Web Console 2	158
9.9	Intel® RAID Web Console 2 Screens	158
9.9.1	LDAP Support	161
9.9.2	Configuring LDAP Support Settings	163
9.9.3	Physical/Virtual View Panel	165
9.9.4	Properties/Operations/Graphical View Panel	165
9.9.5	Event Log Panel	167
9.9.6	Menu Bar/Manage Menu	167
9.9.7	Menu Bar/Go To Menu	168
9.9.8	File Menu/Log Menu	168
9.9.9	File Menu/Tool Menu	168
9.9.10	File Menu/Help Menu	168
9.10	Drive Configuration Tasks	168
9.10.1	Create Virtual Drive	169
9.11	Creating a Virtual Drive Using Simple Configuration	169
9.12	Creating a Virtual Drive Using Advanced Configuration	174
9.13	Creating a Spanned Disk Group	183
9.14	Creating Hot Spares	184
9.15	Setting Adjustable Task Rates	186
9.16	Adding a Drive to a Virtual Disk	188
9.17	Removing a Drive or Changing the RAID Level of a Virtual Disk	191
9.18	Changing Virtual Disk Properties	193
9.19	Deleting a Virtual Disk	194
9.20	Managing Configurations	195
9.20.1	Saving a Configuration to Disk	195
9.20.2	Clearing a Configuration from a Controller	197
9.20.3	Adding a Configuration from a File	198
9.21	Monitoring System Events and Devices	199
9.21.1	Monitoring System Events	199
9.21.2	Monitoring Controllers	201
9.21.3	Monitoring Disk Drives and Other Physical Devices	202
9.21.4	Monitoring Virtual Disks	205
9.21.5	Monitoring Enclosures	206
9.21.6	Monitoring Battery Backup Units	206
9.21.7	Battery Learn Cycle	207

9.21.8	Monitoring Rebuilds and Other Processes	209
9.22	Maintaining and Managing Storage Configurations	211
9.22.1	Initializing a Virtual Disk	211
9.22.2	Running a Consistency Check	212
9.22.3	Scanning for New Drives	214
9.22.4	Rebuilding a Drive	214
9.22.5	Removing a Drive	216
9.22.6	Flashing the Firmware	216
9.23	Enabling RAID Premium Features	216
9.23.1	Enabling Full Disk Encryption feature	216
9.23.2	Enabling Snapshot feature	230
9.23.3	Enabling Super Sized Cache	243
10	StorCLI	248
10.1	Overview	248
10.2	Devices Supported by the StorCLI Tool	248
10.3	Installation	248
10.3.1	Installing StorCLI on Microsoft Windows Operating Systems	249
10.3.2	Installing StorCLI on Linux Operating Systems	249
10.3.3	Installing StorCLI on Ubuntu Operating Systems	249
10.3.4	Installing StorCLI on VMware Operating Systems	250
10.3.5	Installing StorCLI on FreeBSD Operating Systems	250
10.3.6	Installing StorCLI on Microsoft EFI	250
10.3.7	Installing StorCLI on Solaris Operating Systems	250
10.4	StorCLI Command Syntax	251
10.5	Working with the Storage Command Line Tool	253
10.5.1	System Commands	253
10.5.2	Controller Commands	254
10.5.3	Drive Commands	268
10.5.4	Virtual Drive Commands	277
10.5.5	Foreign Configurations Commands	290
10.5.6	BIOS-Related Commands	290
10.5.7	Drive Group Commands	292
10.5.8	Dimmer Switch Commands	293
10.5.9	BBU Commands	294
10.5.10	Enclosure Commands	297
10.5.11	PHY Commands	298
10.5.12	Logging Commands	298
10.6	Frequently Used Tasks	300
10.6.1	Showing the Version of the Storage Command Line Tool	300
10.6.2	Showing StorCLI Help	300
10.6.3	Showing System Summary Information	300
10.6.4	Showing Free Space in a Controller	300
10.6.5	Adding Virtual Drives	300

10.6.6	Setting the Cache Policy in a Virtual Drive	301
10.6.7	Showing Virtual Drive Information	301
10.6.8	Deleting Virtual Drives	302
10.6.9	Flashing Controller Firmware	302
10.7	Unsupported StorCLI Commands	303
Appendix A: Creating a Virtual Drive Using Advanced Configuration		306
Appendix B: Events and Messages		315
Glossary		332

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	43
Figure 10. Adapter Properties Window	53
Figure 11. Manage Volume Window.....	57
Figure 12. Boot Device on SAS Topology Window.....	61
Figure 13. Adapter Properties Window	63
Figure 14. Boot Device on SAS Topology Window.....	66
Figure 15. VD Mgmt.....	71
Figure 16. Physical Drive Management.....	72
Figure 17. Controller Settings – First Screen.....	73
Figure 18. Controller Settings – Second Screen.....	73
Figure 19. Properties	74
Figure 20. Foreign View Menu.....	75
Figure 21. Advanced Software Options	76
Figure 22. Advanced Software Options With Key AXXRPFKSSD2 Installed	76
Figure 23. Create a New Virtual Drive	78
Figure 24. Create Virtual Drive – Advanced	80
Figure 25. Create CacheCade Virtual Drive	82
Figure 26. Modify CacheCade Virtual Drive.....	83
Figure 27. Create CacheCade Virtual Drive	84
Figure 28. Modify CacheCade Virtual Drive.....	85
Figure 29. Associated Virtual Drives.....	85
Figure 30. Message Box for Enabling SSD Caching	86
Figure 31. Message Box for Disabling SSD Caching	87
Figure 32. Manage SSD Caching	87
Figure 33. Message Box for Deleting Virtual Drive	88
Figure 34. Clear Configuration.....	89
Figure 35. Create Security Key.....	90
Figure 36. Change Security Key.....	91
Figure 37. Foreign Configuration Options.....	93
Figure 38. Foreign Configuration – Import	94
Figure 39. Foreign Configuration – Clear.....	94
Figure 40. Manage Preserved Cache	96
Figure 41. Convert JBOD to Unconfigured Good	97
Figure 42. Convert Unconfigured Good to JBOD	98

Figure 43. Virtual Drive Properties	102
Figure 44. Create Virtual Drive – Advanced.....	103
Figure 45. Expand Virtual Drive	104
Figure 46. Erase Virtual Drive	105
Figure 47. Manage Link Speed	106
Figure 48. Manage Power Save.....	107
Figure 49. Manage Power Save Settings – Drive Group	108
Figure 50. Battery Properties	109
Figure 51. Dedicated Hotspare	110
Figure 52. Join Mirror – Choose Option	112
Figure 53. Confirmation Message	113
Figure 54. Confirmation Message	113
Figure 55. Join Mirror – Choose Option	113
Figure 56. Initialize a Virtual Drive.....	114
Figure 57. Consistency Check	115
Figure 58. Copyback Operation	116
Figure 59. Intel® RAID BIOS Console 2 Menu.....	123
Figure 60. Intel® RAID BIOS Console 2 – Controller Selection	124
Figure 61. Controller Properties	124
Figure 62. Additional Controller Properties	125
Figure 63. Intel® RAID BIOS Console 2 – Configuration Types.....	129
Figure 64. Selecting Configuration.....	129
Figure 65. Intel® RAID BIOS Console 2 – Configuration Methods.....	130
Figure 66. Intel® RAID BIOS Console 2 – Add Physical Drives to Array	131
Figure 67. Intel® RAID BIOS Console 2 – Set Array Properties	132
Figure 68. Intel® RAID BIOS Console 2 – Confirm Configuration.....	133
Figure 69. Intel® RAID BIOS Console 2 – Initialization Speed Setting	134
Figure 70. Intel® RAID BIOS Console 2 – Multiple Disk Groups for RAID 10, 50, or 60.....	136
Figure 71. Intel® RAID BIOS Console 2 – Spanning Multiple Arrays.....	137
Figure 72. Intel® RAID BIOS Console 2 – Viewing Completed Settings.....	137
Figure 73. Intel® RAID BIOS Console 2 – Initialization Settings	138
Figure 74. Intel® RAID BIOS Console 2 – RAID 10 Final Screen	139
Figure 75. Intel® RAID BIOS Console 2 – RAID 10 Properties Screen	139
Figure 76. Intel® RAID BIOS Console 2 – RAID 50 Properties Screen	140
Figure 77. Intel® RAID BIOS Console 2 – Choosing a Hot Spare Drive	142
Figure 78. Intel® RAID BIOS Console 2 – Setting a Hot Spare Drive	143
Figure 79. Intel® RAID BIOS Console 2 – Viewing Hot Spare	144
Figure 80. Intel® RAID BIOS Console 2 – Main Screen showing Hot Spare Drive.....	144
Figure 81. Intel® RAID BIOS Console 2 – Event Information Screen	145
Figure 82. Intel® RAID BIOS Console 2 – Selecting Events to View	146
Figure 83. Intel® RAID BIOS Console 2 – Viewing an Event.....	147
Figure 84. Intel® RAID Web Console 2 – Customer Information Screen	150
Figure 85. Setup Type Screen	151
Figure 86. Intel® RAID Web Console 2 – Select Server Screen	158
Figure 87. Intel® RAID Web Console 2 – Login Screen.....	159
Figure 88. Intel® RAID Web console 2 dashboard.....	160
Figure 89. Intel® RAID Web Console 2 – Main Screen.....	161

Figure 90. LDAP Login.....	162
Figure 91. Configure Host LDAP	164
Figure 92. Intel® RAID Web Console 2 – Operations Tab.....	166
Figure 93. Intel® RAID Web Console 2 – Graphical Tab (Optional feature).....	167
Figure 94. Virtual Drive Creation Menu.....	170
Figure 95. Virtual Drive Creation Mode.....	171
Figure 96. Create Virtual Drive Screen.....	172
Figure 97. Create Virtual Drive – Summary Window	173
Figure 98. Option to Create Additional Virtual Drives	173
Figure 99. Option to Close the Configuration Wizard	174
Figure 100. Virtual Drive Creation Menu.....	175
Figure 101. Virtual Drive Creation Mode.....	176
Figure 102. Create Drive Group Settings Screen	177
Figure 103. Span 0 of Drive Group 0.....	178
Figure 104. Span 0 and Span 1 of Drive Group 0.....	179
Figure 105. Virtual Drive Settings Window	180
Figure 106. New Virtual Drive 0.....	181
Figure 107. Create Virtual Drive Summary Window	181
Figure 108. Option to Create Additional Virtual Drives	182
Figure 109. Option to Close the Configuration Wizard	182
Figure 110. Assign Global Hotspare	184
Figure 111. Assign Dedicated Hotspare	185
Figure 112. Select Hotspare Drive.....	186
Figure 113. Set Adjustable Task Rates	187
Figure 114. Starting Modify Drive Group	189
Figure 115. Select RAID level to migrate.....	189
Figure 116. Selecting Drives to Add	190
Figure 117. Changing RAID Level	191
Figure 118. Selecting Drives to Remove	192
Figure 119. Changing RAID Level	193
Figure 120. Set Virtual Disk Properties.....	194
Figure 121. Save Configuration to File	195
Figure 122. Save Configuration Dialog Box.....	196
Figure 123. Clear Configuration.....	197
Figure 124. Add Saved Configuration.....	198
Figure 125. Event Information Window.....	200
Figure 126. Controller Information	201
Figure 127. Physical Drive Information	202
Figure 128. Locating a Physical Drive	203
Figure 129. Patrol Read Configuration	204
Figure 130. Virtual Drive Properties.....	205
Figure 131. Enclosure Information.....	206
Figure 132. Battery Backup Unit Information	207
Figure 133. Battery Backup Unit Operations	208
Figure 134. Group Show Progress Window.....	209
Figure 135. Selecting Initialize	212
Figure 136. Group Consistency Check Window	213

Figure 137. Scan for Foreign Configuration	214
Figure 138. Preparing Drive for Removal.....	215
Figure 139. Check Controller Security status.....	218
Figure 140. Check Drive Security status	219
Figure 141. Enable Drive Security.....	220
Figure 142. Start Security Wizard	221
Figure 143. Enter Security Key Identifier.....	222
Figure 144. Enter Security Key	223
Figure 145. Enter Pass Phrase	224
Figure 146. Confirm Enable Drive Security	225
Figure 147. Check Drive Security Enabled status.....	226
Figure 148. Select Full Disk Encryption	227
Figure 149. Create RAID Virtual Drive with FDE enabled.....	228
Figure 150. Instant Secure Erase.....	229
Figure 151. Confirm Secure Erase.....	230
Figure 152. Enable MegaRAID Recovery	231
Figure 153. Enter the Capacity for Snapshot Respository	232
Figure 154. Confirm Enable Snapshot	233
Figure 155. Snapshot Base is shown.....	234
Figure 156. Enter Snapshot Name.....	235
Figure 157. Create Snapshot	235
Figure 158. Create View.....	236
Figure 159. Set MegaRAID Recovery Properties.....	237
Figure 160. Disable MegaRAID Recovery	238
Figure 161. Confirm Disable Snapshots.....	238
Figure 162. Adapter Selection.....	239
Figure 163. Selecting Snapshot Base	239
Figure 164. Selecting Advanced Operations.....	240
Figure 165. Selecting Rollback.....	241
Figure 166. Selecting a snapshot.....	242
Figure 167. Confirm Page	242
Figure 168. Rollback operation is done.....	243
Figure 169. Create SSC from Dashboard	244
Figure 170. Create SSC Drive Group.....	244
Figure 171. Create SSCD name	245
Figure 172. SSC Summary	246
Figure 173. SSCD status shown	246
Figure 174. Delete SSCD.....	247
Figure 175. Virtual Drive Creation Menu	307
Figure 176. Virtual Drive Creation Mode	308
Figure 177. Create Drive Group Settings Screen.....	309
Figure 178. Span 0 of Drive Group 0	310
Figure 179. Span 0 and Span 1 of Drive Group 0	311
Figure 180. Virtual Drive Settings Window	312
Figure 181. New Virtual Drive 0	313
Figure 182. Create Virtual Drive Summary Window.....	313
Figure 183. Option to Create Additional Virtual Drives.....	314

Figure 184. Option to Close the Configuration Wizard 314

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	13
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	28
Table 13. Ctrl-R Utility Keystrokes	69
Table 14. RAID Premium Feature Key	77
Table 15. Controller Settings	99
Table 16. Intel® RAID BIOS Console 2 Toolbar Icon Descriptions	121
Table 17. Object Identifiers in the StorCli Command Syntax	251
Table 18. Verbs in the StorCli Command Syntax	252
Table 19. Controller Commands Quick Reference Table	254
Table 20. Properties for Show and Set Commands	257
Table 21. Set Patrolread Input Options	262
Table 22. Set CC Input Options	263
Table 23. Set Advanced Software Options Input Options	265
Table 24. Set Security Key Input Options	266
Table 25. Flashing Controller Firmware Input Options	267
Table 26. Physical Drives Commands Quick Reference Table	268
Table 27. Drive Erase Command Options	274
Table 28. Add Hotsparedrive Input Options	276
Table 29. Virtual Drives Commands Quick Reference Table	277
Table 30. Add RAID Configuration Input Options	278
Table 31. Add RAID Configuration Input Options	280
Table 32. Virtual Drive Migration Command Options	286
Table 33. Virtual Drive Migration Table	286
Table 34. Dimmer Switch Input Options	293
Table 35. BBU Mode	295
Table 36. Enclosure Firmware Download Command Options	297
Table 37. Unsupported Commands for Intel® RAID Controller	303
Table 38. MFI Messages	316

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.

1.1 Supported Hardware

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

- Intel® Embedded Server RAID Technology II (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 II 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 II 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 II modes.

Intel® Embedded Server RAID Technology II must be enabled in the server system BIOS before it is available. Intel® Embedded Server RAID Technology II 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 the following 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 Module RMS25KB080, RMS25KB040, RMS25JB080, and RMS25JB040
- Intel® RAID Controller RS25FB044
- 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 Intel® RAID controllers RMS25PB080, RMS25PB040, RMT3PB080, RMS25CB080, RMS25CB040, RMT3CB080, RS25AB080, RS25SB008, RS25DB080, RS25NB008, RS2VB080, RS2VB040, RT3WB080, RS2SG244, RS2WG160, RMS2MH080, RMS2AF080,

RMS2AF040, RS2BL080, RS2BL040, RS2BL080DE, RS2BL080SNGL, RS2PI008, RS2PI008DE, RS2MB044, RS2WC080, RS2WC040, SROMBSASMR, SRCATAWB, SRCASARB, SRCASJLV, SRCASABB8I, SRCASLS4I, SRCASAPH16I, SROMBSASFC, SROMBSASMP2, SROMBSAS18E, SRCAS18E and SRCAS144E.

- The first generation SAS controllers (SRCAS18E, SRCAS144E, SROMBSAS18E) provide XOR RAID modes 0, 1, 5, 10, and 50
- The second generation SAS controller (including SRCATAWB, SRCASARB, SRCASJLV, SRCASABB8I, SRCASLS4I, SRCASAPH16I, SROMBSASFC, SROMBSASMP2, SROMBSASMR) provides XOR RAID modes 0, 1, 5, 6, 10, 50, and 60.
- The Intel® 6Gb/s SAS PCIe Gen 2 / Gen 3 RAID Controllers (including RMS25PB080, RMS25PB040, RMT3PB080, RMS25CB080, RMS25CB040, RMT3CB080, 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.)
- The Intel® RAID SSD Cache Controller RCS25ZB040/RCS25ZB040LX
For more details, refer to the *Technical Product Specification (TPS)* or *Hardware User's Guide (HWUG)* for the RAID controllers.
- 12Gb/s Intel® RAID Controller RS3WC080, RS3DC040, RS3DC080, RS3MC044, RS3SC008, RS3UC080, RS3FC044, and RS3GC008; 12Gb/s Intel® Integrated RAID Module RMS3CC080, RMS3CC040, RMS3HC080, and RMS3JC080

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 through 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.

1.2 Software

Intel[®] Embedded Server RAID Technology II, 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 II, 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 6Gb/s Intel[®] RAID controller, press <Ctrl> + <G> during the server boot to enter the RAID BIOS Console II.
 - For 12Gb/s Intel[®] RAID controller, press <Ctrl> + <R> during the server boot to enter 12Gb/s Intel[®] RAID controller configuration utility.
 - Intel[®] RAID Controller Drivers: Intel provides software drivers for the following operating systems.
 - Microsoft Windows Server 2012*, Microsoft Windows Server 2012* R2, Microsoft Windows Server 2008*, Microsoft Windows Server 2008* R2, 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, 5.0, and 6.0 (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.

1.3 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.

1.3.1 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.

1.3.1.1 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.

1.3.1.2 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.

1.3.2 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.

1.3.3 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.

1.3.3.1 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

1.3.3.2 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.*

1.3.3.3 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.

2.1 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 1E: 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.

2.2 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

2.2.1 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.

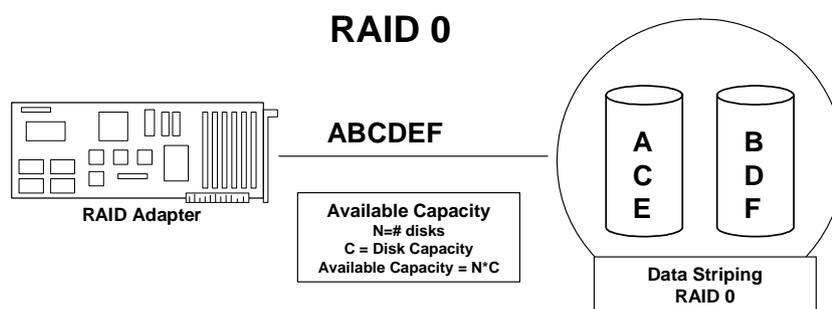


Figure 1. RAID 0 - Data Striping

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

2.2.2 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)

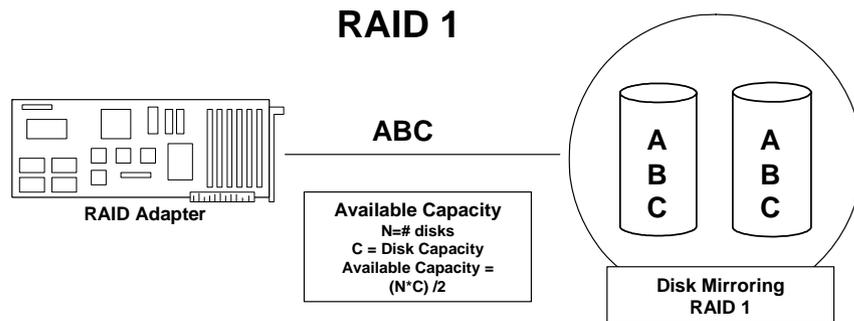


Figure 2. RAID 1 - Disk Mirroring/Disk Duplexing

2.2.3 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

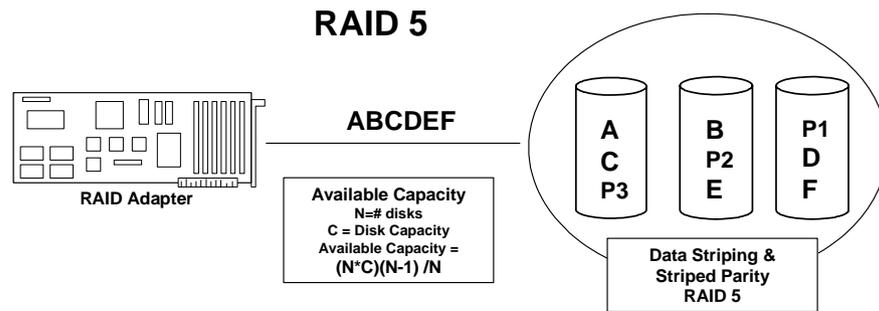


Figure 3. RAID 5 - Data Striping with Striped Parity

2.2.4 RAID 6 - Distributed Parity and Disk Striping

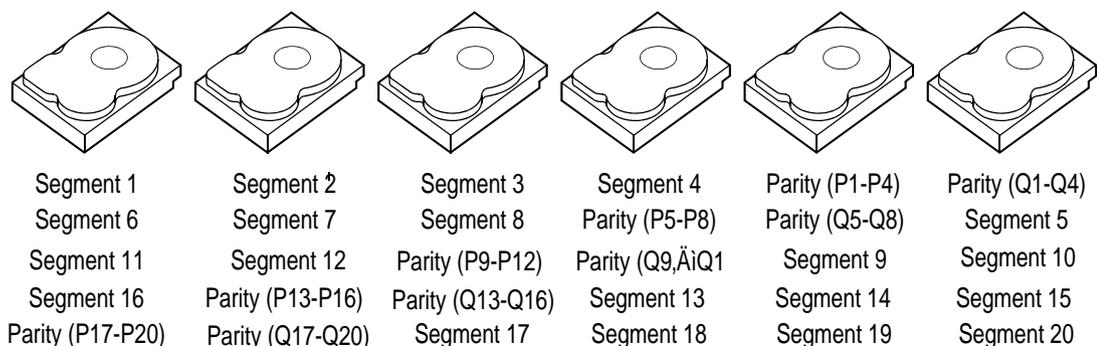
RAID 6 is similar to RAID 5 (disk striping and parity), but instead of one parity block per stripe, there are two. With two independent parity blocks, RAID 6 can survive the loss of two disks in a virtual disk without losing data.

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	<p>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.</p>
Weak Points	<p>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.</p>
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)

2.2.5 RAID 1E

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.

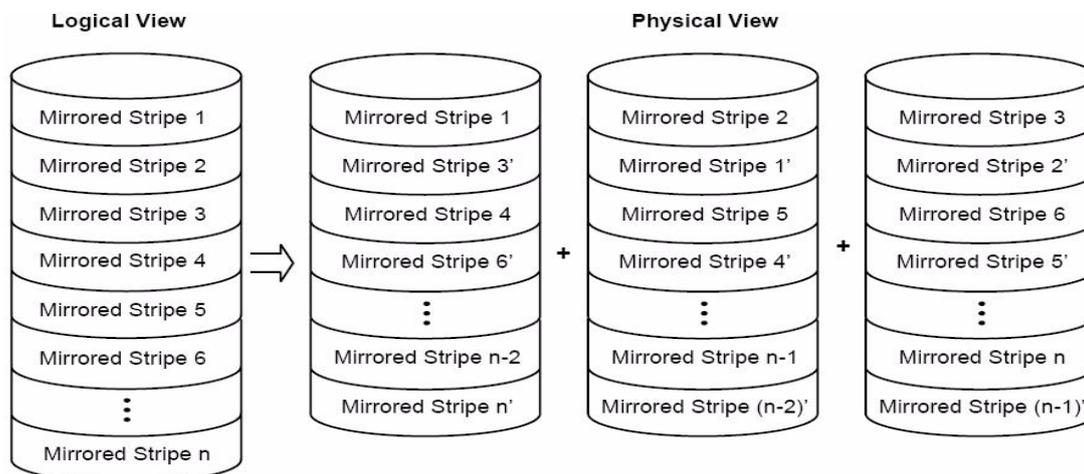


Figure 5. Integrated Mirroring Enhanced with Three Disks

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

2.2.6 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

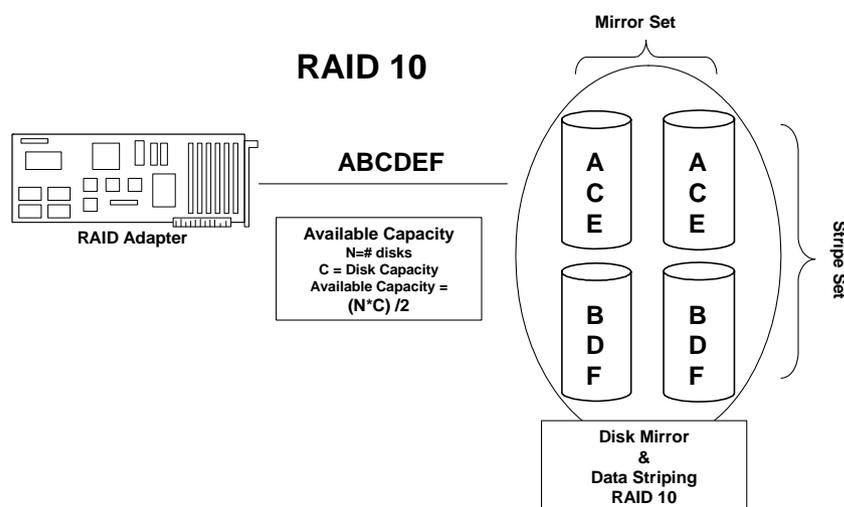


Figure 6. RAID 10 - Combination of RAID 1 and RAID 0

2.2.7 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

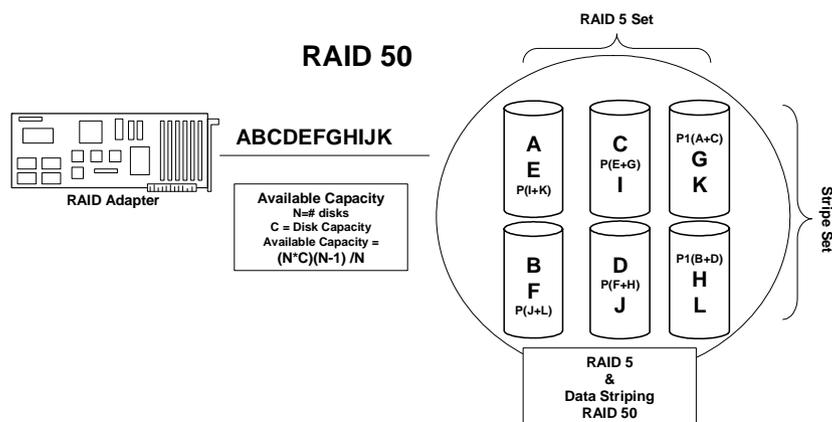


Figure 7. RAID 50 - Combination of RAID 5 and RAID 0

2.2.8 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.

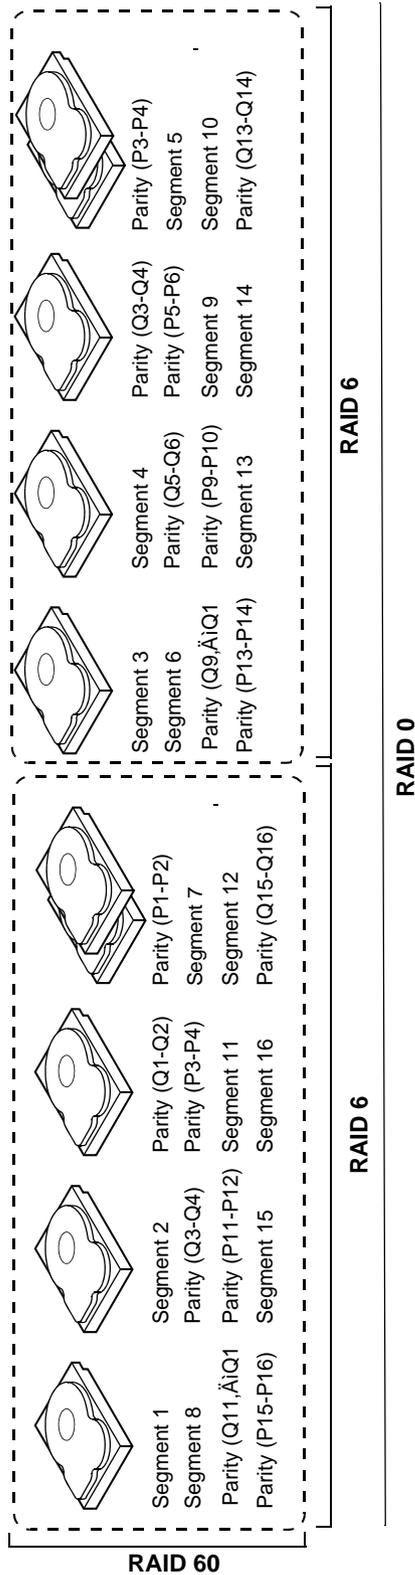
Table 8 provides an overview of RAID 60.

Table 8. RAID 60 Overview

<p>Uses</p>	<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>
<p>Strong Points</p>	<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>
<p>Weak Points</p>	<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>
<p>Drives</p>	<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

2.3 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.

2.3.1 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.

RAID Level	Fault Tolerance
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.
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.

2.3.2 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>

2.3.3 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 6 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.

2.4 RAID Availability

2.4.1 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.

2.4.2 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, 1E, 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.

2.4.3 Rebuilding

If a physical disk fails in an array that is configured as a RAID 1, 1E, 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.

2.4.4 Drive in Foreign State

When newly inserted drives are detected by the RAID controller, and are displayed in either RAID BIOS Console 2 or RAID Web Console 2, their state may show as (Foreign) Unconfigured Good, or (Foreign) Unconfigured Bad. The Foreign state indicates that the RAID controller finds existing RAID configuration on the new drives. Since these drives cannot be configured directly, avoid deleting data on the existing RAID by mistake. Use “Scan for Foreign Configuration” option in RAID Web Console 2, or use “Scan Devices” option in RAID BIOS Console 2, in order to preview, import or clear the existing RAID configurations on the drives. If existing RAID configurations are cleared, their state changes to Unconfigured Good or Unconfigured Bad.

2.4.5 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.*

2.4.5.1 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.

2.4.6 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.

2.4.7 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.

2.4.8 Serial ATA III Features

The SATA bus is a high-speed, internal bus that provides a low pin count (LPC), low voltage level bus for device connections between a host controller and a SATA device.

The following list describes the SATA III features of the RAID controllers:

- Supports SATA III data transfers of 12Gb/s
- Supports STP data transfers of 12Gb/s
- Provides a serial, point-to-point storage interface
- Simplifies cabling between devices
- Eliminates the master-slave construction used in parallel ATA
- Allows addressing of multiple SATA II targets through an expander
- Allows multiple initiators to address a single target (in a fail-over configuration) through an expander

2.4.9 UEFI 2.0 Support

UEFI 2.0 provides MegaRAID customers with expanded platform support. The MegaRAID UEFI 2.0 driver, a boot service device driver, handles block IO requests and SCSI pass-through (SPT) commands, and offers the ability to launch pre-boot MegaRAID management applications through a driver configuration protocol (DCP). The UEFI driver also supports driver diagnostic protocol, which allows administrators to access pre-boot diagnostics.

2.4.10 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.

2.4.11 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 or AXXRPFKSSD2 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) or Intel[®] RAID Premium Feature Keys AXXRPFKSSD2, AXXRPFKDE2, and AXXRPFKSNSH2 Installation Guide (G29824-00x) for description of the Premium Feature Key (PFK).

2.4.12 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.

2.4.13 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.

2.4.14 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.

2.4.15 Shield State

Physical devices in RAID firmware transit between different states. If the firmware detects a problem or a communication loss for a physical drive, the firmware transitions the drive to a bad (FAILED or UNCONF BAD) state. To avoid transient failures, an interim state called the Shield State is introduced before marking the drive as being in a bad state.

The Shield State is an interim state of a physical drive for diagnostic operations. The results of the diagnostic tests determine if the physical drive is good or bad. If any of the diagnostic tests fail, the physical drive transitions to a bad state (FAILED or UNCONF BAD).

The three possible Shield States are:

- Unconfigured - Shielded
- Configured - Shielded
- Hotspare - Shielded

Physical View and Logical View in either RAID BIOS Console 2 or RAID Web Console 2 can reflect drive Shield State. Other drive states include:

- Unconfigured Good

- Online
- Hotspare
- Failed
- Rebuilding
- Unconfigured Bad
- Missing
- Offline
- None

2.4.16 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, 1E, 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 1E, RAID 5, RAID 6, RAID 10, RAID 50, RAID 60

3 RAID Utilities

3.1 Intel® Embedded Server RAID Technology II 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 II 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.

3.2 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.
- 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.
- 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.

3.3 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

3.4 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.

3.4.1 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.

3.5 Intel[®] Intelligent RAID Controller Features

3.5.1 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 as mentioned in the next point), 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.

3.5.2 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 through 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.

3.5.3 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.

3.5.4 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.

3.5.5 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.

3.5.6 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.

4.1 RAID Driver Installation for Microsoft Windows*

4.1.1 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.

4.1.2 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.

4.2 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.

4.3 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.

4.4 RAID Driver Installation for Novell NetWare*

4.4.1 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.

4.4.2 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.

4.5 RAID Driver Installation for Solaris* 10

4.5.1 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.

4.5.2 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:

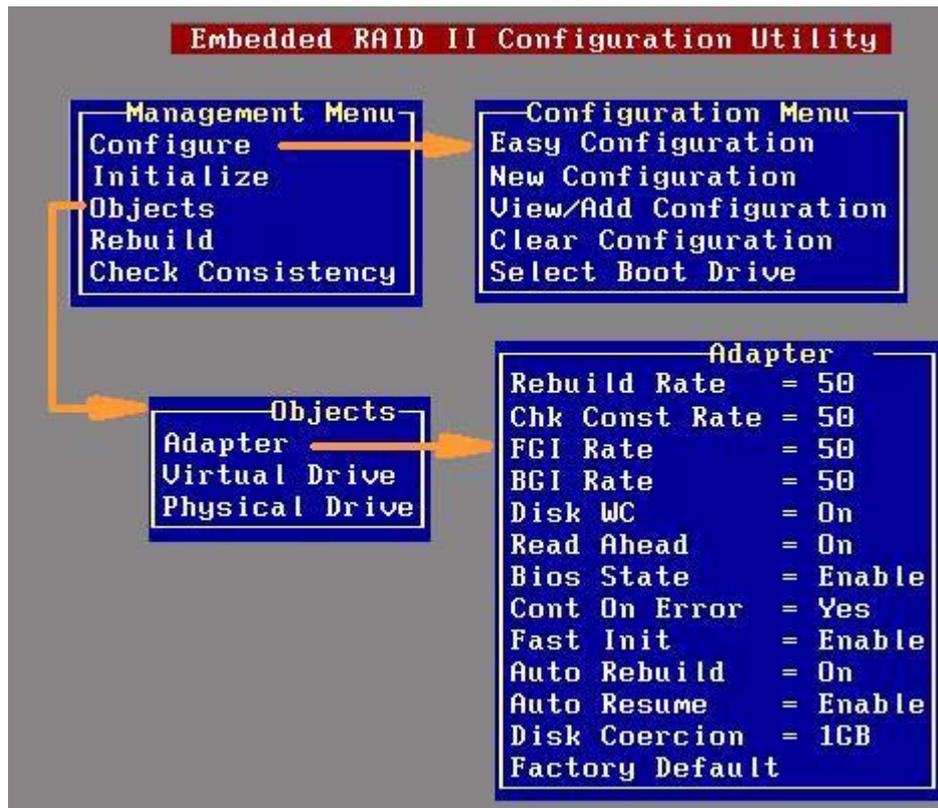


Figure 9. Intel® Embedded Server RAID BIOS Configuration Utility Screen

5.1 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
- Access Policy

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.

Note: If you create an SSD virtual drive and set the access policy to 'Read-only', it is strongly recommended that you reboot the system for the changes to take effect, or else, you will still have access to create files, and delete them.

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.

5.2 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.

5.3 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 43.

5.3.1 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.

5.3.2 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.

5.4 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.

5.4.1 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**.

5.5 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**.

5.6 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.

5.7 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.

5.8 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.

5.9 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 Controller Configuration Utility

This chapter explains how to create mirrored volumes and integrated striping volumes with the Intel SAS-3 BIOS Configuration Utility (SAS3 BIOS CU).

6.1 Overview

This section provides an overview of Intel Integrated RAID features that support the creation of mirrored volumes and striped volumes.

6.1.1 Integrated Mirroring and Integrated Mirroring Enhanced Features

Integrated Mirroring, Integrated Mirroring + Striping, and Integrated Mirroring Enhanced volumes support the following features:

- Configurations of one or two mirrored volumes on each Intel SAS-3 controller. Each volume can consist of two mirrored disks for an Integrated Mirroring volume; three to ten mirrored disks for an Integrated Mirroring Enhanced volume; or four, six, eight, or ten mirrored disks for an Integrated Mirroring + Striping volume.
- (Optional) Two global hot spare disks per Intel SAS-3 controller to automatically replace failed disks in mirrored volumes.
- Ability of mirrored volumes to run in optimal mode or in degraded mode if one mirrored disk in an Integrated Mirroring volume fails or if one or more mirrored disks fail in an Integrated Mirroring + Striping volume or Integrated Mirroring Enhanced volume.
- Support for hot swapping.
- Support for OCE for RAID 1 volumes. OCE permits you to increase the size of a RAID 1 volume by replacing the existing disk drives with higher-capacity disk drives. Data is protected during the expansion process, and the RAID 1 volume remains online.
- Presentation of a single virtual drive to the operating system for each mirrored volume.
- Support for both SAS disks and SATA disks, although you cannot combine the two types of disks in the same volume. However, an Intel SAS-3 controller can support one volume with SATA disks and a second volume with SAS disks.
- Automatic background initialization after a volume is created.
- Consistency checking.

- Fusion-MPT architecture.
- Menu-driven, BIOS-based configuration utility.
- Error notification, in which the drivers update an OS-specific event log.
- Support for SCSI Enclosure Services (SES) status LED.
- Write journaling, which permits automatic synchronization of potentially inconsistent data after unexpected powerdown situations.
- Use of metadata to store volume configuration on disks in a mirrored volume.
- Automatic background resynchronization while host I/O transactions continue.
- Background media verification, which makes sure that data on mirrored volumes is always accessible.

6.1.2 Integrated Striping Features

Integrated Striping supports the following features:

- Support for RAID volumes with two disks to ten disks
- Support for two Integrated Striping volumes with up to 14 drives total on a SAS-3 controller
- Support for combining one Integrated Striping volume and one Integrated Mirroring, Integrated Mirroring + Striping, or Integrated Mirroring Enhanced volume on a single controller
- Support for both SAS and SATA drives, although you cannot combine the two types of drives in one volume
- Fusion-MPT architecture
- Easy-to-use SAS-3 BIOS configuration utility
- Error notification
- Disk write caching, which is enabled by default on all Integrated Striping volumes
- Use of metadata to store volume configurations on disks
- OS-specific event log
- Error display inside the Fusion-MPT BIOS
- SES status LED support for drives used in Integrated Striping volumes

6.2 Creating Mirrored Volumes

This section explains how to create Integrated Mirroring volumes, Integrated Mirroring + Striping volumes, and Integrated Mirroring Enhanced volumes with the Intel SAS-3 BIOS Configuration Utility (SAS3 BIOS CU).

6.2.1 Mirrored Volume Configuration Overview

The Intel SAS3 BIOS CU is a menu-driven utility program that enables you to easily configure and manage Integrated RAID volumes. You can use the SAS3 BIOS CU to create one or two mirrored volumes on each Intel SAS-3 controller, with up to two optional global hot spare disks. You must connect all disks in a mirrored volume to the same Intel SAS-3 controller.

Although you can use disks of different sizes in mirrored volumes, the smallest disk in the volume determines the logical size of all disks in the volume. In other words, the volume does not use the excess space of the higher-capacity member disks. For example, if you create an Integrated Mirroring Enhanced volume with two 100-GB disks and two 120-GB disks, the volume uses only 100 GB on each of the 120-GB disks.

See [6.1.1, Integrated Mirroring and Integrated Mirroring Enhanced Features](#) for more information about the features of Integrated Mirroring, Integrated Mirroring + Striping, and Integrated Mirroring Enhanced volumes.

6.2.2 Creating Mirrored Volumes

The SAS3 BIOS CU is part of the Fusion-MPT BIOS. When the BIOS loads during the startup sequence and you see the message about the Intel Configuration Utility, press Ctrl-C to start the SAS3 BIOS CU. When you start the SAS3 BIOS CU, the message changes to the following:

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

After a brief pause, the main menu (Adapter List window) of the SAS3 BIOS CU appears. On some systems, however, the following message appears next:

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

In this case, the SAS3 BIOS CU loads after the system completes its power-on self-test.

You can configure one or two Integrated Mirroring, Integrated Mirroring + Striping, and Integrated Mirroring Enhanced volumes on each Intel SAS-3 controller. Alternatively, you can configure one mirrored volume and one Integrated Striping volume on the same controller, up to a maximum of 14 disk drives for the two volumes. (The maximum number includes one or two optional hot spare disks for the mirrored volume or volumes.) Additional information about configuring a RAID volume follows:

- All physical disks in a volume must be either SATA (with extended command set support) or SAS (with SMART support). You cannot combine SAS and SATA disks 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 in the volume must have 512-byte blocks and must not have removable media.
- Integrated Mirroring volumes must have two disks, Integrated Mirroring Enhanced volumes contain three disks to ten disks, and Integrated Mirroring + Striping volumes can have four, six, eight, or ten disks.

Note: Intel strongly recommends that you create global hot spare disks for all mirrored volumes to increase the level of data protection. If a disk in a mirrored volume fails, the Integrated RAID firmware rebuilds it using one of the global hot spares, and the data is safe. If you create two mirrored volumes on an Intel SAS-3 controller, either of the volumes can use the global hot spares if a disk fails.

6.2.2.1 Creating an Integrated Mirroring Volume

Follow these steps to create a two-disk Integrated Mirroring (RAID 1) volume with the SAS3 BIOS CU. The steps begin with the Adapter List window that appears when the SAS3 BIOS CU starts.

1. On the Adapter List window, use the arrow keys to select an Intel SAS-3 adapter, and then press Enter.

The Adapter Properties window appears, as the following figure shows.

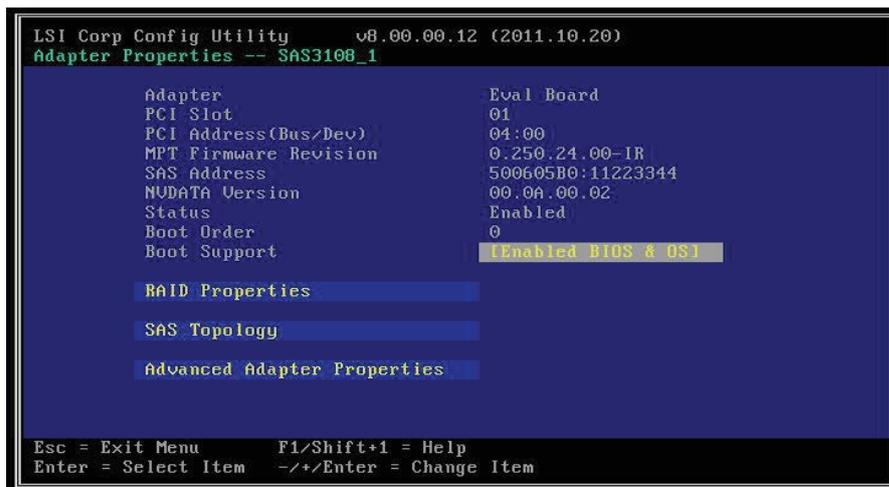


Figure 10. Adapter Properties Window

2. Use the arrow keys to select **RAID Properties**, and then press Enter.

The Select New Volume Type window appears.

3. Select **Create RAID 1 Volume**.

The Create New Volume window appears.

4. Move the cursor to the *RAID Disk* column and select a line that has a *No* entry in this column, indicating that the disk is not already part of the volume you are creating. To add the disk to the new volume, change the *No* to *Yes* by pressing the space bar.

This disk is the *Primary* disk in the volume.

Caution: The SAS3 BIOS CU deletes all existing data from the disks drives when you select them to use in a mirrored volume.

5. Move the cursor to another line and press the space bar to add the second disk to the volume.

This disk is the *Secondary* disk in the volume.

6. Press C to create the volume.

A menu window appears.

7. From the menu options, select **Save changes then exit this menu**.

A message appears briefly, and then the SAS3 BIOS CU returns to the Adapter Properties window. Initialization of the new volume continues in the background.

Note: To create a second Integrated Mirroring volume, repeat these instructions starting with step 2. Alternatively, follow the instructions in the following section to create an Integrated Mirroring Enhanced or Integrated Mirroring + Striping volume.

Note: See the instructions in Section 6.2.3, *Managing Hot Spare Disks*, if you want to create one or two global hot spares.

6.2.2.2 Creating an Integrated Mirroring Enhanced or Integrated Mirroring + Striping Volume

Integrated Mirroring Enhanced volumes contain three physical disks to ten physical disks. Data is written to a disk and is mirrored on an adjacent disk. Integrated Mirroring + Striping volumes can have a minimum of four physical disks and a maximum of 10 physical disks, in even numbers. In an Integrated Mirroring Enhanced or Integrated Mirroring + Striping volume, the data is both mirrored and striped.

Follow these steps to create an Integrated Mirroring Enhanced (RAID 1E) or Integrated Mirroring + Striping (RAID 10) volume with the SAS3 BIOS CU. The steps begin with the Adapter List window that appears when the configuration utility starts.

1. On the Adapter List window, use the arrow keys to select an Intel SAS-3 adapter, and then press Enter.

The Adapter Properties window appears, as shown in [Figure 10](#).

2. Use the arrow keys to select **RAID Properties**, and then press Enter.

The Select New Volume Type window appears.

3. Select **Create RAID 1E/10 Volume**.

The Create New Volume window appears.

4. Move the cursor to the *RAID Disk* column and select a line that has a *No* entry in this column, which indicates that the disk is not already part of the volume you are creating. To add the disk to the new volume, change the *No* to *Yes* by pressing the space bar.

Caution: The SAS3 BIOS CU deletes all existing data from the disks drives when you select them to use for a mirrored volume.

5. Move the cursor to another line and press the space bar to add another disk to the volume.

If you select an odd number of disks, the SAS3 BIOS CU creates an Integrated Mirroring Enhanced volume. If you select an even number of disks, it creates an

Integrated Mirroring + Striping volume. As you add disks, the Volume Size field changes to reflect the size of the new volume.

6. Press C to create the volume.

A menu window appears.

7. From the menu options, select **Save changes then exit this menu**.

A message appears briefly, and then the SAS3 BIOS CU returns to the Adapter Properties window. Initialization of the new volume continues in the background.

Note: To create a second Integrated Mirroring Enhanced or Integrated Mirroring + Striping volume, repeat the previous instructions.

Note: See the instructions in Section 6.2.3, [Managing Hot Spare Disks](#), if you want to create one or two global hot spares.

6.2.2.3 Expanding an Integrated Mirroring Volume with OCE

Use the OCE feature to expand the capacity of a two-disk Integrated Mirroring (RAID 1) volume by replacing the original disks with two higher-capacity disk drives while the volume remains online. This process maintains data integrity at all times, even if one of the disks fails during the replacement process. The new disks must have at least 50 GB more capacity than the disks they are replacing, and they must use the same protocol (SAS or SATA) as the disks they are replacing.

Follow these steps to expand an existing RAID 1 volume with OCE:

1. Physically replace one of the two volume disk drives with a drive that has at least 50 GB more capacity.
If necessary, you can identify the disks in the volume by following the instructions in Section 6.2.4.5, [Locating Disk Drives in a Volume](#).
2. Wait until synchronization completes on the new disk and the volume returns to the Optimal state, as indicated in the Adapter Properties window of the SAS3 BIOS CU.
3. Physically replace the other volume disk drive with a drive that has at least 50 GB more capacity.
4. Again, wait until synchronization completes on the new disk and the volume returns to the Optimal state.
5. In the Adapter List window of the SAS3 BIOS CU, use the arrow keys to select the Intel SAS adapter with the RAID 1 volume, and then press Enter.
The Adapter Properties window appears.
6. Use the arrow keys to select **RAID Properties**, and then press Enter.
The Select New Volume Type window appears.
7. Select **View Existing Volume**.
The View Volume window appears. If necessary, press Alt + N to switch to the RAID 1 volume with the new, higher-capacity disk drives.

8. Select **Manage Volume**.
The Manage Volume window appears.
9. Select **Online Capacity Expansion**.
A menu window appears with a warning message and with options to start the expansion process or quit.
10. Press Y to start the expansion.
The RAID Properties window appears when the expansion process completes.
11. Run a commercial tool specific to the operating system to move or increase the size of the partition on the newly expanded RAID 1 volume.

6.2.3 Managing Hot Spare Disks

You can create one or two global hot spare disks to protect the data on mirrored volumes on an Intel SAS-3 controller. You can also delete hot spare disks.

6.2.3.1 Creating Hot Spare Disks

Follow these steps to add global hot spare disks to an existing volume. The steps begin with the Adapter List window that appears when the configuration utility starts.

1. In the Adapter List window, use the arrow keys to select the Intel SAS-3 adapter on which you want to create hot spare disks, and then press Enter.
The Adapter Properties window appears.
2. Use the arrow keys to select **RAID Properties**, and then press Enter.
The Select New Volume Type window appears.
3. Select **View Existing Volume**.
The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.
4. Select **Manage Volume**.
The Manage Volume window appears, as shown in the following figure.



Figure 11. Manage Volume Window

5. Select **Manage Hot Spares**.

The Manage Hot Spares window appears.

6. Identify a disk that is not part of a RAID volume (that is, the value in the *Drive Status* column is not *RAID*) and that is not already identified as a hot spare disk.

A global hot spare disk must have 512-byte blocks and nonremovable media. The disk type must be either SATA with extended command set support or SAS with SMART support.

7. Select the *Hot Spr* (Hot Spare) field for this disk, and press the space bar.

The *Hot Spare* status changes to *Yes*.

8. (Optional) Repeat the preceding step to select a second global hot spare disk.

9. Press C to create the hot spare disk.

A menu window appears. An error message appears if the selected disk is not at least as large as the smallest disk used in the existing volume or volumes. An error message also appears if you try to add a SATA disk as a hot spare for volumes that use SAS disks, or if you try to add a SAS disk as a hot spare for volumes that use SATA disks.

10. Select **Save changes then exit this menu** to create the hot spare disk or disks.

The SAS3 BIOS CU pauses while it configures the global hot spares.

6.2.3.2 Deleting a Hot Spare Disk

Follow these steps to delete a global hot spare disk:

1. Access the Manage Hot Spares window by following the first five steps of the previous section.
2. Select a hot spare disk for deletion, and press C.

3. Select **Save changes then exit this menu** to complete the deletion of the hot spare disk.

The configuration utility pauses while it removes the global hot spare.

6.2.4 Other Configuration Tasks

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

6.2.4.1 Viewing Volume Properties

Follow these steps to view the RAID properties of a mirrored volume:

1. In the SAS3 BIOS CU, select an Intel SAS-3 adapter from the adapter list.

The Adapter Properties window appears.

2. Select **RAID Properties**.

The Select New Volume Type window appears.

3. Select **View Existing Volume**.

The View Volume window appears, showing information about the volume and each disk in it. The window includes global hot spare information, if any exists.

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

6.2.4.2 Running a Consistency Check

Use the Consistency Check command to verify that the data is synchronized on the mirrored disks in the volume.

Follow these steps to run a consistency check on a selected mirrored volume:

1. In the Adapter List window, use the arrow keys to select an Intel SAS adapter.

The Adapter Properties window appears.

2. Use the arrow keys to select **RAID Properties**, and then press Enter.

The Select New Volume Type window appears.

3. Select **View Existing Volume**.

The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.

4. Select **Manage Volume**.

The Manage Volume window appears.

5. Select **Consistency Check** on the Manage Volume window.

A menu window appears.

6. Press Y to start the consistency check.

The consistency check runs in the background. If it encounters any data mismatches, it stores the information in a bad block table.

6.2.4.3 Activating a Volume

A volume can become inactive if, for example, you remove it from one controller or computer and install it on a different one. The *Activate Volume* option permits you to reactivate an inactive volume.

Follow these steps to activate a selected volume:

1. In the Adapter List window, use the arrow keys to select an Intel SAS adapter and press Enter.

The Adapter Properties window appears.

2. Select **RAID Properties**, and then press Enter.

The Select New Volume Type window appears.

3. Select **View Existing Volume**.

The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.

4. Select **Manage Volume**.

The Manage Volume window appears.

5. Select **Activate Volume** on the Manage Volume window.

A menu window appears.

6. Press Y to activate the volume.

The volume becomes active after a pause.

6.2.4.4 Deleting a Volume

Caution: *Before you delete a volume, be sure to back up all data on the volume that you want to keep.*

Follow these steps to delete a selected volume:

1. In the Adapter List window, use the arrow keys to select an Intel SAS adapter.

The Adapter Properties window appears.

2. Use the arrow keys to select **RAID Properties**, and then press Enter.

The Select New Volume Type window appears.

3. Select **View Existing Volume**.

The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.

4. Select **Manage Volume**.
The Manage Volume window appears.
5. Select **Delete Volume**.
A menu window appears.
6. Either press Y to delete the volume, or press N to cancel the deletion process.
After a pause, the utility deletes the volume. If there is another remaining volume and one or two hot spare disks, the BIOS checks the hot spare disks to determine if they are compatible with the remaining volume. If they are not compatible (too small or wrong disk type), the BIOS deletes them also.

6.2.4.5 Locating Disk Drives in a Volume

You can use the SAS3 BIOS CU to locate and identify a specific physical disk drive in a disk enclosure by flashing the drive's LED. You can also flash the LEDs of all the disk drives in a RAID volume, if they are in a disk enclosure.

When you add a disk drive to a new mirrored volume, the LED on the disk drive starts flashing. The LED stops flashing when you finish creating the volume.

To locate disk drives by flashing their LEDs, follow these steps:

1. Select the desired SAS-3 controller on the Adapter List window, and press Enter.
The Adapter Properties window appears.
2. Highlight **SAS Topology**, and press Enter.
The SAS Topology window appears.
3. Select the disk in the *Device Identifier* column, and press Enter.
The LED on the disk flashes until you press a key to stop it.
4. To identify all the disk drives in a volume, select the volume in the left column of the SAS Topology window, and press Enter.
The LEDs flash on all disk drives in the volume until you press a key to stop them.

Note: The LEDs on the disk drives flash as previously described if the firmware configuration is correct and the drives are in a disk enclosure.

6.2.4.6 Selecting a Boot Disk

You can select a boot disk in the SAS Topology window. The next time you boot the computer, the firmware moves this disk to scan ID 0, making it the new boot disk. This feature makes it easier to set BIOS boot device options and to keep the boot device constant during device additions and removals. You can also select an alternative boot device. If the BIOS cannot find the preferred boot device when it loads, it attempts to boot from the alternative device.

Follow these steps to select a boot disk:

1. In the SAS3 BIOS CU, select an adapter from the adapter list.
2. Select the **SAS Topology** option.

If a device is currently designated as the boot device, the *Device Info* column on the SAS Topology window lists the word *Boot*, as shown in the following figure.

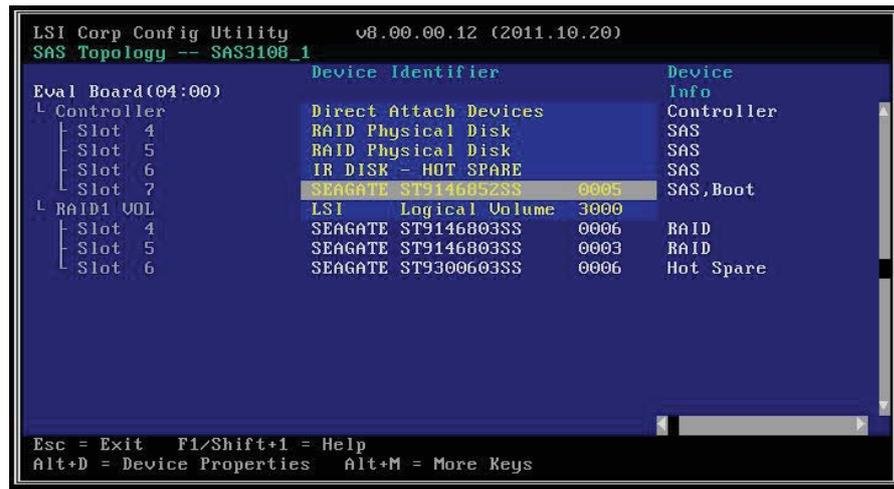


Figure 12. Boot Device on SAS Topology Window

If a device is currently designated as the alternative boot device, the *Device Info* column shows the word *Alt*.

3. To select the preferred boot disk, move the cursor to the disk, and press Alt + B.
4. To remove the boot designator, move the cursor to the current boot disk, and press Alt + B.

This controller no longer has a disk designated as boot.

5. To change the boot disk, move the cursor to the new boot disk, and press Alt + B. The Boot designator moves to this disk.
6. To select an alternative boot disk, move the cursor to the disk, and press Alt + A.

Note: To change the alternative boot device from one disk to another, follow step 4 and step 5 in this procedure, but press Alt + A instead of Alt + B.

6.3 Creating Integrated Striping Volumes

This section explains how to create Integrated Striping volumes using the Intel SAS-3 BIOS Configuration Utility (SAS3 BIOS CU).

6.3.1 Integrated Striping Configuration Overview

The Intel SAS3 BIOS CU is a menu-driven utility program that enables you to easily configure and manage Integrated RAID volumes. Use the SAS3 BIOS CU to create one or two Integrated Striping volumes on each Intel SAS-3 controller. Each volume contains two drives to ten drives. All disks in an Integrated Striping volume must be connected to the same Intel SAS-3 controller.

Although you can use disks of different sizes in Integrated Striping volumes, the smallest disk in the volume determines the logical size of all disks in the volume. In other words, the firmware does not use the excess space of the higher-capacity member disks. For example, if you create an Integrated Striping volume with two 100-GB disks and two 120-GB disks, the firmware uses only 100 GB on each of the 120-GB disks for the volume. The supported stripe size is 64 KB.

See [6.1.2, Integrated Striping Features](#) for more information about Integrated Striping volumes.

6.3.2 Creating Integrated Striping Volumes

The SAS3 BIOS CU is part of the Fusion-MPT BIOS. When the BIOS loads during boot and you see the message about the Intel Configuration Utility, press Ctrl-C to start the SAS3 BIOS CU. After you start the SAS3 BIOS CU, the message changes to the following:

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

After a brief pause, the main menu of the SAS3 BIOS CU appears. On some systems, however, the following message appears next:

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

In this case, the SAS3 BIOS CU loads after the system completes its power-on self-test.

Each Intel controller can support one or two Integrated RAID volumes. The volumes can include two Integrated Striping (RAID 0) volumes, two mirrored volumes, or one volume of each type. The two volumes can have a maximum of 14 disk drives. (This configuration includes one or two hot spare disks for mirrored volumes.)

The following guidelines apply when creating an Integrated Striping volume:

- All physical disks in the volume must be either SATA (with extended command set support) or SAS (with SMART support). You cannot combine SAS and SATA disks in the same volume. However, it is possible to configure one volume with SAS disks and one volume with SATA disks on the same controller.
- Disks in the volume must have 512-byte blocks and must not have removable media.
- Integrated Striping volumes must have at least 2 disks and no more than 10 disks. Integrated Striping volumes do not support hot spare disks.

Follow these steps to configure an Integrated Striping volume with the SAS3 BIOS CU. The steps begin with the Adapter List window that appears when the SAS3 BIOS CU starts.

1. On the Adapter List window, select an Intel SAS-3 adapter, and press Enter. The Adapter Properties window appears, as the following figure shows.



Figure 13. Adapter Properties Window

2. Select **RAID Properties**, and press Enter. The Select New Volume Type window appears.
3. Select **Create RAID 0 Volume**. The Create New Volume window appears.
4. Move the cursor to the *RAID Disk* column, and select a line that has a *No* entry in this column, which indicates that the disk is not already part of the volume you are creating. To add the disk to the new volume, change the *No* to *Yes* by pressing the space bar.
5. Move the cursor to another line and press the space bar to add another disk to the volume.
6. Continue adding disks in this way until you reach the desired number of disks.
7. Press C to create the volume. A menu appears.
8. From the menu options, select **Save changes then exit this menu**. A message appears briefly, and then the SAS3 BIOS CU returns to the Adapter Properties window. Initialization of the new volume continues in the background.

Note: Repeat the previous instructions to create a second Integrated Striping volume, if desired, and if enough additional disks are available.

6.3.3 Other Configuration Tasks

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

6.3.3.1 Viewing Volume Properties

Follow these steps to view the RAID properties of a volume:

1. In the SAS3 BIOS CU, select an Intel SAS-3 adapter from the adapter list.
The Adapter Properties window appears.
2. Select **RAID Properties**.
The Select New Volume Type window appears.
3. Select **View Existing Volume**.
The View Volume window appears, showing information about the volume and each disk in it.
4. If the currently displayed volume is not the one you want, press Alt + N to view another volume on the adapter.

6.3.3.2 Activating a Volume

A volume can become inactive if, for example, you remove it from one controller or computer and install it on a different one. The *Activate Volume* option permits you to reactivate an inactive volume.

Follow these steps to activate a selected volume:

1. In the Adapter List window, use the arrow keys to select an Intel SAS adapter and press Enter.
The Adapter Properties window appears.
2. Select **RAID Properties**, and then press Enter.
The Select New Volume Type window appears.
3. Select **View Existing Volume**.
The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.
4. Select **Manage Volume**.
The Manage Volume window appears.
5. Select **Activate Volume** on the Manage Volume window.
A menu window appears.
6. Press Y to activate the volume.
The volume becomes active after a pause.

6.3.3.3 Deleting a Volume

Caution: Before you delete a volume, be sure to back up the data.

Follow these steps to delete a selected volume:

1. In the Adapter List window, use the arrow keys to select an Intel SAS adapter.
The Adapter Properties window appears.
2. Use the arrow keys to select **RAID Properties**, and then press Enter.
The Select New Volume Type window appears.
3. Select **View Existing Volume**.
The View Volume window appears. If necessary, press Alt + N to switch to another volume on this adapter.
4. Select **Manage Volume**.
The Manage Volume window appears.
5. Select **Delete Volume**.
A menu window appears.
6. Either press Y to delete the volume, or press N to cancel the deletion process.
After a pause, the utility deletes the volume.

6.3.3.4 Locating Disk Drives in a Volume

Use the SAS3 BIOS CU to locate and identify a specific physical disk drive in a disk enclosure by flashing the drive's LED. Alternatively, use the SAS3 BIOS CU to flash the LEDs of all the disk drives in a RAID volume if they are in a disk enclosure.

When you add a disk drive to a new mirrored volume, the LED on the disk drive starts flashing. The LED stops flashing when you finish creating the volume.

To locate disk drives by flashing their LEDs, follow these steps:

1. Select the desired SAS-3 controller on the Adapter List window and press Enter.
The Adapter Properties window appears.
2. Highlight *SAS Topology*, and press Enter.
The SAS Topology window appears.
3. Select the disk in the *Device Identifier* column, and press Enter.
The LED on the disk flashes until you press a key to stop it.
4. To identify all the disk drives in a volume, select the volume in the left column of the SAS Topology window, and press Enter.
The LEDs flash on all disk drives in the volume until you press a key to stop them.

Note: The LEDs on the disk drives flash, as previously described, if the firmware configuration is correct and the drives are in a disk enclosure.

6.3.3.5 Selecting a Boot Disk

You select a boot disk in the SAS Topology window. The next time you boot the computer, the firmware moves this disk to scan ID 0, making it the new boot disk. This feature makes it easier to set BIOS boot device options and to keep the boot device constant during device additions and removals. Optionally, you can select an alternative boot device. If the BIOS cannot find the preferred boot device when it loads, it attempts to boot from the alternative device.

Follow these steps to select a boot disk:

1. In the SAS3 BIOS CU, select an adapter from the Adapter List.
2. Select the **SAS Topology** option. If a device is currently designated as the boot device, the *Device Info* column on the SAS Topology window lists the word *Boot*, as the following figure shows.



Figure 14. Boot Device on SAS Topology Window

If a device is currently designated as the alternative boot device, the *Device Info* column shows the word *Alt*.

3. To select the preferred boot disk, move the cursor to the disk, and press Alt + B.
4. To remove the boot designator, move the cursor to the current boot disk, and press Alt + B.

This controller no longer has a disk designated as boot.

5. To change the boot disk, move the cursor to the new boot disk, and press Alt + B. The *Boot* designator moves to this disk.
6. To select an alternative boot disk, move the cursor to the disk and press Alt + A.

Note: To change the alternative boot device from one disk to another, follow step 4 and step 5 in this procedure, but press Alt + A instead of Alt + B.

7 12Gb/s Intel[®] RAID Controller Configuration Utility

This chapter describes the Ctrl-R Utility, a BIOS configuration utility, that lets you create and manage RAID configurations on Intel SAS controllers. You can configure the drive groups and drives on the system before the operating system has been installed.

7.1 Overview

The Ctrl-R Utility resides in the SAS controller BIOS and operates independently of the operating system.

You can use the Ctrl-R Utility to perform tasks such as these:

- Create drive groups and virtual drives for storage configurations
- View controller, physical drive, virtual drive, enclosure, and battery backup unit (BBU) properties, and change parameters
- Delete virtual drives
- Modify power settings
- Import and clear foreign configurations
- Initialize virtual drives
- Check configurations for data consistency
- Create CacheCade virtual drives

7.2 Starting the Ctrl-R Utility

When you boot the system, perform the following steps to start the Ctrl-R Utility:

1. When the host computer is booting, press and hold the Ctrl key, and press the R key when the following text appears on the dialog:

```
Copyright© Intel Corporation
```

```
Press <Ctrl><R> for Ctrl-R
```

2. Based on the controllers on the system, one of the two following scenarios occurs:
 - If the system has multiple SAS controllers, a controller selection dialog appears. Select a controller and press Enter. The Ctrl-R Utility main menu screen appears.

- If the system has only one SAS controller, the Ctrl-R Utility main menu screen appears.

7.3 Exiting the Ctrl-R Utility

To exit the Ctrl-R Utility, perform these steps:

1. Perform one of these actions:
 - If you are not in a dialog, press Esc once.
 - If you are in a dialog, press Esc twice (once to exit the dialog, and the second time to exit the utility).

A confirmation message box appears.
2. Press **OK** to exit the utility.

7.4 Ctrl-R Utility Keystrokes

The following table lists the keystrokes that you can use in the Ctrl-R Utility to navigate between the screens.

Table 13. Ctrl-R Utility Keystrokes

Keystroke	Action
F1	Displays help for the particular screen that you are in.
F2	Displays a list of commands that can be performed for the selected device. This key stroke is available only in the VD Mgmt, the PD Mgmt, and the Foreign View menus. The commands that are enabled are highlighted in white and the disabled commands are highlighted in black. Note: <i>Based on the configurations that you make, commands are enabled or disabled.</i>
F5	Refreshes the screen that you currently are in.
F11	Switches between controllers.
F12	Displays a list of all the available controllers. You can also scroll to the next controller.
<Ctrl><N>	Displays the next menu screen.
<Ctrl><P>	Displays the previous menu screen.
<Ctrl><S>	Shortcut key for the Apply button in the Controller Settings screens.
<Tab>	Moves the cursor to the next control.
<Shift><Tab>	Moves the cursor to the previous control on a screen or a dialog.

Keystroke	Action
<Enter>	Lets you to select a menu item, a button, a check box, and values in a list box.
<Esc>	Closes a screen or a window. Press Esc twice to exit from the Ctrl-R Utility.
Up Arrow	Moves the cursor to the next menu selection.
Down Arrow	Moves the cursor to the lower menu items or to a lower level menu.
Right Arrow	Opens a submenu, moves from a menu heading to the first submenu, or moves to the first item in a submenu. The right arrow also closes a menu list in a popup window.
Left Arrow	Closes a submenu, moves from a menu item to the menu heading, or moves from a submenu to a higher level menu.
Spacebar	Lets you select a menu item, a button, and a check box.

7.5 Ctrl-R Utility Menus

The Ctrl-R Utility contains the following menus:

- **VD Mgmt**
- **PD Mgmt**
- **Ctrl Mgmt**
- **Properties**
- **Foreign View**

7.5.1 VD Mgmt Menu

The **VD Mgmt** menu is the first menu screen that appears when you start the Ctrl-R Utility.

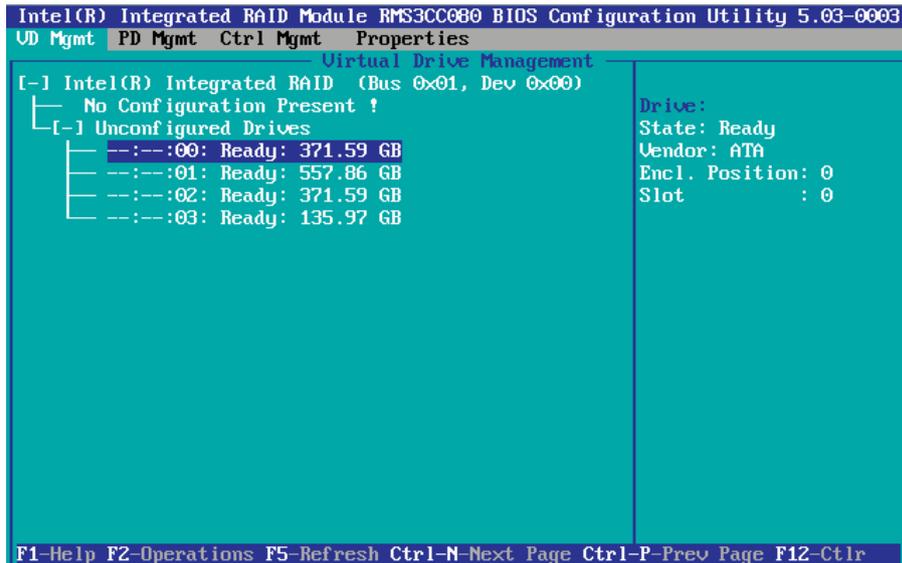


Figure 15. VD Mgmt

This screen shows information on the configuration of controllers, drive groups, and virtual drives.

The right panel of the screen shows attributes of the selected device.

In the **Virtual Drive Management** screen, you can perform tasks, such as creating and initializing virtual drives; performing a consistency check; deleting, expanding, and erasing virtual drives; importing or clearing foreign configurations; and creating CacheCade virtual drives.

Note: Based on the controller settings that you make, options will be enabled or disabled.

7.5.2 PD Mgmt Menu

The **PD Mgmt** menu shows information about all the physical drives connected to the selected controller. This menu also shows information about enclosures, the number of physical drives in an enclosure, and all of the direct-attached drives under a backplane node.

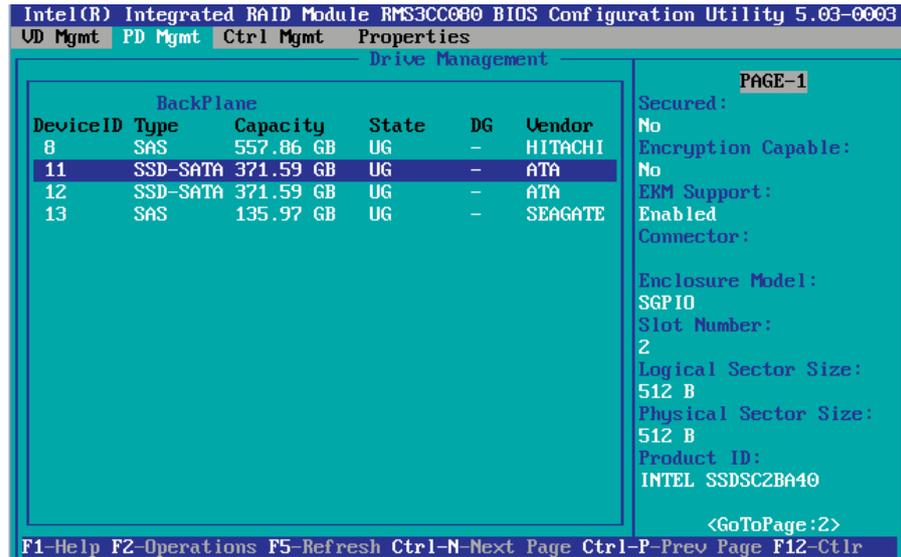


Figure 16. Physical Drive Management

The right panel of the screen shows additional attributes of the selected device.

In the **Drive Management** screen, you can perform tasks, such as rebuilding a failed drive, making a drive offline, or making it a global hot spare drive.

7.5.3 Ctrl Mgmt Menu

The **Ctrl Mgmt** menu lets you change the settings of the selected controller. The **Ctrl Mgmt** menu consists of two screens.

In the first screen (as shown in the following figure), you can change controller options, such as **Enable Controller BIOS**, **Maintain PD Fail History**, and **Enable Stop CC on Error**. You also can perform tasks, such as enabling or silencing an alarm and entering values for Rebuild Rate and Patrol Rate.

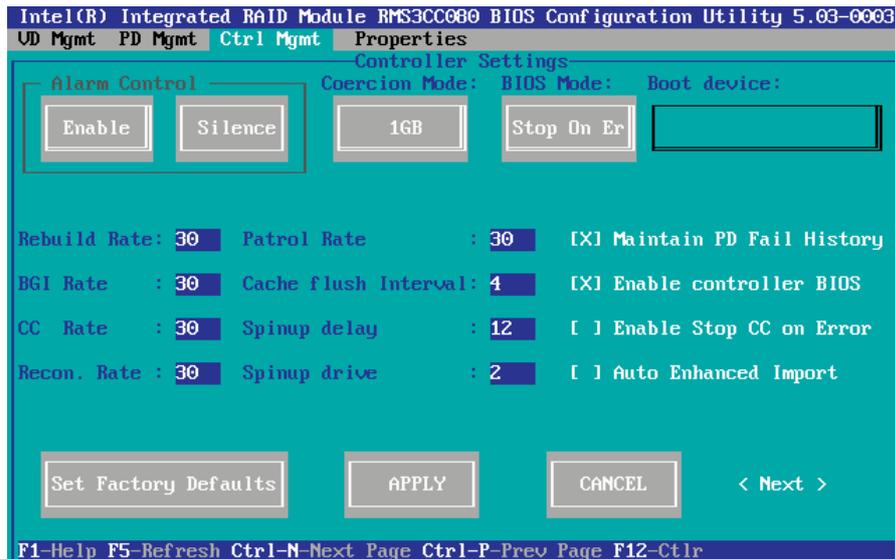


Figure 17. Controller Settings – First Screen

In the second screen (as shown in the following figure), you can perform tasks, such as changing the link speed, the power save, and the battery settings of the controller.

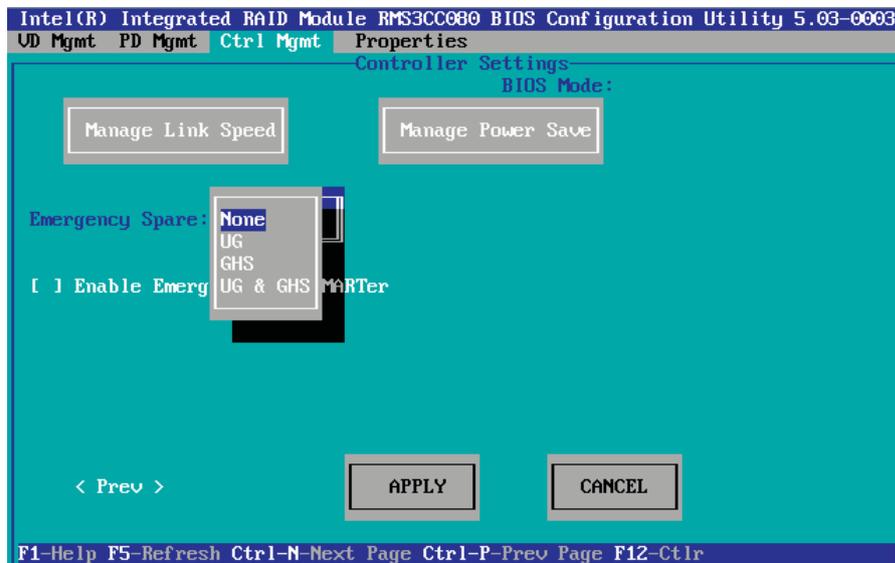


Figure 18. Controller Settings – Second Screen

7.5.4 Properties Menu

The **Properties** menu shows all of the properties of the active controller. The **Properties** menu consists of two screens. The information shown in these screens is read only.

In the first screen (as shown in the following figure), you can view properties, such as controller status, firmware version, BIOS version, and metadata size.

```
Intel(R) Integrated RAID Module RMS3CC000 BIOS Configuration Utility 5.03-0003
UD Mgmt PD Mgmt Ctrl Mgmt Properties
Properties
Product Name      : Intel(R) Integrated RAID Module RMS3CC000
Controller Status : Optimal
Serial No        : SR349P0136
ROC Temperature  : 72 Celsius
Package         : 24.2.0-0012
FW Version       : 4.220.00-2968
BIOS Version     : 6.14.05.0_4.16.00.00_0x06030500
Boot Block Version : 3.02.00.00-0000
Battery status   : Missing
Security Capable : No
Controller ID    : 0
PCI Bus         : 0x01
PCI Device      : 0x00
PCI Function    : 0x00
PCI Slot ID     : 0x05
Metadata Size   : 512 MB
Data Protection Support : Yes
Data Protection Enabled : Yes
Emergency Spare : Unconfigured Good & Global Hot Spare

< Prev > [ Page:01 ] < Next >
F1-Help F5-Refresh Ctrl-N-Next Page Ctrl-P-Prev Page F12-Ctrl
```

Figure 19. Properties

To view additional properties, you can navigate to **Next** and press Enter. The second **Properties** screen shows information, such as maximum cache size, drive standby time, and power saving properties.

To go back to the previous **Properties** screen, navigate to **Prev**, and press Enter.

7.5.5 Foreign View Menu

If one or more physical drives in a configuration are removed and reinserted, the controller considers the drives as foreign configurations.

The **Foreign View** menu is shown only when the controller detects a foreign configuration. If no foreign configuration exists, the **Foreign View** menu is not shown.

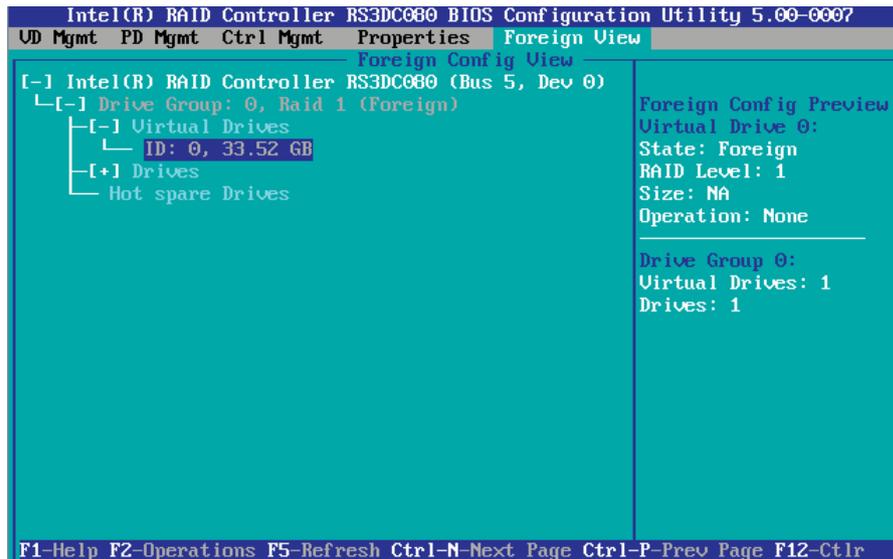


Figure 20. Foreign View Menu

You can use the **Foreign Config View** screen to view information about the foreign configuration, such as drive groups, virtual drives, physical drives, and hot spares.

The **Foreign View** menu lets you import foreign configurations to the RAID controller or clear the foreign configurations.

7.6 Controller Advanced Software Options

7.6.1 Advanced Software Options

Perform the following steps to see the advanced software options that the controller supports.

1. In the **VD Mgmt** screen, navigate to the controller and press the F2 key.
2. Navigate to **Advanced Software Options**, and press Enter.

The **Manage MegaRAID Advanced Software Options** dialog appears, as shown in the following figure.

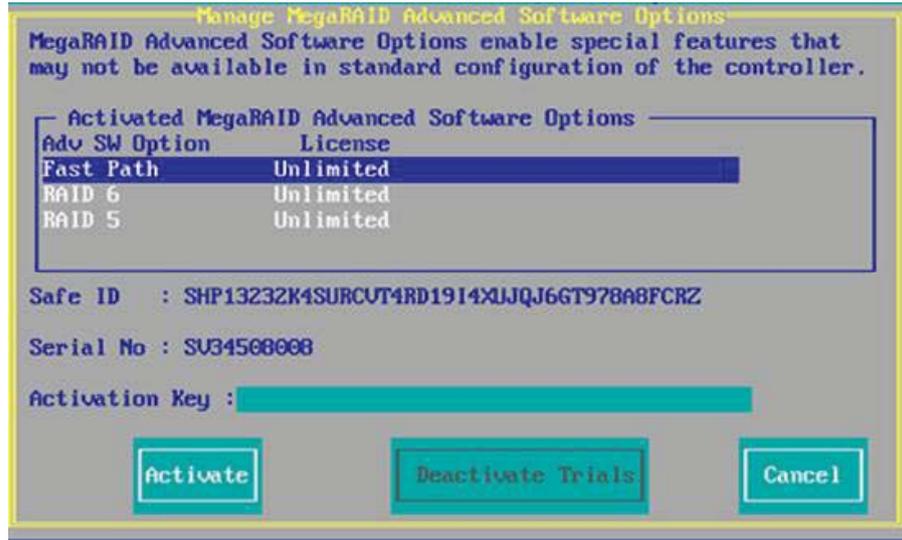


Figure 21. Advanced Software Options

The **Adv SW Option** column shows the list of Advanced Software Options available in the controller.

You can ignore **License** column, **Safe ID** and **Serial No** and **Activation Key** fields and **Activate** button. Intel RAID Controllers do not need software licenses to enable advanced software options.

To enable some premium features present in the RAID controller, you need to install Intel RAID Premium Feature Keys on the controller. For example, if you install the premium feature key AXXRPFKSSD2 on the RAID controller, you can see a premium option “CacheCade Pro 2.0” shown in the **Adv SW Option** column.

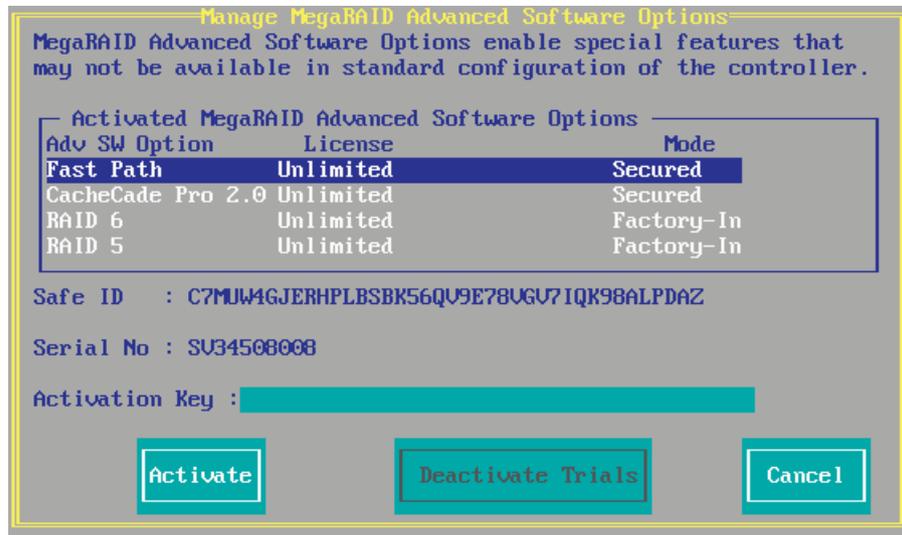


Figure 22. Advanced Software Options With Key AXXRPFKSSD2 Installed

7.6.2 RAID Premium Feature Key

Below is a table listing all premium feature keys that are available now.

For the latest support information, refer to the *Tested Hardware and Operating System List (THOL)* of each RAID product available at <http://www.intel.com/support>.

For more details of premium feature key, refer to the *Intel RAID Controller Configuration Guide* at <http://www.intel.com/support>.

Table 14. RAID Premium Feature Key

Image	Product Code MM UPC	Description
	AXXRPFKSNSH 907051 00735858214377	Intel® RAID Rapid Recovery Snapshot
	AXXRPFKDE 907050 00735858214360	Intel® RAID Drive Encryption Management
	AXXRPFKSSD 918399 00735858214353	Intel® RAID SSD Cache with Fast Path *I/O
	AXXRPFKSNSH2 915318 00735858221481	Intel® RAID Rapid Recovery Snapshot
	AXXRPFKDE2 915317 00735858221474	Intel® RAID Drive Encryption Management
	AXXRPFKSSD2 919499 00735858221467	Intel® RAID SSD Cache with Fast Path *I/O

7.7 Creating a Storage Configuration

You can use the Ctrl-R Utility to configure RAID drive groups and virtual drives to create storage configurations on systems with Intel SAS controllers.

1. In the **VD Mgmt** screen, navigate to the controller and press the F2 key.
2. Press Enter.

The **Create New VD** screen appears.

Note: You can use the **Create New VD** dialog to create virtual drives for **Unconfigured Good** drives. To create virtual drives for existing drive groups, navigate to a drive group and press the F2 key to view the **Add VD in Drive Group** dialog. The fields in the **Add VD in Drive Group** dialog are the same as in the **Create New VD** dialog.

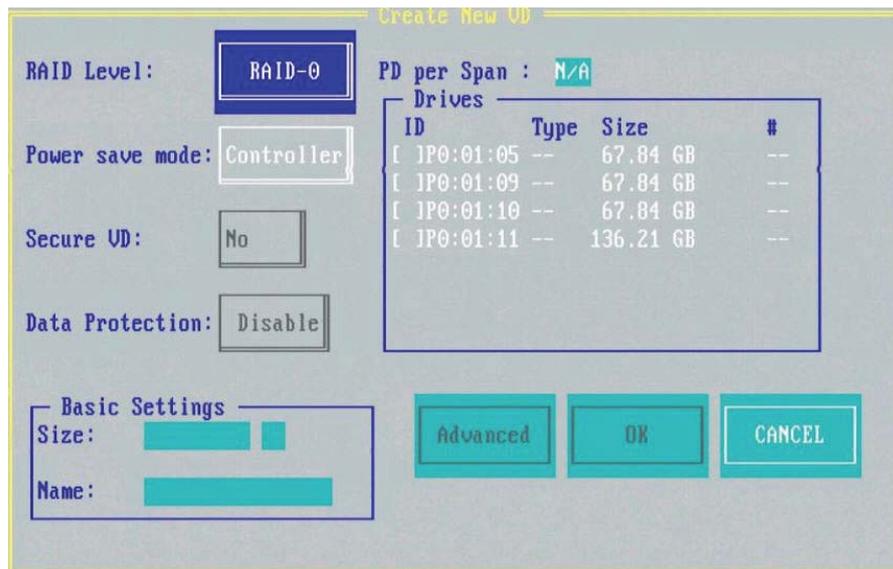


Figure 23. Create a New Virtual Drive

Note: If your system detects any JBODs the **Convert JBOD to Unconfigured Good** dialog (Figure 41) appears before the **Create New VD** dialog. The **Convert JBOD to Unconfigured Good** dialog lets you convert the JBODs to **Unconfigured Good**.

3. Select a RAID level for the drive group from the **RAID Level** field.
4. Select a power save mode for the drive group from the **Power save mode** field. The options available are **Auto**, **Max**, and **Controller defined**.

This field is enabled only if power saving on configured drives is supported on the controller.

Power Save (Dimmer Switch feature) is a technology that conserves energy by placing certain unused drives into a Power Save mode. In Power-Save mode, the

drives use less energy. The fan and the enclosure require less energy to cool and house the drives, respectively. Also, this technology helps avoid application time-outs caused by spin-up delays and drive wear caused by excessive spin-up/down cycles.

5. You can encrypt data and use drive-based key management for your data security solution. This option protects the data in the event of theft or loss of drives. Select a value from the **Secure VD** field. The options available are **Yes** and **No**.
6. You can choose whether you want to use the data protection feature on the newly created virtual drive. Select a value from the **Data Protection** field. The options available are **Yes** and **No**. The **Data Protection** field is enabled only if the controller has data protection physical drives connected to it.
7. You can change the sequence of the physical drives in the **Drives** box. All the available unconfigured good drives appear in the **Drives** box. Select the physical drives in the sequence that you prefer. Based on your selection, the sequence number appears in the **#** column.
8. You can enter a size less than the maximum size of the drive group, if you want to create other virtual drives on the same drive group. The maximum size of the drive group appears in the **Size** field. The size entered can be in MB, GB, or TB and should be mentioned only in uppercase. Before entering a size, ensure that you have deleted the previous default value by using the Backspace key.
9. Enter a name for the virtual drive in the **Name** field. The name given to the virtual drive cannot exceed 15 characters.

You may press **Advanced** to set additional properties for the newly created virtual drive. For more information, see [Section “Selecting Additional Virtual Drive Properties” on page 80](#).

10. Press **OK**.

A dialog appears, asking you whether you want to initialize the virtual drive you just created.

11. To initialize the virtual drive, press **OK**.

The **Create New VD** dialog appears again.

12. Press **Advanced**.

The **Create Virtual Drive – Advanced** dialog appears.

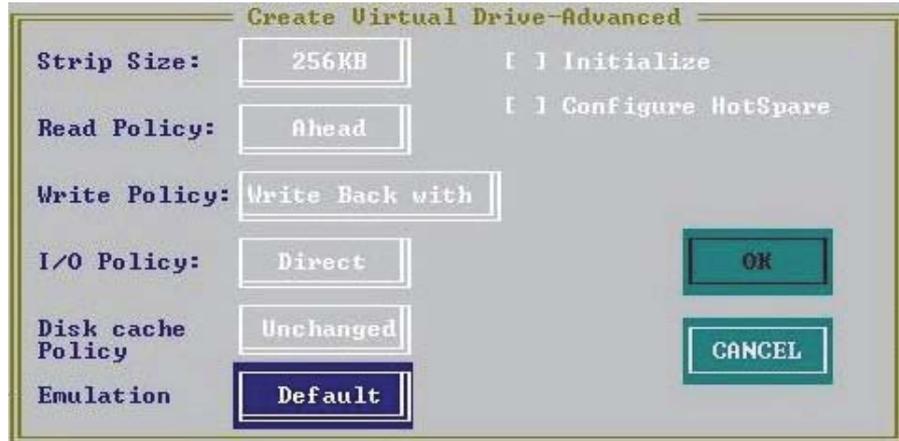


Figure 24. Create Virtual Drive – Advanced

13. Select **Initialize**, and press **OK**.

The new virtual drive is created and initialized.

7.7.1 Selecting Additional Virtual Drive Properties

This section describes the following additional virtual drive properties that you can select while you create virtual drives. Change these parameters only if you have a specific reason for doing so. It is usually best to keep them at their default settings.

- **Strip Size** – The strip size is the portion of the stripe that resides on a single virtual drive in the drive group. Strip sizes of 64 KB, 128 KB, 256 KB, 512 KB, or 1 MB are supported.
- **Read Policy** – Specify one of the following options to specify the read policy for this virtual drive:
 - **Normal** – Read ahead capability lets the controller read sequentially ahead of requested data and store the additional data in cache memory, thereby anticipating that the data will be needed soon. This process speeds up reads for sequential data, but there is little improvement when the computer accesses random data.
 - **Ahead** – Disables the read ahead capability.
- **Write Policy** – Select one of the following options to specify the write policy for this virtual drive:
 - **Write Thru** – In this mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all the data in a transaction. This option eliminates the risk of losing cached data in case of a power failure.
 - **Write Back** – In this mode, the controller sends a data transfer completion signal to the host when the controller cache has received all the data in a transaction.

- **Write Back with BBU** – In this mode, the controller enables write back caching when the BBU is installed and charged. This option provides a good balance between data protection and performance.

Note: The write policy depends on the status of the BBU. If the BBU is not present, is low, is failed, or is being charged, the current write policy switches to write through, which provides better data protection.

- **I/O Policy** – The I/O policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.
 - **Cached** – In this mode, all reads are buffered in cache memory. **Cached I/O** provides faster processing.
 - **Direct** – In this mode, reads are not buffered in cache memory. Data is transferred to the cache and the host concurrently. If the same data block is read again, it comes from cache memory. **Direct I/O** makes sure that the cache and the host contain the same data.
- **Disk cache policy** – Select a cache setting for this virtual drive:
 - **Enable** – Enable the drive cache.
 - **Disable** – Disable the drive cache.
 - **Unchanged** – Updating the drive cache policy to **Unchanged** may enable /disable the drive cache based on the WCE (Write Cache Policy) bit of the save mode page of the drive.
- **Emulation** – Lets you to set the emulation type on a virtual drive to default or none. The force option forces the emulation to be set on a controller even when MFC settings do not support it. The possible options are **Default**, **Disabled**, or **Forced**.
- **Initialize** – Select to initialize the virtual drive. Initialization prepares the storage medium for use. Fast initialization will be performed on the virtual drive.
- **Configure Hot Spare** – Select to configure physical drives as hot spares for the newly created virtual drive.

This option is enabled only if there are additional drives and if they are eligible to be configured as hot spares. This option is not applicable for RAID 0. If you select this option and after the Virtual drive is created, a dialog appears. The dialog asks you to choose the physical drives that you want to configure as hot spares.

7.7.2 Creating a CacheCade Virtual Drive

The MegaRAID CacheCade software provides you with read caching capability.

Perform the following steps to create a CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Create CacheCade Virtual Drive**, and press Enter.

The **Create CacheCade Virtual Drive** dialog appears.



Figure 25. Create CacheCade Virtual Drive

3. Enter a name for the CacheCade virtual drive in the **Name** field.
4. Select a SSD from the **Select SSD** box. The size of the SSD is reflected in the **Size** field (in the **Basic Settings** box).
5. Press **OK**.

A message appears, stating that the CacheCade virtual drive has been created.

7.7.3 Modifying a CacheCade Virtual Drive

You can modify an existing CacheCade virtual drive by renaming it.

Perform the following steps to modify the CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the CacheCade virtual drive, and press the F2 key.
2. Navigate to **Properties**, and press Enter.

The following dialog appears.

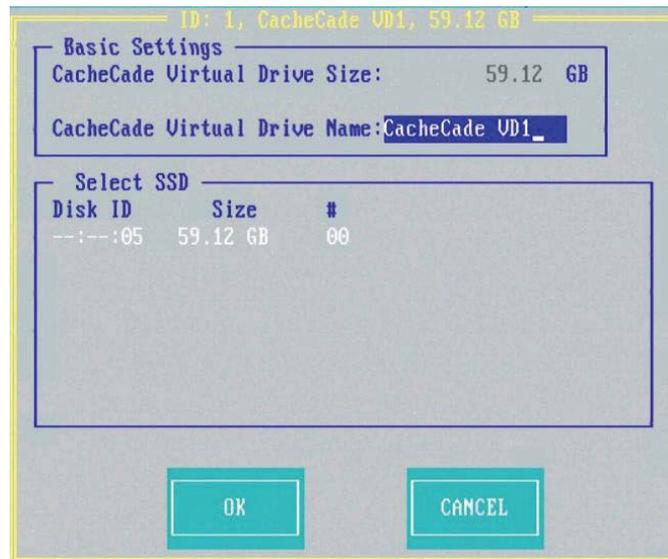


Figure 26. Modify CacheCade Virtual Drive

3. You can rename a CacheCade virtual drive in the **CacheCade Virtual Drive Name** field.
4. Press **OK**.

7.7.4 Creating a CacheCade Pro 2.0 Virtual Drive

The MegaRAID CacheCade Pro 2.0 provides you with read and write capability.

Perform the following steps to create a CacheCade Pro 2.0 virtual drive:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Create CacheCade Virtual Drive**, and press Enter.

The **Create CacheCade Virtual Drive** dialog appears.

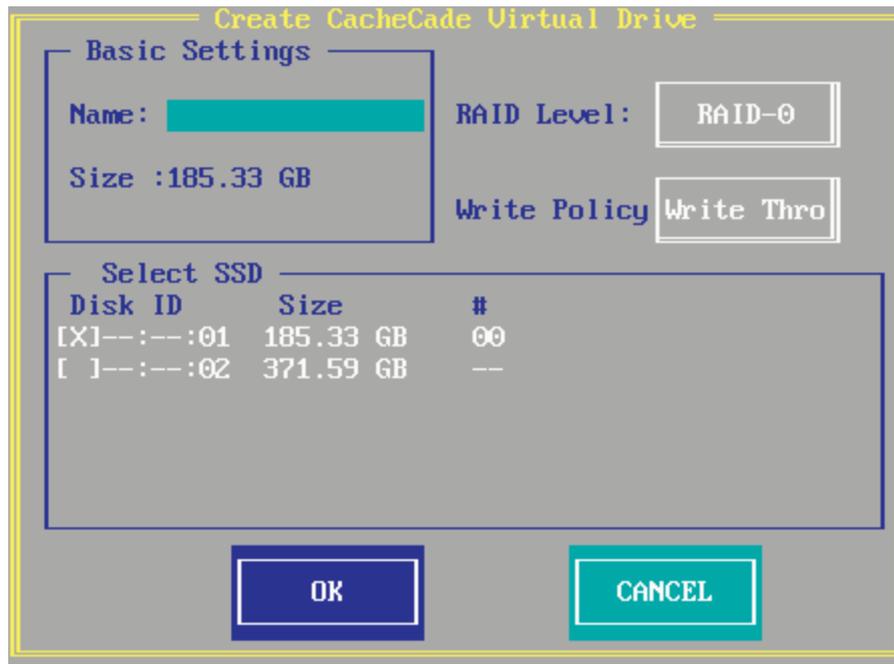


Figure 27. Create CacheCade Virtual Drive

3. Enter a name for the CacheCade virtual drive in the **Name** field.
4. Select a SSD from the **Select SSD** box.
5. Press **OK**.

A message appears, stating that the CacheCade virtual drive has been created.

7.7.5 Modifying a CacheCade Pro 2.0 Virtual Drive

You can modify the name and the write policy of an existing CacheCade virtual drive any time after a CacheCade virtual drive is created.

Perform the following steps to modify the CacheCade virtual drive:

1. In the **VD Mgmt** screen, navigate to the CacheCade virtual drive, and press the F2 key.
2. Navigate to **Properties**, and press Enter.

The following dialog appears.

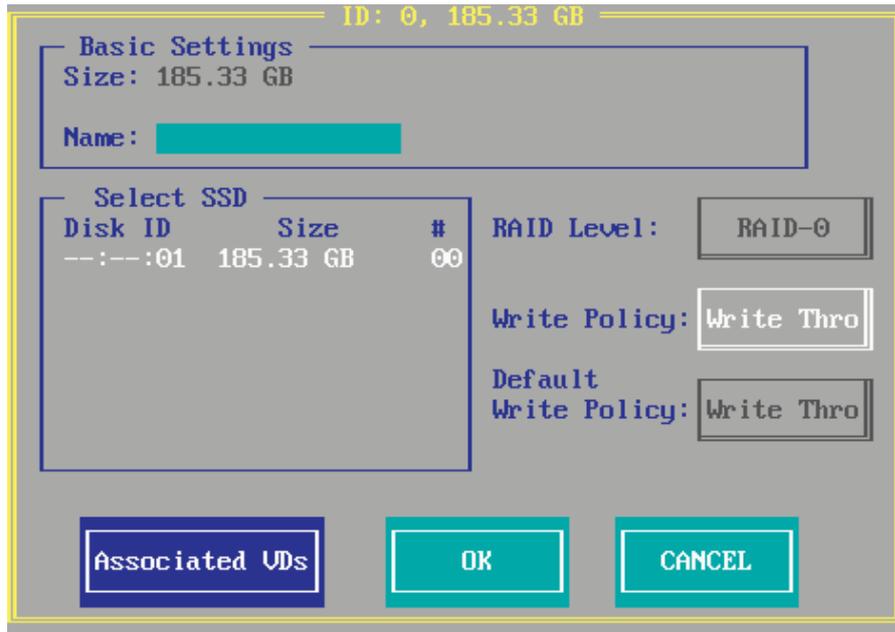


Figure 28. Modify CacheCade Virtual Drive

3. You can rename a CacheCade virtual drive in the **CacheCade Virtual Drive Name** field.
4. You can also modify the write policy by selecting one from the **Write Policy** field.
5. Press **OK**.

To view the virtual drives associated with this CacheCade virtual drive, click **Associated VDs** in the dialog above.

The **Associated Virtual Drives** dialog appears.

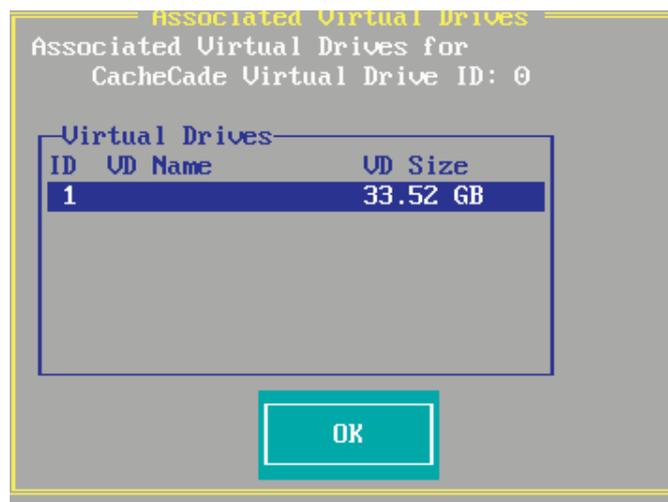


Figure 29. Associated Virtual Drives

You can view the ID, the name, and the size of the associated virtual drives.

7.7.6 Enabling SSD Caching on a Virtual Drive

You can enable SSD caching on a virtual drive. When you enable SSD caching on a virtual drive, that virtual drive becomes associated with an existing or with a future CacheCade SSD Caching virtual drive. This option is only available when the virtual drive's caching is currently disabled.

Perform the following steps to enable SSD caching on a virtual drive:

1. In the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Select **Enable Caching** and press Enter.

The following message dialog appears.

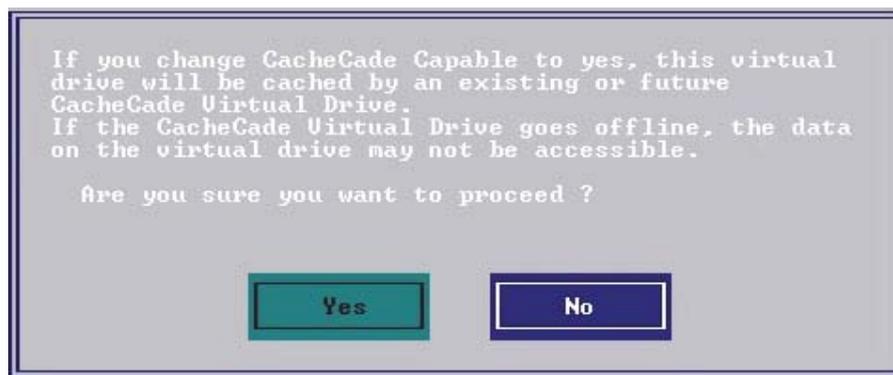


Figure 30. Message Box for Enabling SSD Caching

The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.

3. Click **Yes** to enable caching for that virtual drive.

7.7.7 Disabling SSD Caching on a Virtual Drive

You can disable caching on a virtual drive. When you disable SSD caching on a virtual drive, any associations that the selected virtual drive has with a CacheCade SSD Caching virtual drive are removed. This option is only available when the virtual drive's caching is currently enabled.

Perform the following steps to disable SSD Caching on a virtual drive:

1. In the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Select **Disable Caching** and press Enter.

The following message dialog appears.

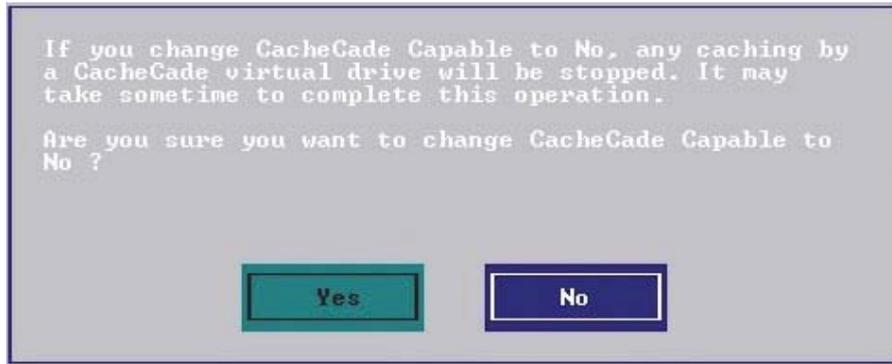


Figure 31. Message Box for Disabling SSD Caching

3. Click **Yes** to disable caching for that virtual drive.

7.7.8 Enabling or Disabling SSD Caching on Multiple Virtual Drives

You can enable or disable SSD caching on multiple virtual drives at one go.

Perform the follow steps to enable or disable SSD caching on multiple drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Select **Manage SSD Caching** and press Enter.

The **Manage SSD Caching** dialog appears.

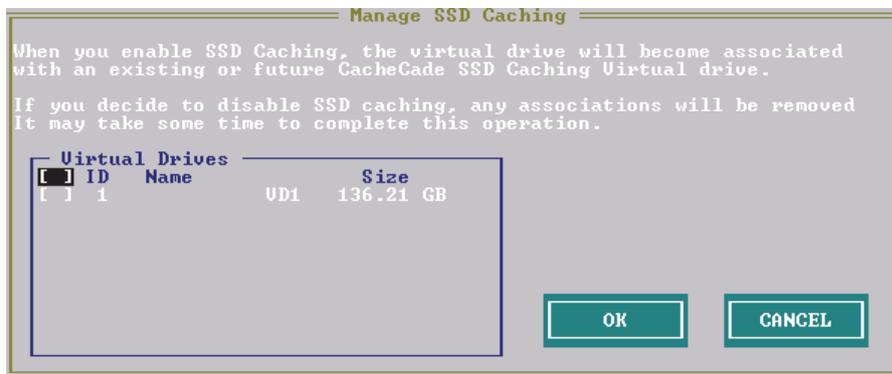


Figure 32. Manage SSD Caching

The virtual drives that have SSD caching enabled, have the check boxes next to them selected. The virtual drives that have SSD caching disabled, have deselected check boxes.

3. Select or deselect a check box to change the current setting of a virtual drive.
4. Click **OK** to enable/disable SSD caching on the selected virtual drives.

7.7.9 Deleting a Virtual Drive with SSD Caching Enabled

You can delete a virtual drive that has SSD caching enabled on it.

Perform the following steps to delete the virtual drive:

1. In the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Select **Delete VD**, and click **Yes**.

The following message dialog appears.



Figure 33. Message Box for Deleting Virtual Drive

***Note:** If you select the **Force delete to complete quickly** check box to delete the virtual drive, the data is not flushed before deleting the virtual drive. In this scenario, if you create this virtual drive after deleting it, there will be no data available.*

3. Press **Yes** to proceed with the delete operation.

7.8 Clearing the Configuration

You can clear all the existing configuration on virtual drives by deleting the virtual drives.

Perform the following steps to clear configuration:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Clear Configuration** and press Enter.

The following dialog appears.

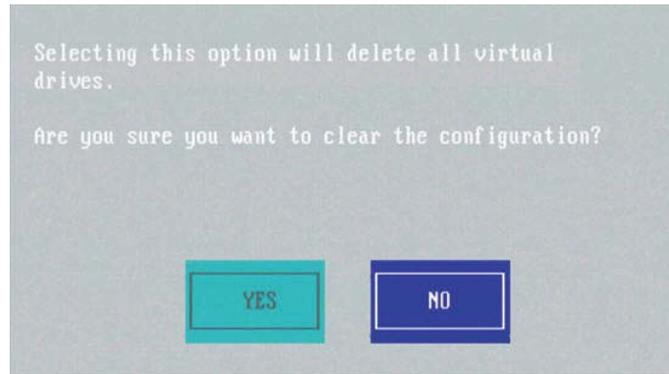


Figure 34. Clear Configuration

3. Press **Yes** to delete all the virtual drives.

7.9 Intel SafeStore Encryption Services

Intel SafeStore Encryption Services can encrypt data on the drives and use the drive-based key management to provide data security. This solution protects data in the event of theft or loss of physical drives. If you remove a self-encrypting drive from its storage system or the server in which it resides, the data on that drive is encrypted, and becomes useless to anyone who attempts to access it without the appropriate security authorization.

7.9.1 Enabling Drive Security

This section describes how to enable, change, and disable the drive security, and how to import a foreign configuration by using the SafeStore Encryption Services advanced software.

To enable security on the drives, you need to perform the following actions to set drive security:

- Enter a security key identifier.
A security key identifier appears whenever you have to enter a security key.
- Enter a security key.
After you create a security key, you can create secure virtual drives by using the key. You must use the security key to perform certain operations.

You can improve security by entering a password. To provide additional security, you can request for the password whenever anyone boots the server.

Perform the following steps to enable drive security:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.

3. Navigate to **Enable Security**, and press Enter.

The **Create Security Key** dialog appears.

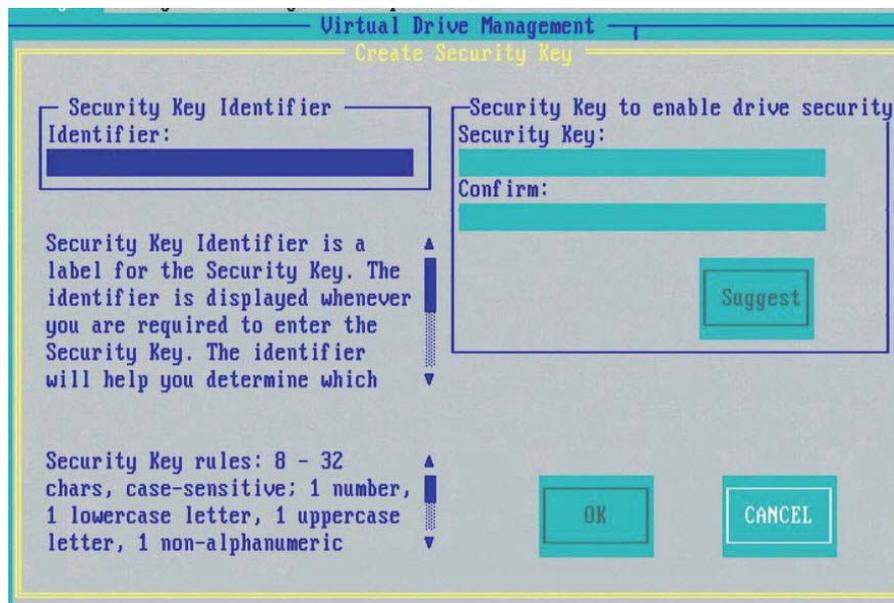


Figure 35. Create Security Key

4. Either use the default security key identifier, or enter a new security key identifier.

Note: After you create a security key, the **Enable Security** option is disabled. This option is reenabled only after you delete the existing key.

5. Either click **Suggest** to ask the system to create a security key, or you can enter a new security key.
6. Reenter the new security key to confirm it.

Caution: *If you forget the security key, you lose access to your data. Be sure to record your security key information. You might need to enter the security key to perform certain operations.*

The security key is case sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric character (for example, < > @ +). The space character is not permitted.

Note: *Non-U.S. keyboard users must be careful not to enter double-byte character set (DBCS) characters in the security key field. The firmware works with the ASCII character set only.*

7.9.2 Changing Security Settings

Perform the following steps to change the encryption settings for the security key identifier, security key, and password:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.
3. Select **Change Security Settings**, and press Enter.

The **Change Security Key** dialog appears.

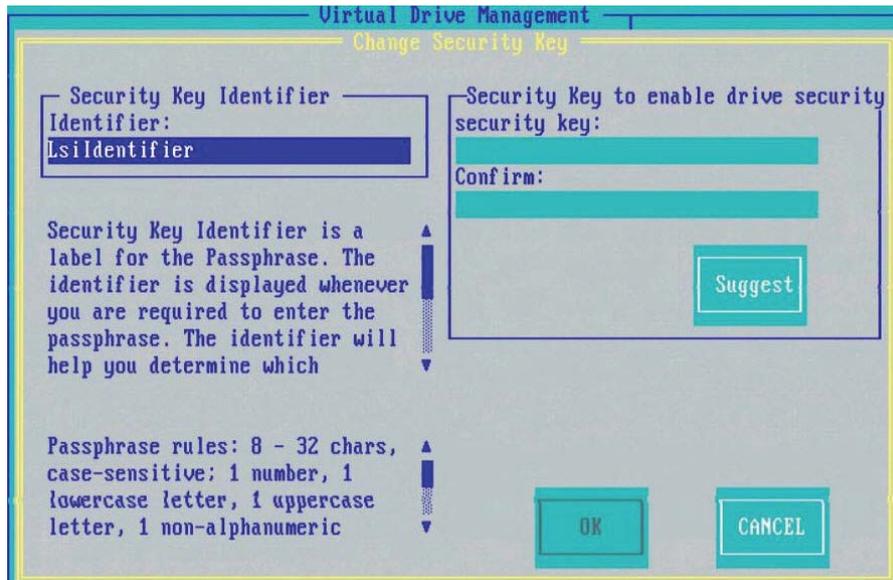


Figure 36. Change Security Key

4. Either keep the existing security key identifier, or enter a new security key identifier.

Note: *If you change the security key, you need to change the security key identifier. Otherwise, you cannot differentiate between the security keys.*

5. Either click **Suggest** to ask the system to create a security key, or you can enter a new security key.
6. Reenter the new security key to confirm it.

Caution: *If you forget the security key, you lose access to your data. Be sure to record your security key information. You might need to enter the security key to perform certain operations.*

The security key is case sensitive. It must be between 8 and 32 characters and contain at least one number, one lowercase letter, one uppercase letter, and one nonalphanumeric character (for example, < > @ +). The space character is not permitted.

Note: Non-U.S. keyboard users must be careful not to enter DBCS characters in the Security Key field. The firmware works with the ASCII character set only.

7.9.3 Disabling Drive Security

If you disable drive security, your existing data is not secure and you cannot create any new secure virtual drives. Disabling drive security does not affect data security on foreign drives. If you removed any drives that were previously secured, you still need to enter the password when you import them. Otherwise, you cannot access the data on those drives.

If there are any secure drive groups on the controller, you cannot disable drive security. A warning dialog appears if you attempt to do so. To disable drive security, you first must delete the virtual drives on all the secure drive groups.

Perform the following steps to disable drive security:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Drive Security**, and press Enter.
3. Select **Disable Security**.
A message box appears.
4. To disable drive security, press **Yes** to delete the security key.

Note: If you disable drive security, you cannot create any new encrypted virtual drives and the data on all encrypted unconfigured drives will be erased. Disabling drive security does not affect the security or data of foreign drives.

7.9.4 Importing or Clearing a Foreign Configuration

A foreign configuration is a RAID configuration that already exists on a replacement set of drives that you install in a computer system. You can use the Ctrl-R Utility to import the foreign configuration to the RAID controller or to clear the foreign configuration so that you can create a new configuration by using these drives.

To import a foreign configuration, you must perform the following tasks:

- Enable security to permit importation of locked foreign configurations. You can import unsecured or unlocked configurations when security is disabled.
- If a locked foreign configuration is present and security is enabled, enter the security key, and unlock the configuration.
- Import the foreign configuration.

If one or more drives are removed from a configuration, by a cable pull or drive removal for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

Verify whether any drives are left to import because the locked drives can use different security keys. If any drives remain, repeat the import process for the remaining drives. After all the drives are imported, there is no configuration to import.

Note: When you create a new configuration, the Ctrl-R Utility shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you first must clear the configuration on those drives.

You can import or clear a foreign configuration from the **VD Mgmt** menu or from the **Foreign View** menu.

Perform the following steps to import or clear a foreign configuration from the **VD Mgmt** menu:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Foreign Config**, and press Enter.

The foreign configuration options **Import** and **Clear** appear.

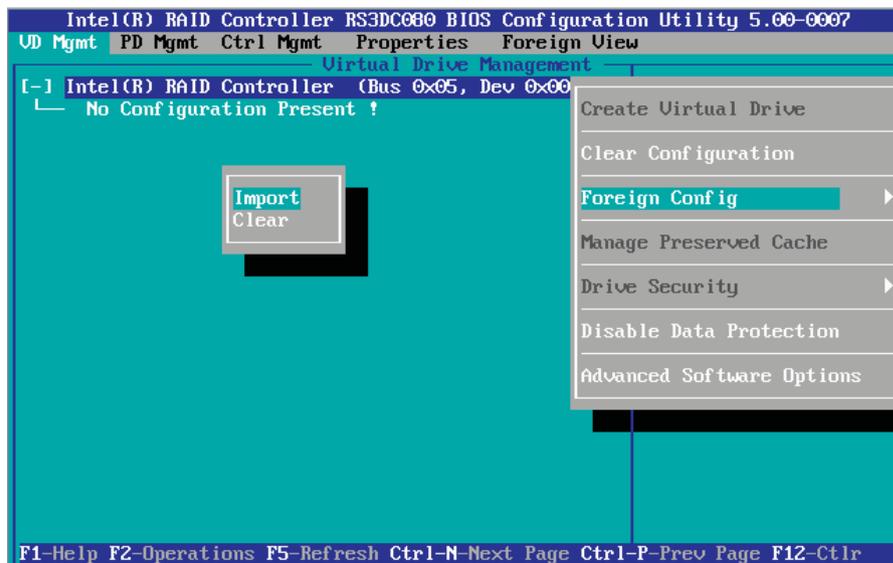


Figure 37. Foreign Configuration Options

3. Navigate to the command you want to perform.
 - To import a foreign configuration, go to [step 4](#).
 - To clear a foreign configuration, go to [step 6](#).
4. To import a foreign configuration, select **Import**, and press Enter.

The following dialog appears.

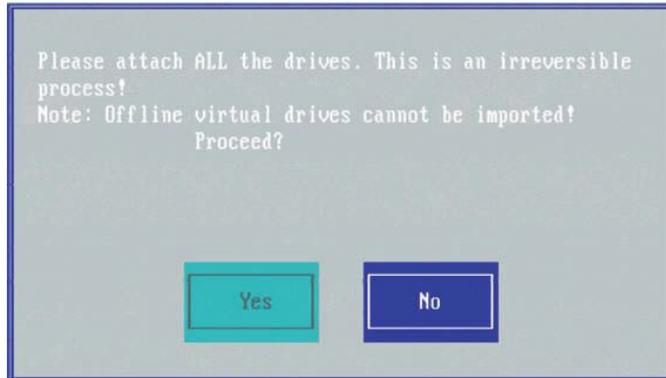


Figure 38. Foreign Configuration – Import

5. Press **Yes** to import the foreign configuration from all the foreign drives. Repeat the import process for any remaining drives.

Because locked drives can use different security keys, you must verify whether there are any remaining drives to be imported.

***Note:** When you create a new configuration, the Ctrl-R Utility shows only the unconfigured drives. Drives that have existing configurations, including foreign configurations, do not appear. To use drives with existing configurations, you first must clear the configuration on those drives.*

6. To clear a foreign configuration, select **Clear**, and press Enter.

The following dialog appears.



Figure 39. Foreign Configuration – Clear

7. Press **OK** to clear a foreign configuration.

***Note:** The operation cannot be reversed after it is started. Imported drives appear as Online in the Ctrl-R Utility.*

7.9.4.1 Foreign Configurations in Cable Pull and Drive Removal Scenarios

If one or more drives are removed from a configuration, by a cable pull or drive removal, for example, the configuration on those drives is considered a foreign configuration by the RAID controller.

The following scenarios can occur with cable pulls or drive removals.

Note: *To import the foreign configuration in any of the following scenarios, you must have all the drives in the enclosure before you perform the import operation.*

- Scenario 1: If all the drives in a configuration are removed and reinserted, the controller considers the drives to have foreign configurations.
Import or clear the foreign configuration. If you select **Import**, automatic rebuilds will occur in redundant virtual drives.

Note: *Start a consistency check immediately after the rebuild is complete, to ensure data integrity for the virtual drives.*

- Scenario 2: If some of the drives in a configuration are removed and reinserted, the controller considers the drives to have foreign configurations.
Import or clear the foreign configuration. If you select **Import**, automatic rebuilds will occur in redundant virtual drives.

Note: *Start a consistency check immediately after the rebuild is complete, to ensure data integrity for the virtual drives.*

- Scenario 3: If all the drives in a virtual drive are removed, but at different times, and reinserted, the controller considers the drives to have foreign configurations.
Import or clear the foreign configuration. If you select **Import**, all drives that were pulled before the virtual drive became offline will be imported and will be automatically rebuilt. Automatic rebuilds will occur in redundant virtual drives.
- Scenario 4: If the drives in a nonredundant virtual drive are removed, the controller considers the drives to have foreign configurations.
Import or clear the foreign configuration. No rebuilds will occur after the import operation because there is no redundant data to rebuild the drives.

7.10 Discarding Preserved Cache

If the controller loses access to one or more virtual drives, the controller preserves the data from the virtual drive. This preserved cache, is preserved until you import the virtual drive or discard the cache.

Certain operations, such as creating a new virtual drive, cannot be performed if preserved cache exists.

Caution: *If there are any foreign configurations, import the foreign configuration before you discard the preserved cache. Otherwise, you might lose data that belongs to the foreign configuration.*

Perform the following steps to discard the preserved cache:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Manage Preserved Cache**, and press Enter.

The **Manage Preserved Cache** dialog appears.

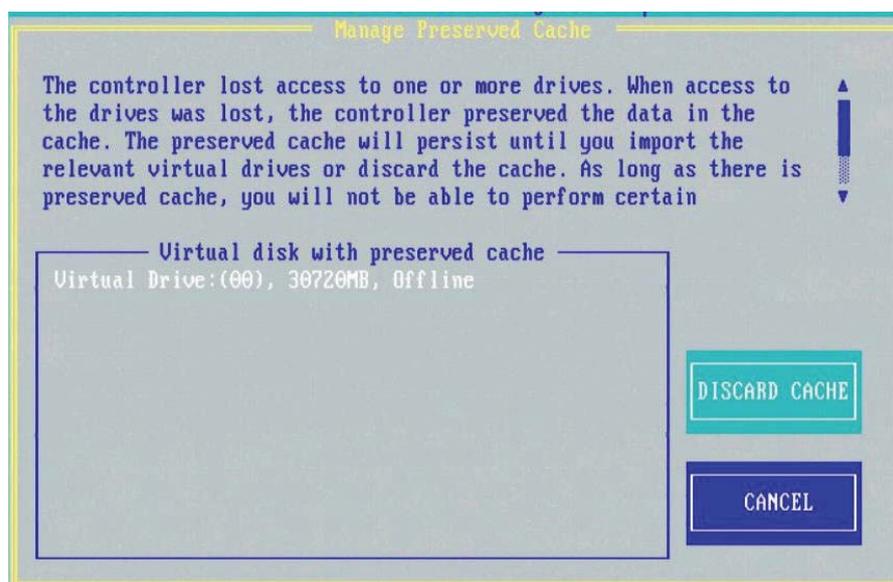


Figure 40. Manage Preserved Cache

3. Press **Discard Cache** to discard the preserved cache from the virtual drive. A message box appears, asking you to confirm your choice.
4. Press **OK** to continue.

7.11 Converting JBOD Drives to Unconfigured Good Drives

You can convert multiple JBODs to Unconfigured Good drives (from the **VD Mgmt** screen) or you can convert a particular JBOD drive to an Unconfigured Good drive (from the **Drive Management** screen).

Perform the following steps to convert multiple JBODs to Unconfigured Good drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Make Unconfigured Good**, and press Enter.

The **Convert JBOD to Unconfigured Good** dialog appears, which shows all JBODs available in the system.

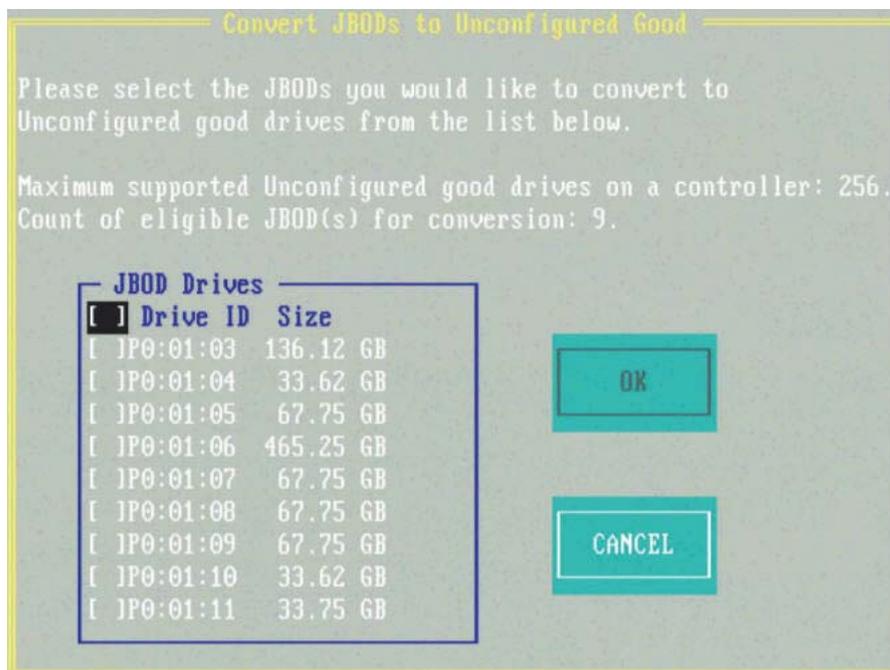


Figure 41. Convert JBOD to Unconfigured Good

3. Select the JBODs which you want configured as Unconfigured Good drives.
To select or deselect all the JBODs at one go, select the top most square brackets in the **JBOD Drives** box.
4. Press **OK**.
The selected JBODS are converted to Unconfigured Good drives.

Perform the following steps to convert a particular JBOD drive to an Unconfigured Good drive:

1. In the **Drive Management** screen, navigate to a JBOD drive, and press the F2 key.
2. Navigate to **Make Unconfigured Good**, and press Enter.
3. Press **Yes** in the confirmation message box to proceed.

7.12 Converting Unconfigured Good Drives to JBOD Drives

You can convert multiple Unconfigured Good drives to JBOD drives (from the **VD Mgmt** screen) or you can convert a particular Unconfigured Good drive to a JBOD drive (from the **Drive Management** screen).

Perform the following steps to convert multiple Unconfigured Good drives to JBOD drives:

1. In the **VD Mgmt** screen, navigate to the controller, and press the F2 key.
2. Navigate to **Make JBOD**, and press Enter.

The **Convert Unconfigured Good to JBOD** dialog appears, which shows all Unconfigured Good drives available in the system.

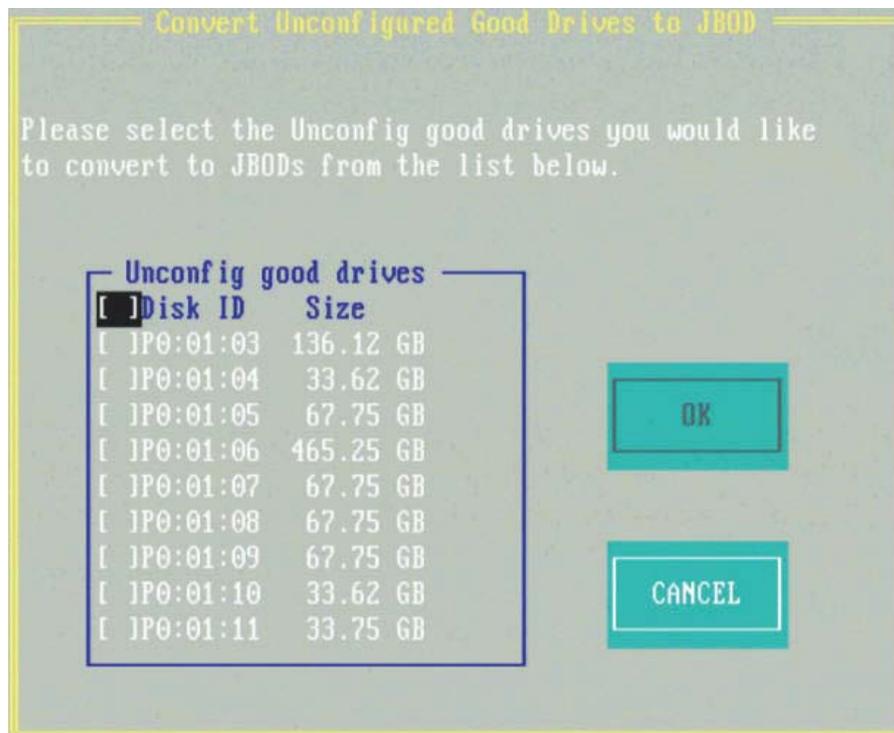


Figure 42. Convert Unconfigured Good to JBOD

3. Select the Unconfigured Good drives which you want configured as JBODs.
To select or deselect all the Unconfigured Good drives at one go, select the top most square brackets in the **Unconfig good drives** box.
4. Press **OK**.
The selected Unconfigured Good drives are converted to JBOD drives.

Perform the following steps to convert a particular Unconfigured Good drive to a JBOD drive:

1. In the **Drive Management** screen, navigate to a Unconfigured Good drive, and press the F2 key.
2. Navigate to **Make JBOD**, and press Enter.
3. Press **OK** in the message confirmation box to continue.

7.13 Viewing and Changing Device Properties

This section explains how you can use the Ctrl-R Utility to view and change the properties for controllers, virtual drives, drive groups, physical drives, and BBUs.

7.13.1 Viewing Controller Properties

The Ctrl-R Utility shows information for one Intel SAS controller at a time. If your system has multiple Intel SAS controllers, you can view information for a different controller by pressing the F12 key and selecting a controller from the list.

Navigate to the **Properties** menu to view the properties of the active controller.

The information in the **Properties** screen (Figure 19) is read only. Most of this information is self-explanatory. To view additional properties, navigate to **Next**, and press Enter.

7.13.2 Modifying Controller Properties

You can change the properties of the controller in the **Ctrl Mgmt** menu.

Perform the following steps to change the controller properties:

1. Navigate to the **Ctrl Mgmt** menu to view the first **Controller Settings** screen.
2. You may change the values of the properties for the editable fields.

To change additional properties, such as link speed, battery properties, and power settings, press **Next** to go to the second **Controller Settings** screen.

3. Press **Apply**.

The following table describes all entries and options listed on both the **Controller Settings** screen. Leave these options at their default settings to achieve the best performance, unless you have a specific reason for changing them.

Table 15. Controller Settings

Option	Description
Alarm Control	Select this option to enable, disable, or silence the onboard alarm tone generator on the controller.
Coercion Mode	Use this option to force drives of varying capacities to the same size so they can be used in a drive group. The coercion mode options are None, 128MB-way, and 1GB-way. The number you choose depends on how much the drives from various vendors vary in their actual size.

Option	Description
BIOS Mode	Specifies the following options to set the BIOS boot mode: <ul style="list-style-type: none"> • Stop on Error: Shows the errors encountered during boot up and waits for your input. The firmware does not proceed with the boot process till you take some action • Ignore Error: Ignores errors and the firmware proceeds with boot. • Pause on Error: The firmware may halt due to hardware faults. If the firmware encounters no hardware faults, then the boot up continues. • SafeMode Error: Boots the controller to run in safe mode.
Boot Device	Use this option to select the boot device from the list of virtual drives and JBODs.
Rebuild Rate	Use this option to select the rebuild rate for drives connected to the selected controller. The rebuild rate is the percentage of system resources dedicated to rebuilding a failed drive. The higher the number, the more system resources that are devoted to a rebuild. The range of rebuild rate is between 0 and 100 percent.
BGI Rate	Use this option to select the amount of system resources dedicated to background initialization of virtual drives connected to the selected controller. The range of background initialization (BGI) rate is between 0 and 100 percent.
CC Rate	Use this option to select the amount of system resources dedicated to consistency checks of virtual drives connected to the selected controller. The range of Consistency Check (CC) rate is between 0 and 100 percent.
Recon. Rate	Use this option to select the amount of system resources dedicated to reconstruction of drives connected to the selected controller. The range of Recon rate is between 0 and 100 percent.
Patrol Rate	Use this option to select the rate for patrol reads for drives connected to the selected controller. The patrol read rate is the percentage of system resources dedicated to running a patrol read. The range of patrol read is between 0 to 100 percent.
Cache Flush Interval	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Cache Flush Interval is between 0 to 100 seconds.
Spinup Delay	Use this option to control the interval (in seconds) between spinup of drives connected to this controller. The delay prevents a drain on the system's power supply that would occur if all drives spun up at the same time. The range of Spinup Delay is between 0 to 255 seconds.
Spinup Drive	Use this option to control the interval at which the contents of the onboard data cache are flushed. The range of Spinup Drive is between 0 to 255 seconds.
Maintain PD Fail History	Use this option to maintain the history of all drive failures.
Enable Controller BIOS	Use this option to enable or disable the BIOS for the selected controller. If the boot device is on the selected controller, the BIOS must be enabled. Otherwise, the BIOS should be disabled, or you might be unable to use a boot device elsewhere.

Option	Description
Enable Stop CC on Error	Use this option to stop a consistency check when the controller BIOS encounters an error.
Auto Enhanced Import	Use this option to import automatically at boot time.
Set Factory Defaults	Use this option to load the default Ctrl-R Utility settings.
Manage Link Speed	Use this option to change the link speed between the controller and the expander, or between a controller and a drive that is directly connected to the controller.
Manage Power Save	Use this option to reduce the power consumption of drives that are not in use, by spinning down the unconfigured good drives, hot spares and configured drives.
Manage Battery	Use this option to view information about the BBU, if the selected controller has a BBU.
Emergency Spare	Use this option to commission unconfigured good drives or global hotspares as emergency spare drives. You can select from the options None , UG (Unconfigured Good), GHS (Global Hotspare), or UG and GHS (Unconfigured Good and Global Hotspare).
Enable Emergency for SMARTer	Use this option to commission emergency hot spare drives for predictive failure analysis events.

7.13.3 Viewing and Changing Virtual Drive Properties

The Ctrl-R Utility shows the properties, policies, and the operations for virtual drives.

To view these items for the currently selected virtual drive and to change some of these settings, perform the following steps:

1. In the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Press Enter.

The **Virtual Drive Properties** dialog appears.

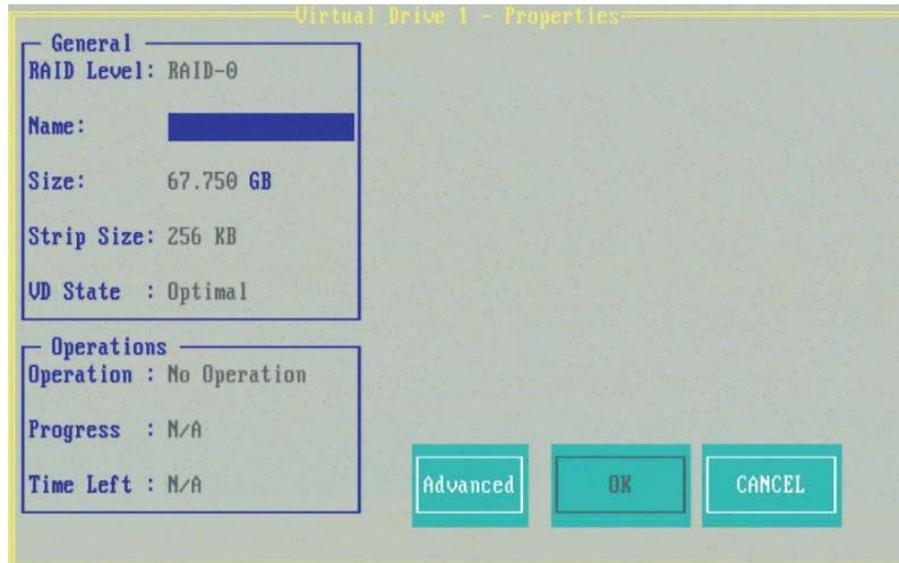


Figure 43. Virtual Drive Properties

The **General** box shows the virtual drive's RAID level, name, state, size, and strip size.

The **Operations** box lists any operation (performed on the virtual drive) in progress, along with its progress status and the time remaining for the operation to be completed.

3. You may change the settings for the fields that are enabled in this dialog.

Caution: Before you change a virtual drive configuration, back up any data on the virtual drive that you want to save.

4. Press **OK** to save your changes.
5. Press **Advanced** to view additional virtual drive properties.

The **Create Virtual Drive – Advanced** dialog appears.

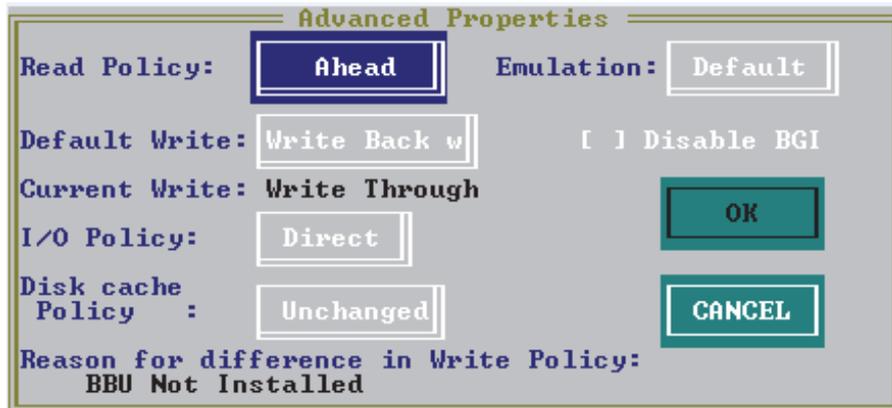


Figure 44. Create Virtual Drive – Advanced

You can view the virtual drive policies that were defined when the storage configuration was created.

You can also select **Initialize** and/or **Configure Hot Spare** to initialize the newly created virtual drive or to configure the virtual drive as a hot spare, respectively.

If you select **Initialize** and/or **Configure Hot Spare**, messages pertaining to initializing the virtual drive or to configuring the virtual drive as a hot spare appear only after you exit the **Virtual Drive Properties** dialog.

7.13.4 Deleting a Virtual Drive

You can delete any virtual drive on the controller if you want to reuse that space for a new virtual drive. The Ctrl-R Utility lists configurable drive groups where there is space to configure. If multiple virtual drives are defined on a single drive group, you can delete a virtual drive without deleting the entire drive group.

Caution: *Back up any data that you want to keep before you delete a virtual drive.*

Perform the following steps to delete a virtual drive:

1. In the **VD Mgmt** screen, navigate to the virtual drive, and press the F2 key.
2. Navigate to **Delete VD**, and press Enter.
A message box appears, asking you to confirm the deletion.
3. Press **OK** to delete the virtual drive.

7.13.5 Deleting a Virtual Drive Group

You can delete a virtual drive group. On deleting a drive group, all the virtual drives in that drive group also are deleted.

Caution: *Back up any data that you want to keep before you delete a virtual drive group.*

Perform the following steps to delete a drive group:

1. In the **VD Mgmt** screen, navigate to a drive group, and press the F2 key.
2. Navigate to **Delete Drive Group**, and press Enter.

The drive group is deleted and is removed from the **VD Mgmt** screen.

7.13.6 Expanding a Virtual Drive

You can increase the size of a virtual drive to occupy the remaining capacity in a drive group.

Perform the following steps to expand the size of a virtual drive:

1. In the **VD Mgmt** screen, select the virtual drive whose size you want to expand and press the F2 key.
2. Navigate to **Expand VD**, and press Enter.

The **Expand Virtual Drive** dialog appears.

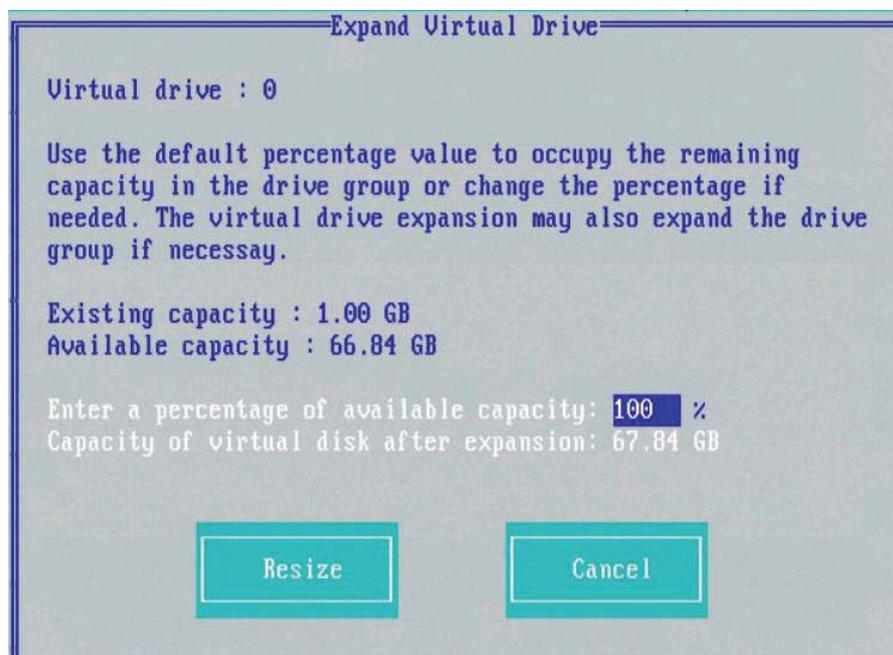


Figure 45. Expand Virtual Drive

3. Enter the percentage of the available capacity that you want the virtual drive to use. For example, if 100 GB of capacity is available and you want to increase the size of the virtual drive by 30 GB, select 30 percent.
4. Press **Resize** to determine the capacity of the virtual drive after expansion. The virtual drive expands by the selected percentage of the available capacity.

7.13.7 Erasing a Virtual Drive

Virtual drive erase operates on a specified virtual drive and overwrites all user-accessible locations. It supports nonzero patterns and multiple passes. Virtual drive erase optionally deletes the virtual drive and erases the data within the virtual drive's logical base address range. Virtual drive erase is a background operation that posts events to notify users of their progress.

Perform the following steps to perform the virtual drive erase operation:

1. In the **VD Mgmt** screen, select a virtual drive, and press the F2 key.
2. Navigate to **Erase VD**, and press Enter.

A menu appears displaying the following modes:

- **Simple** – Specifies a single-pass erase operation that writes pattern A to the virtual drive.
- **Normal** – Specifies a three-pass erase operation that first overwrites the virtual drive content with random values, then overwrites it with pattern A, and then overwrites it with pattern B.
- **Thorough** – Specifies a nine-pass erase operation that repeats the **Normal** erase three times.
- **Stop Erase**: Stops the erase operation that has already been started. This option is disabled at first. Once the erase operation begins, this option is enabled.

3. Select a mode and press Enter.

A message box appears.



Figure 46. Erase Virtual Drive

4. To delete the virtual drive after the erase operation has been completed, select the **Delete Virtual Drive after Erase operation** check box.

5. Press **Yes** for the erase operation to start.

Once the Drive Erase operation has started, the **Simple**, **Normal**, and **Thorough** options are disabled and the **Stop Erase** option is enabled.

7.13.8 Managing Link Speed

The Managing Link Speed feature lets you change the link speed between the controller and an expander, or between the controller and a drive that is directly connected to the controller.

All phys in a SAS port can have different link speeds or can have the same link speed.

You can select a link speed setting. However, if phys in a SAS port have different link speed settings and if a phy is connected to a drive or an expander, the firmware overrides the link speed setting you have selected. Instead, the firmware uses the common maximum link speed among all the phys.

Perform the following steps to change the link speed:

1. In the **Controller Settings** screen, press **Next**.
The second **Controller Settings** screen appears.
2. Press **Manage Link Speed**.

The **Manage Link Speed** dialog appears.

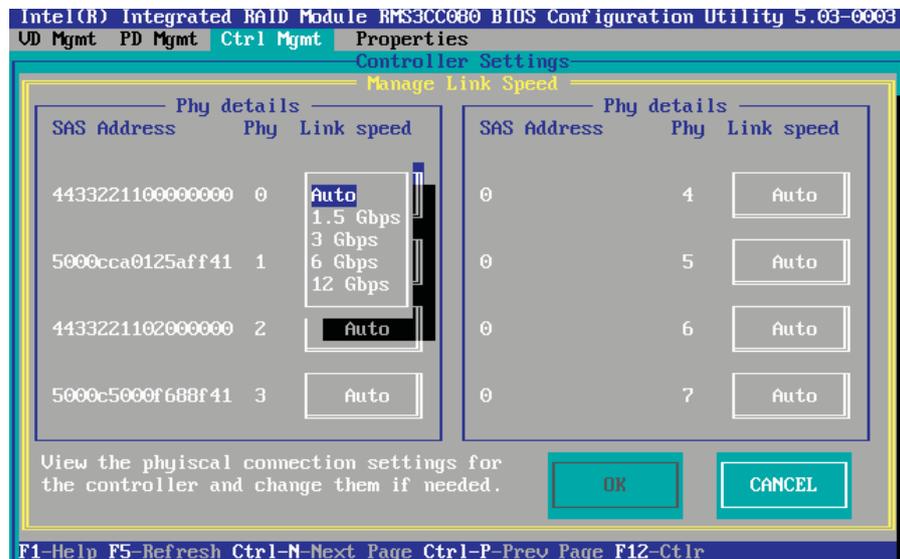


Figure 47. Manage Link Speed

- The **SAS Address** column shows the SAS address that uniquely identifies a device in the SAS domain.
- The **Phy** column shows the system-supported phy link values. The phy link values are from 0 through 7.

- The **Link Speed** column shows the phy link speeds.
- 3. Select the desired link speed by using the drop-down list.
The link speed values are Auto,1.5Gb/s, 3Gb/s, 6Gb/s, or 12Gb/s.

Note: By default, the link speed in the controller is Auto or the value last saved by you.

- 4. Press **OK**.
A message box appears, asking you to restart your system for the changes to take effect.
- 5. Press **OK**.
The link speed value is now reset. The change takes place after you restart the system.

7.13.9 Managing Power Save Settings for the Controller

You can change the controller’s power-save settings by using the Dimmer Switch enhancement (Power-Save mode).

Perform the following steps to change the power save settings:

- 1. Navigate to the second **Controller Settings** screen.
- 2. Navigate to **Manage Power Save**, and press Enter.

The **Manage Power Save** dialog appears.



Figure 48. Manage Power Save

- 3. Select the **Spin down Unconfigured drives** check box to let the controller enable the unconfigured drives to enter the Power-Save mode.

4. Select the **Spin down Hot Spares** check box to let the controller enable the Hot spare drives to enter the Power-Save mode.
5. Select the drive standby time from the **Drive Standby Time** drop-down list.

Note: The **Drive Standby Time** drop-down list is enabled only if any of the preceding check boxes are checked. The drive standby time can be 30 minutes, 1 hour, 90 minutes, or 2 hours through 24 hours.

6. Press **OK**.
A message box appears, asking you to save the power-save settings.
7. Press **Yes** to save the settings.

7.13.10 Managing Power Save Settings for the Drive Group

You can change the power save settings for a selected drive group.

Perform the following steps to change the power save settings for a drive group:

1. Navigate to a drive group in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Manage Power Save Settings** and press Enter.

The **Manage Power Save Settings** dialog appears.



Figure 49. Manage Power Save Settings – Drive Group

3. Select a power save mode from the **Select power save mode** drop-down list.
A description of the selected mode appears in the dialog.
4. Press **OK**.

7.13.11 Managing BBU Information

If your SAS controller has a BBU, you can view information about it and change some settings.

A learning cycle is a battery calibration operation that the controller performs periodically to determine the condition of the battery.

Perform the following steps to view and change the battery settings:

1. Navigate to the second **Controller Settings** screen and select **Manage Battery**.

The **Battery Properties** dialog appears. Most of the battery properties are read only.

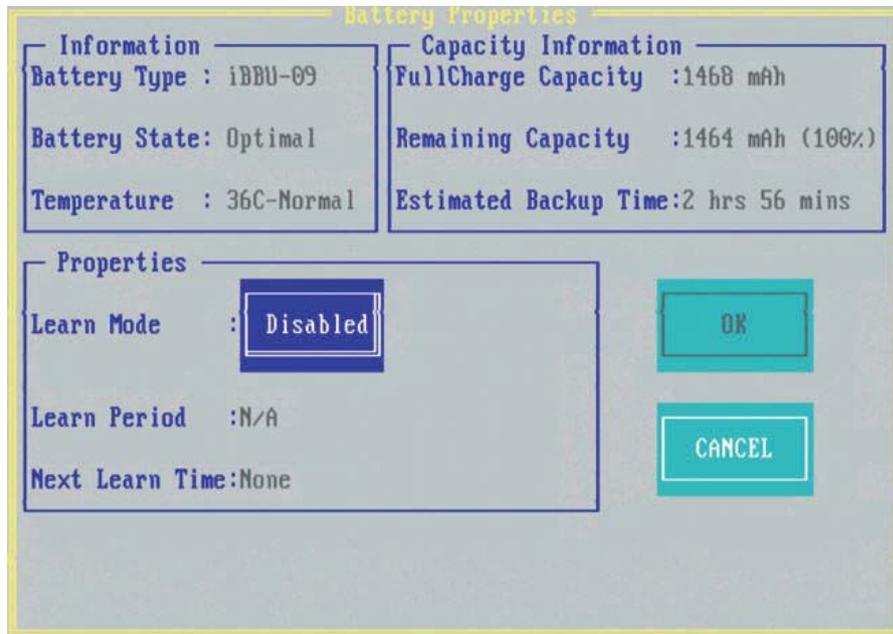


Figure 50. Battery Properties

If the **Battery State** field has a value other than **Optimal**, the **Non-Optimal Reason** field appears at the bottom of the **Battery Properties** dialog. The **Non-Optimal Reason** field is a read-only field and states a reason for the non optimal state of the battery.

2. Select a battery learn mode from the **Learn Mode** drop-down list.

The values in the drop-down list differ based on whether the battery supports transparent learn cycles.

— If the battery supports transparent learn, the following values appear in the **Learn Mode** drop-down list:

- G **Transparent** – The firmware tracks the time since the last learning cycle and performs a learn cycle when it is due.
- G **Disabled** – The firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.
- G **Unknown** – The firmware warns about a pending learning cycle. You can start a learning cycle manually. After the learning cycle completes, the firmware resets the counter and warns you when the next learning cycle time is reached.

- If the battery does not support transparent learn, the following values appear in the **Learn Mode** drop-down list:
 - G **Automatic** – The firmware tracks the time since the last learning cycle and performs a learn cycle when due. Write caching need not be disabled.
 - G **Disabled** – The firmware does not monitor or initiate a learning cycle. You can schedule learning cycles manually.
 - G **Disabled (Warning Only)** – The firmware never initiates a battery learn cycle but notifies you through events when a learn cycle is needed.
3. Press **OK** to change the learn mode.

7.13.12 Managing Dedicated Hot Spares

A dedicated hot spare is used to replace failed drives only in a selected drive group that contains the hot spare. You can create or delete dedicated hot spares in the **Virtual Drive Management** screen.

Perform the following steps to create or delete dedicated hot spares:

1. Navigate to a drive group in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Manage Dedicated Hotspare**, and press Enter.

The **Dedicated Hotspare** dialog appears, which shows a list of all hot spares that are available to create dedicated hot spares.

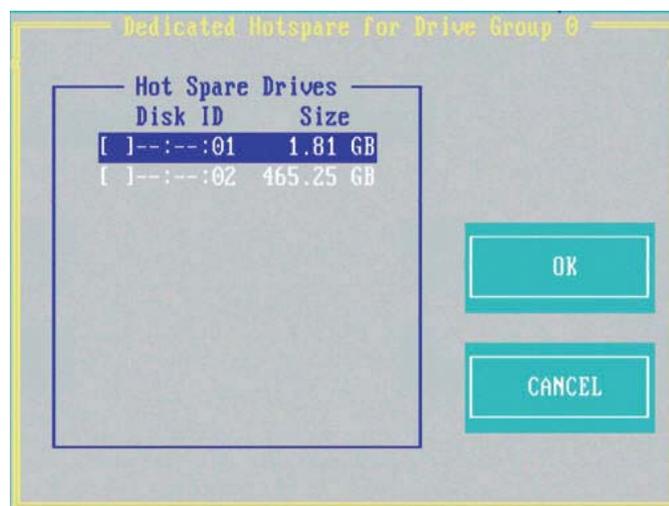


Figure 51. Dedicated Hotspare

3. Perform one of these steps:
 - To create a dedicated hot spare, select a drive and press **OK**.
 - To delete a dedicated hot spare, deselect the hot spare and press **OK**.

7.13.13 Securing a Drive Group

If a drive group is created with FDE drives (security enabled drives) and at the time of creation, the security is set to **No**; later, you can secure that drive group using encryption.

Perform the following steps to secure a drive group:

1. Navigate to the **VD Mgmt** screen, navigate to the drive group that you want to secure, and press the F2 key.
2. Navigate to **Secure Drive Group**, and press Enter.
A message box appears asking for your confirmation.
3. Press **Yes** to secure the drive group.

Note: After a virtual drive is secured, you will not be able to remove the encryption without deleting the virtual drive.

7.13.14 Setting LED Blinking

You can use the **Locate** option to make the LEDs blink on the physical drives used by a virtual drive. You can choose to start or stop the LED blinking.

Perform the following steps to start or stop LED blinking:

1. Navigate to the **Drive Management** screen (in the **PD Mgmt** menu).
2. Select a physical drive, and press the F2 key.
3. Navigate to **Locate**, and press Enter.
The **Start** and the **Stop** options appear.
4. Perform one of these actions:
 - Select **Start**, and press Enter to start LED blinking.
 - Select **Stop**, and press Enter to stop LED blinking.

Note: Both the **Start** and **Stop** options of **Locate** only work if the drive is installed in a drive enclosure.

7.13.15 Performing a Break Mirror Operation

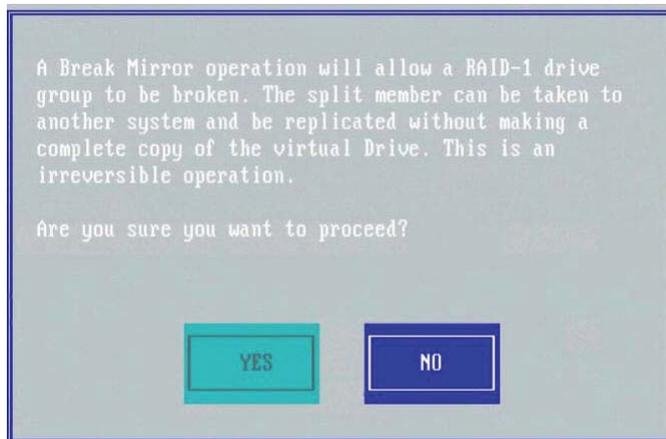
You can perform a Break Mirror operation on a drive group. The Break Mirror operation enables a RAID 1 configured drive group to be broken into two volumes. You can use one of the volumes in another system and replicate it without making a copy of the virtual drive.

Perform the following steps to perform a break mirror operation:

1. Navigate to the **VD Mgmt** screen, navigate to a drive group on which you want to perform the break mirror operation, and press the F2 key.

2. Navigate to **Break Mirror**, and press Enter.

The following message box appears, asking for your confirmation.



3. Press **Yes** to proceed.

7.13.16 Performing a Join Mirror Operation

You can perform a join mirror operation on a drive group to continue using the modified virtual drive or to reuse the original virtual drive.

Perform the following steps to perform a join mirror operation:

1. Navigate to the **VD Mgmt** screen, navigate to a drive group on which you want to perform the join mirror operation, and press the F2 key.
2. Navigate to **Join Mirror**, and press Enter.

The following dialog appears.

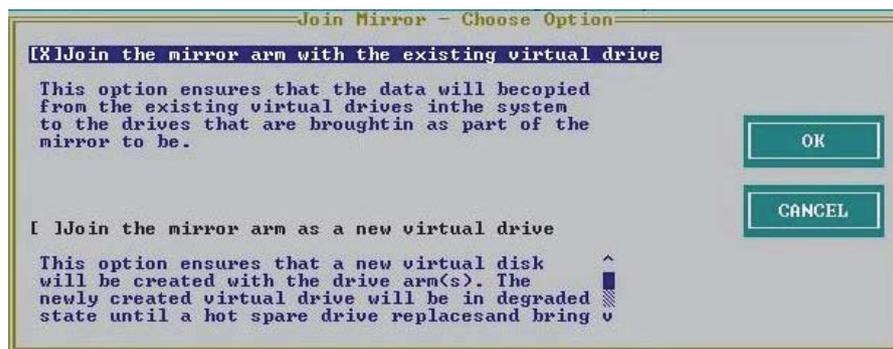


Figure 52. Join Mirror – Choose Option

3. Select one of the options and press **OK**.

If you select **Join the mirror arm with the existing virtual drive**, the following confirmation dialog appears.



Figure 53. Confirmation Message

If you select **Join the mirror arm as a new virtual drive**, the following confirmation dialog appears.



Figure 54. Confirmation Message

4. Press **Yes** to proceed.

The following dialog appears.



Figure 55. Join Mirror – Choose Option

5. Select one of the options and press **OK**.

7.14 Managing Storage Configurations

This section describes how to use the Ctrl-R Utility to maintain and manage storage configurations.

7.14.1 Initializing a Virtual Drive

When you create a new virtual drive, the Ctrl-R Utility asks whether you would like to initialize the virtual drive. If you do not want to initialize the virtual drive at that stage, you can initialize the drive later.

Perform the following steps to initialize a virtual drive:

1. Navigate to the **VD Mgmt** screen, navigate to a virtual drive, and press the F2 key.
2. Select **Initialization**, and press Enter.

The two initialization options, **Fast Init** and **Slow Init**, appear.

3. Select one of the two options, and press Enter.

A confirmation dialog appears.



Figure 56. Initialize a Virtual Drive

4. Press **Yes** to begin initialization.

Caution: Initialization erases all data on the virtual drive. Make sure to back up any data you want to keep before you initialize a virtual drive. Make sure the operating system is not installed on the virtual drive you are initializing.

7.14.2 Running a Consistency Check

You should periodically run a consistency check on fault-tolerant virtual drives (RAID 1, 5, 6, 10, 50, or 60 configurations; RAID 0 does not provide data redundancy). A consistency check scans the virtual drive to determine whether the data has become corrupted and needs to be restored.

For example, in a system with parity, checking consistency means computing the data on one drive and comparing the results with the contents of the parity drive. You must run a consistency check if you suspect that the data on the virtual drive might be corrupted.

Caution: *Make sure to back up the data before you run a consistency check, if you think the data might be corrupted.*

Perform the following steps to run a consistency check:

1. Navigate to a virtual drive in the **VD Mgmt** screen, and press the F2 key.
2. Navigate to **Consistency Check**, and press Enter.
3. Navigate to **Start**, and press Enter.

The consistency check starts and checks the redundant data in the virtual drive.

If you attempt to run a consistency check on a virtual drive that has not been initialized, a confirmation dialog appears, asking for your confirmation.



Figure 57. Consistency Check

4. Press **Yes** to run the consistency check.

7.14.3 Rebuilding a Physical Drive

If a drive in a redundant virtual drive (RAID 1, 5, 6, 10, 50, or 60) fails, you must rebuild that drive on a hot spare drive to prevent data loss.

Perform the following steps to rebuild a physical drive:

1. Navigate to the **Drive Management** screen (in the **PD Mgmt** menu), and press the F2 key.
2. Select **Rebuild**, and press Enter.

The rebuild operation starts.

7.14.4 Performing a Copyback Operation

You can perform a copyback operation on a selected drive.

The copyback operation copies data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation often creates or restores a specific physical configuration for a drive group (for example, a specific arrangement of drive group members on the device I/O buses).

Perform the following steps to perform the copyback operation:

1. Navigate to the **Drive Management** screen, navigate to a physical drive, and press the F2 key.
2. Navigate to **Copyback**, and press Enter.

The following dialog appears.

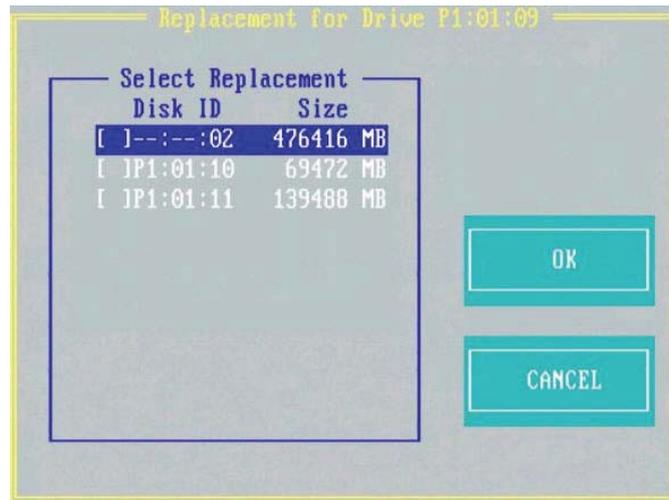


Figure 58. Copyback Operation

3. Select the replacement drive to which you want the data copied.
4. Press **OK**.

The copyback operation is performed on the selected drive.

7.14.5 Removing a Physical Drive

You might sometimes need to remove a non-failed drive that is connected to the controller. Preparing a physical drive for removal spins the drive into a power save mode.

Perform the following steps to prepare a physical drive for removal:

1. Navigate to the **Drive Management** screen, and press the F2 key.
2. Select **Prepare for Removal**, and press Enter.

The physical drive is now in a power save mode.

If you change your mind and do not want to remove the drive, navigate to **Undo Removal**, and press Enter.

7.14.6 Creating Global Hot Spares

A global hot spare is used to replace a failed physical drive in any redundant array, as long as the capacity of the global hot spare is equal to or larger than the coerced capacity of the failed physical drive.

You can designate the hot spare to have enclosure affinity. In an enclosure affinity, if drive failures are present on a split backplane configuration, the hot spare first is used on the backplane in which it resides.

Perform the following steps to create global hot spares:

1. Navigate to the **Drive Management** screen, navigate to a physical drive that you want to change to a hot spare, and press the F2 key.
2. Select **Make Global HS**, and press Enter.

The physical drive is changed to a global hot spare. The status of the physical drive as a global hot spare appears in the **Drive Management** screen.

7.14.7 Removing a Hot Spare Drive

Perform these steps to remove a hot spare drive:

1. Navigate to the **Drive Management** screen, navigate to a hot spare drive that you want to remove, and press the F2 key.
2. Select **Remove Hot Spare drive**, and press Enter.

The hot spare drive is removed.

7.14.8 Making a Drive Offline

If a drive is part of a redundant configuration and you want to use it in another configuration, you can remove the drive from the first configuration and change the drive state to Unconfigured Good.

Caution: *After you perform this procedure, all data on that drive is lost.*

Perform the following steps to remove the drive from the configuration without harming the data on the virtual drive:

1. Navigate to the **Drive Management** screen, select a physical drive, and press the F2 key.
2. Navigate to **Place Drive Offline**, and press Enter.

The drive status changes to Unconfigured Good.

Caution: After you perform this step, the data on this drive is no longer valid.

7.14.9 Making a Drive Online

You can change the state of a physical drive to online. In an online state, the physical drive works normally and is a part of a configured virtual drive.

Perform the following steps to make a physical drive online:

1. Navigate to the **Drive Management** screen, select a physical drive, and press the F2 key.
2. Navigate to **Place Drive Online**, and press Enter.

The state of the physical drive changes to Online.

7.14.10 Instant Secure Erase

You can erase data on SED drives by using the **Instant Secure Erase** option in the **PD Mgmt** menu.

Perform the following steps to erase data on SED drives:

1. Navigate to the **Drive Management** screen, select a physical drive and press the F2 key.
2. Navigate to **Instant Secure Erase**, and press Enter.
A confirmation dialog appears, asking whether you would like to proceed.
3. Press **Yes** to proceed.

7.14.11 Erasing a Physical Drive

You can securely erase data on Non SEDs (normal HDDs) by using the **Drive Erase** option in the **PD Mgmt** menu.

For Non-SEDs, the erase operation consists of a series of write operations to a drive that overwrites every user-accessible sector of the drive with specified patterns. It can be repeated in multiple passes using different data patterns for enhanced security. The erase operation is performed as a background task.

Perform the following steps to erase data on Non SEDs:

1. Navigate to the **Drive Management** screen, select a physical drive and press the F2 key.
2. Navigate to **Drive Erase**, and press Enter.

A menu appears displaying the following modes:

- **Simple:** Specifies a single pass operation that writes pattern A to the physical drive.

- **Normal:** Specifies a three pass erase operation that first overwrites the physical drive content with random values, then overwrites it with pattern A and then overwrites it with pattern B.
 - **Thorough:** Specifies a nine pass erase operation that repeats Normal erase three times.
 - **Stop Erase:** This option is disabled. This option is disabled at first. Once the erase operation begins, this options is enabled.
3. Select a mode and press Enter.
On selecting **Simple**, **Normal**, or **Thorough**, a confirmation dialog appears.
 4. Press **Yes** on the confirmation dialog to proceed with the drive erase operation.
Once the Drive Erase operation has started, you are intimated with the progress of the operation. Also, the **Simple**, **Normal**, and **Thorough** modes are disabled and the **Stop Erase** mode is enabled.

8 6Gb/s Intel® RAID Controller Configuration 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.

8.1 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.

8.2 Detailed Configuration Steps using the Intel® RAID BIOS Console 2

8.2.1 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 128.

8.2.2 Screen and Option Descriptions

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

8.2.2.1 Toolbar Options

Table 16 describes the Intel® RAID BIOS Console 2 toolbar icons.

Table 16. 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.

**Table 16. Intel® RAID BIOS Console 2 Toolbar
Icon Descriptions (Cont.)**

Icon	Description
 A square icon with a grey background, containing a blue and white arrow pointing to the right, symbolizing exit.	Exit the Intel® RAID BIOS Console 2 utility.
 A square icon with a grey background, containing a yellow bell with a diagonal line through it, symbolizing silence or alarm.	Silence the alarm.

8.2.2.2 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

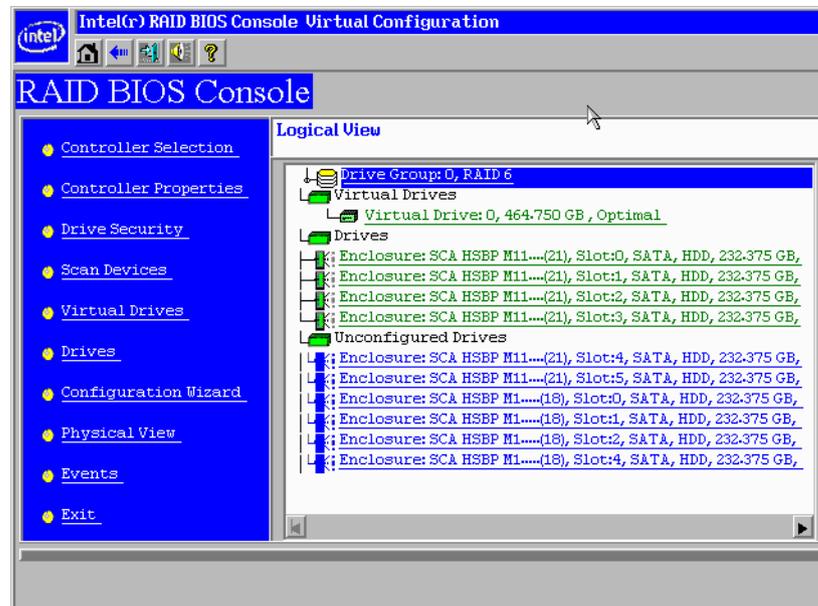


Figure 59. Intel® RAID BIOS Console 2 Menu

8.2.2.3 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 60. Intel® RAID BIOS Console 2 – Controller Selection

8.2.2.4 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	NVRAM Size	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 61. 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.

8.2.2.4.1 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**.

Properties			
Battery Backup	Present	Coercion Mode	None
Set Factory Defaults	No	S.M.A.R.T Polling	300 seconds
Cluster Mode	Disabled	Alarm Control	Enabled
Rebuild Rate	30	Patrol Read Rate	30
BGI Rate	30	Cache Flush Interval	4
CC Rate	30	Spinup Drive Count	2
Reconstruction Rate	30	Spinup Delay	6
Controller BIOS	Enabled	StopOnError	Enabled
NCQ	Enabled	Drive Powersave	Enabled

Figure 62. 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.

8.2.2.5 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.

8.2.2.6 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.

8.2.2.7 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.

8.2.2.8 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 128 provides detailed steps for using the Configuration Wizard.

8.2.2.9 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 315 for events and message descriptions.

8.2.2.10 Physical View/Logical View Option

This option toggles between Physical View and Logical View.

8.2.2.11 Exit

This option allows you to exit and reboot the system.

8.2.2.12 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.

8.3 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.

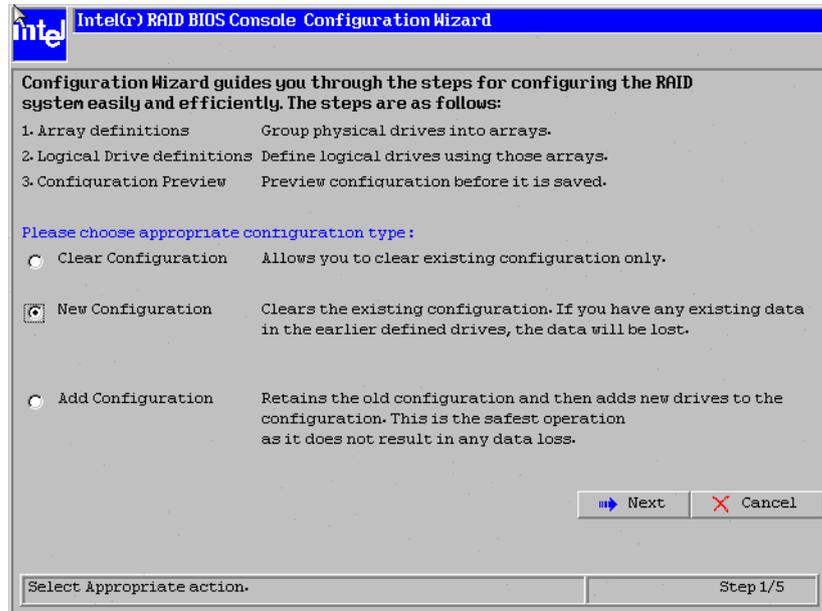


Figure 63. Intel® RAID BIOS Console 2 – Configuration Types

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

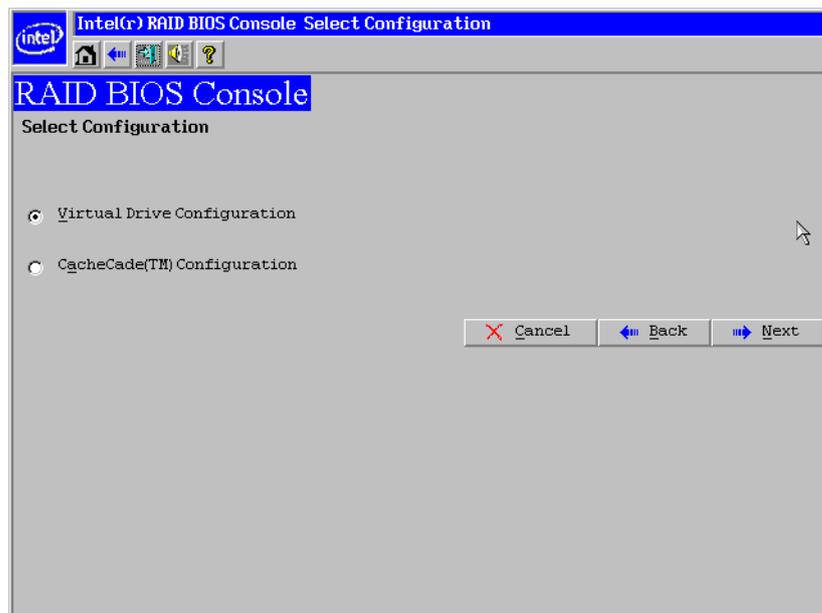


Figure 64. Selecting Configuration

4. Choose the configuration method and click Next.

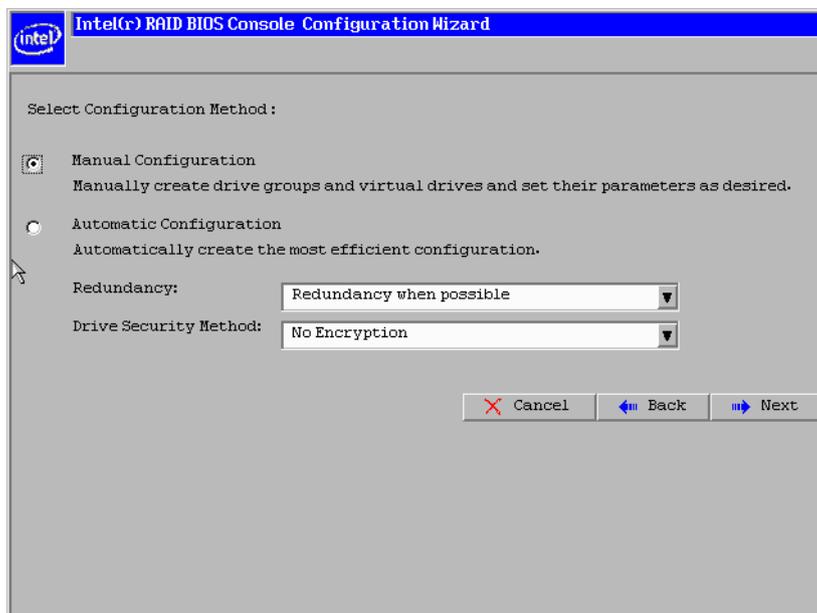


Figure 65. Intel® RAID BIOS Console 2 – Configuration Methods

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.

8.4 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 65).
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.

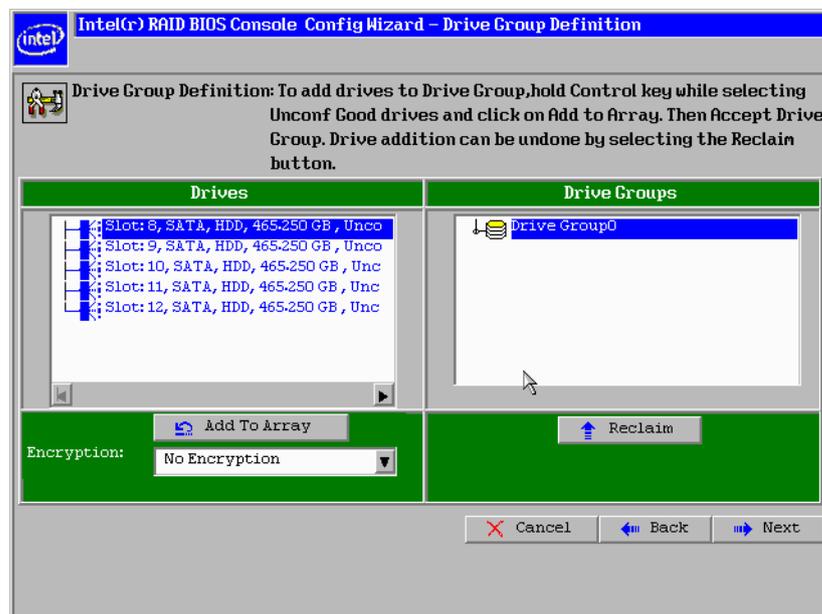


Figure 66. Intel® RAID BIOS Console 2 – Add Physical Drives to Array

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 141](#).

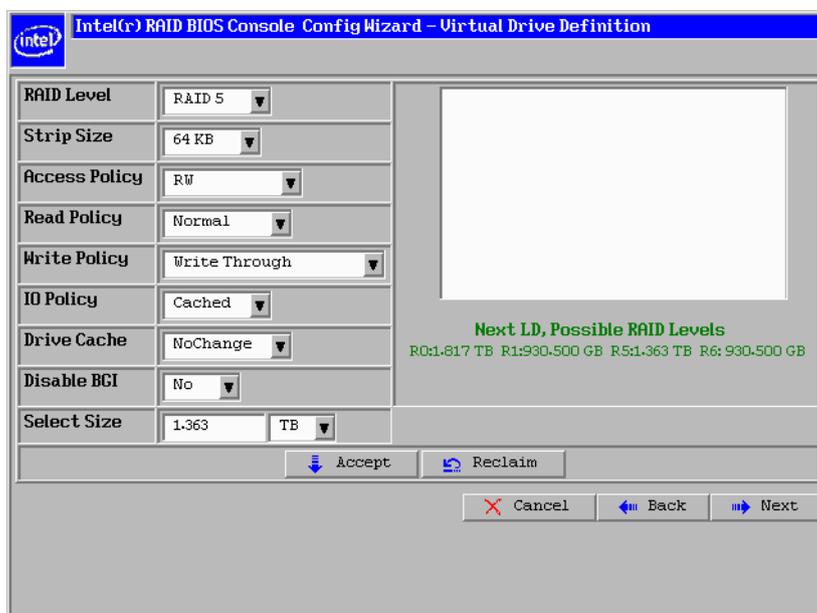


Figure 67. Intel® RAID BIOS Console 2 – Set Array Properties

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.

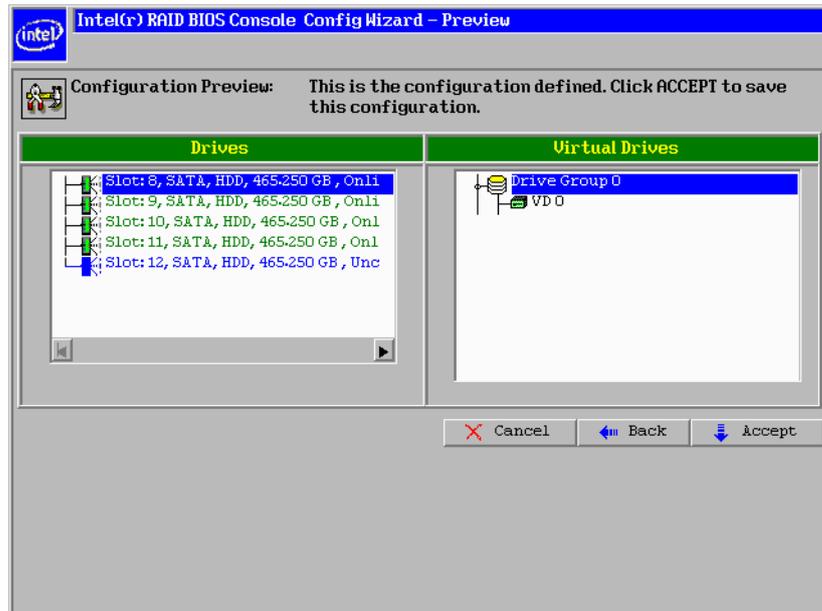


Figure 68. Intel® RAID BIOS Console 2 – Confirm Configuration

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.

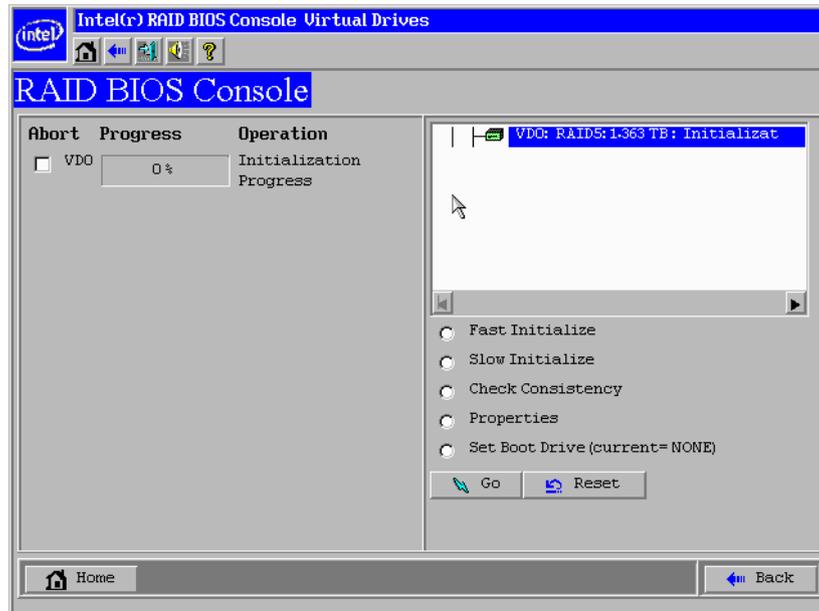


Figure 69. Intel® RAID BIOS Console 2 – Initialization Speed Setting

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.

8.5 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 65](#)).
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.

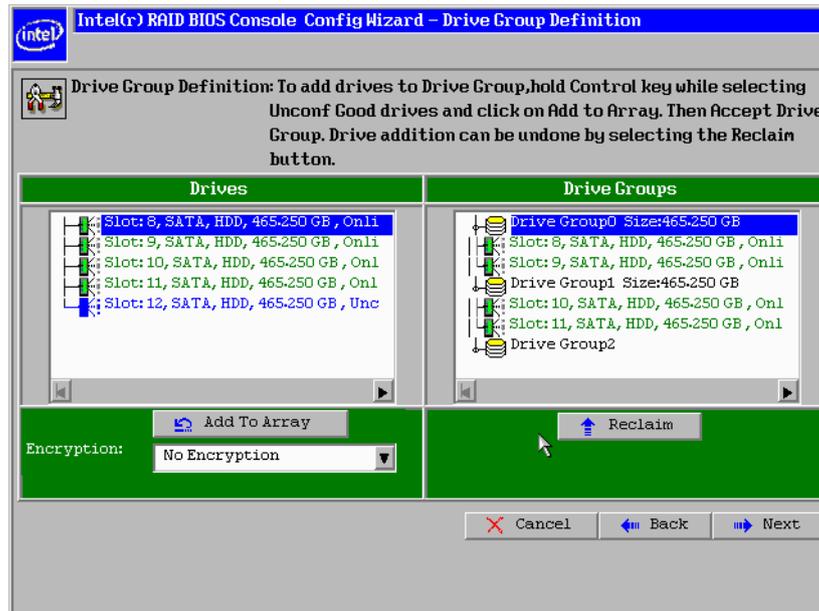


Figure 70. Intel® RAID BIOS Console 2 – Multiple Disk Groups for RAID 10, 50, or 60

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 141](#).
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.

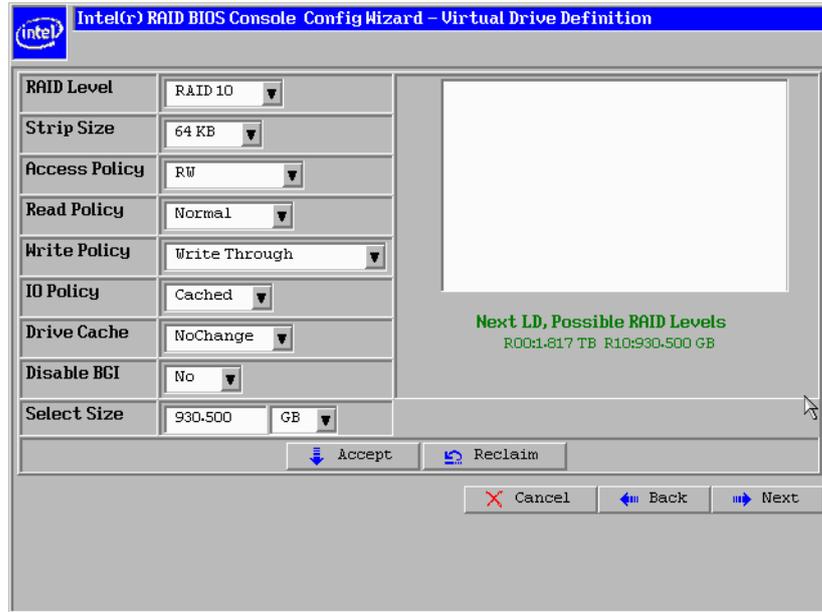


Figure 71. Intel® RAID BIOS Console 2 – Spanning Multiple Arrays

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).

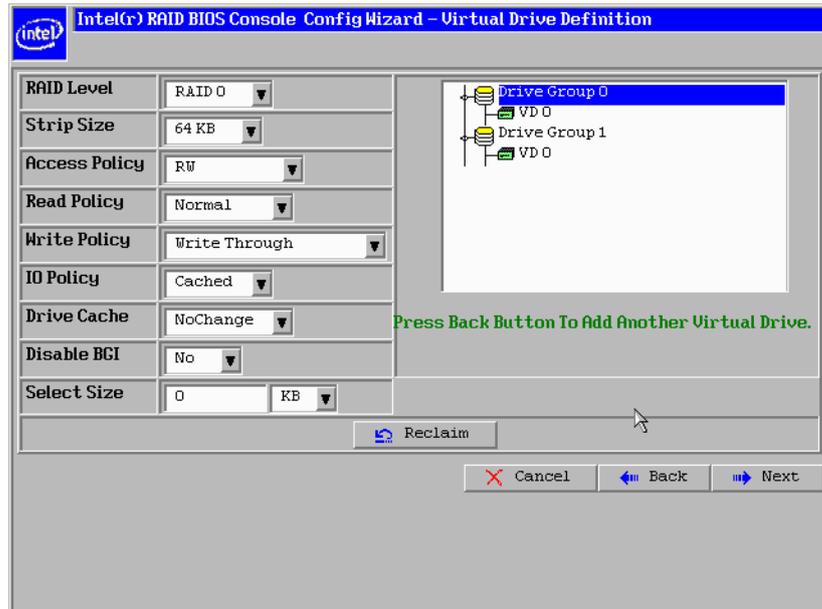


Figure 72. Intel® RAID BIOS Console 2 – Viewing Completed Settings

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.

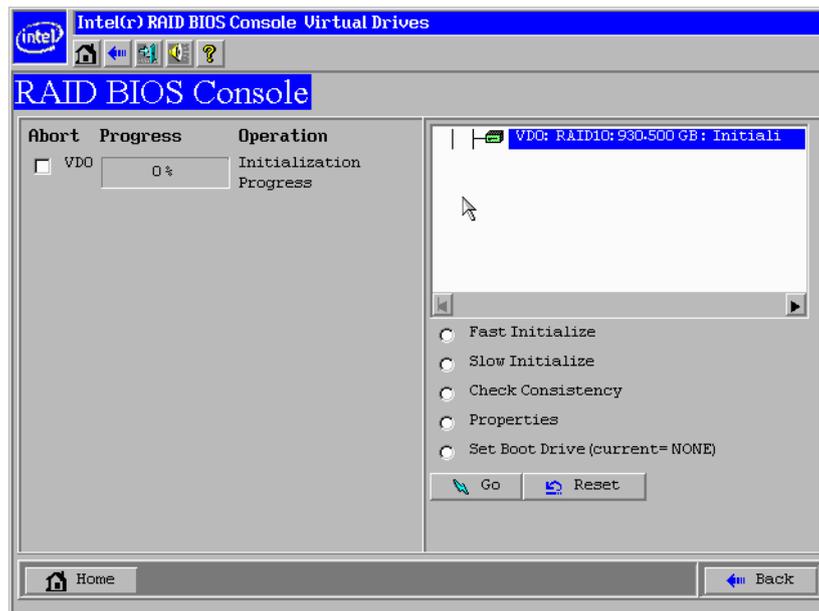


Figure 73. Intel® RAID BIOS Console 2 – Initialization Settings

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.

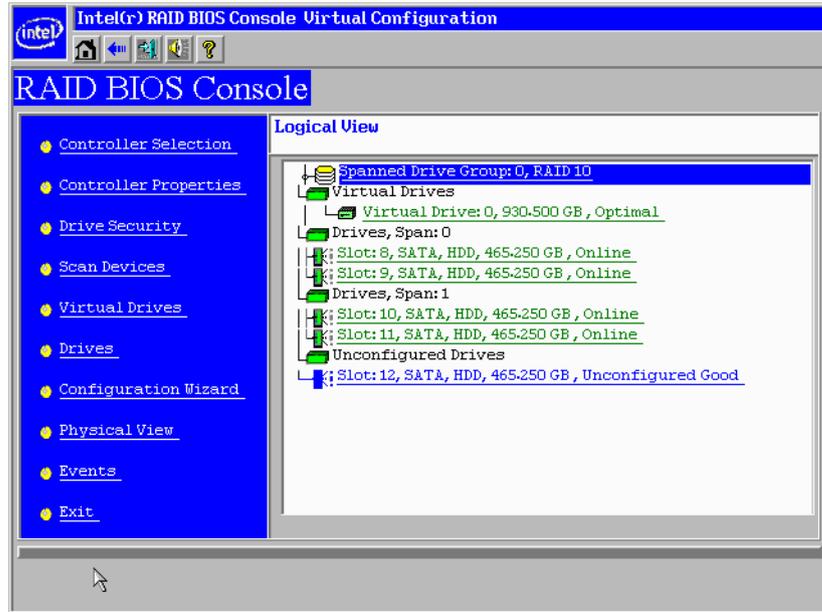


Figure 74. Intel® RAID BIOS Console 2 – RAID 10 Final Screen

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.

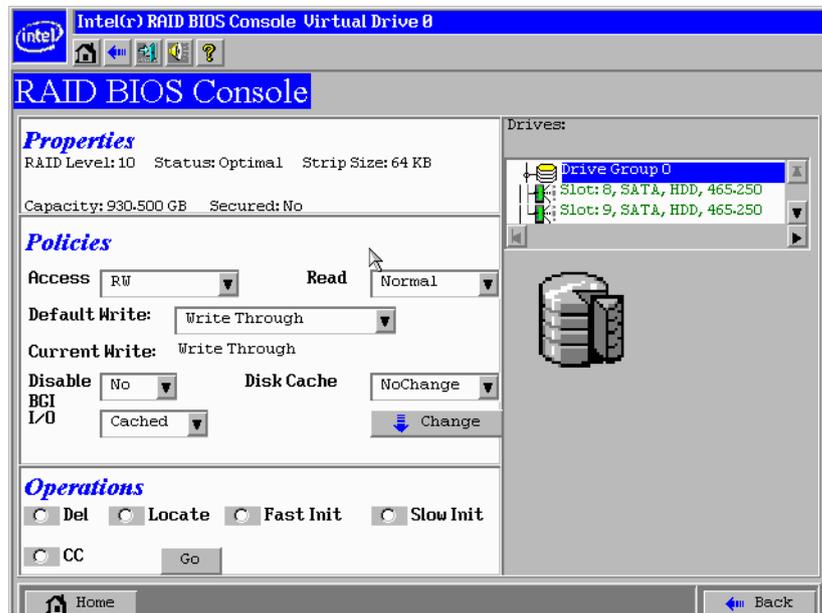


Figure 75. Intel® RAID BIOS Console 2 – RAID 10 Properties Screen

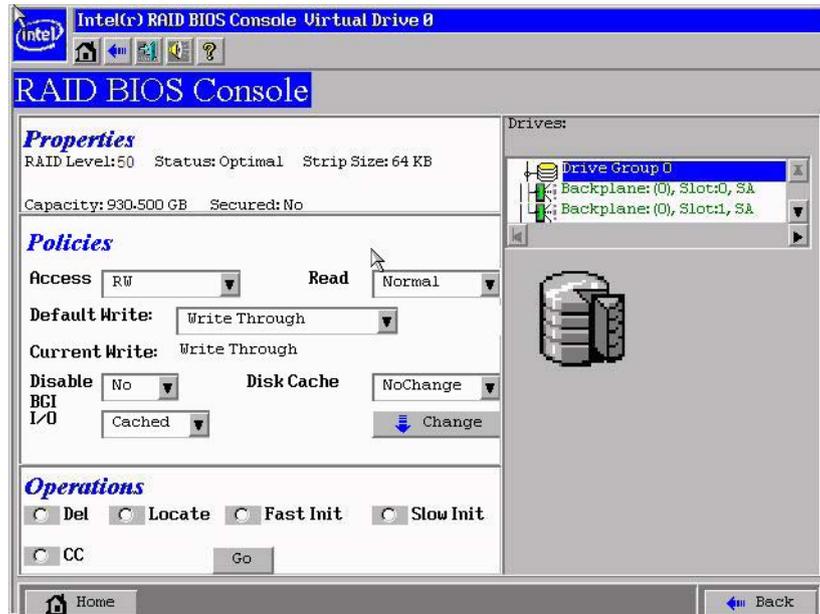


Figure 76. Intel® RAID BIOS Console 2 – RAID 50 Properties Screen

8.6 Setting Drive Parameters

The following fields are displayed in the VD Definition screen (see [Figure 67](#) and [Figure 71](#)), 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.

8.7 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.

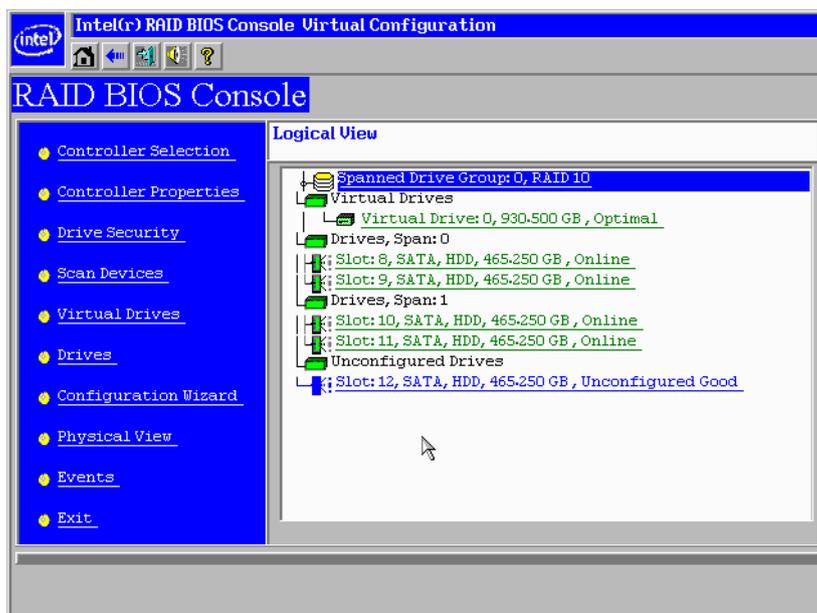


Figure 77. Intel® RAID BIOS Console 2 – Choosing a Hot Spare Drive

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.

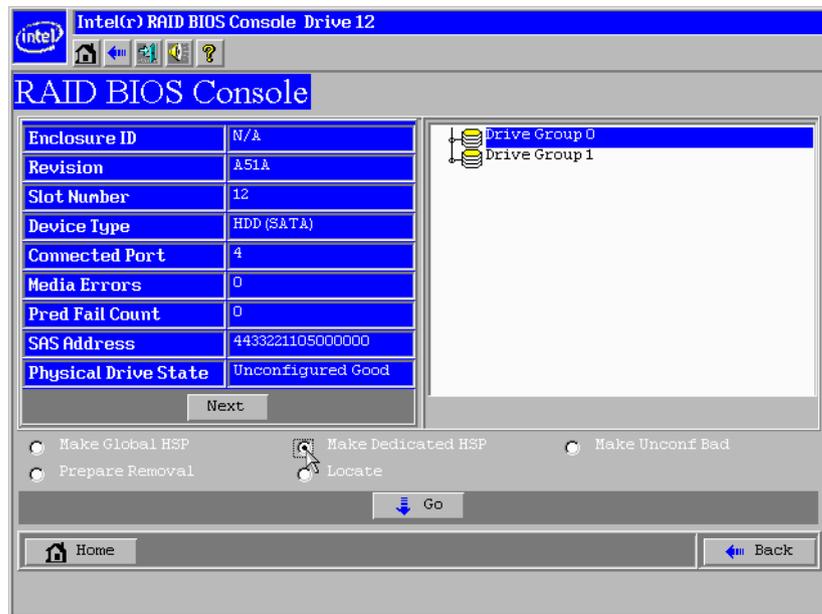


Figure 78. Intel® RAID BIOS Console 2 – Setting a Hot Spare Drive

4. Click **Go** to create the hot spare.
The Drive State changes to **HOTSPARE**, as shown below.

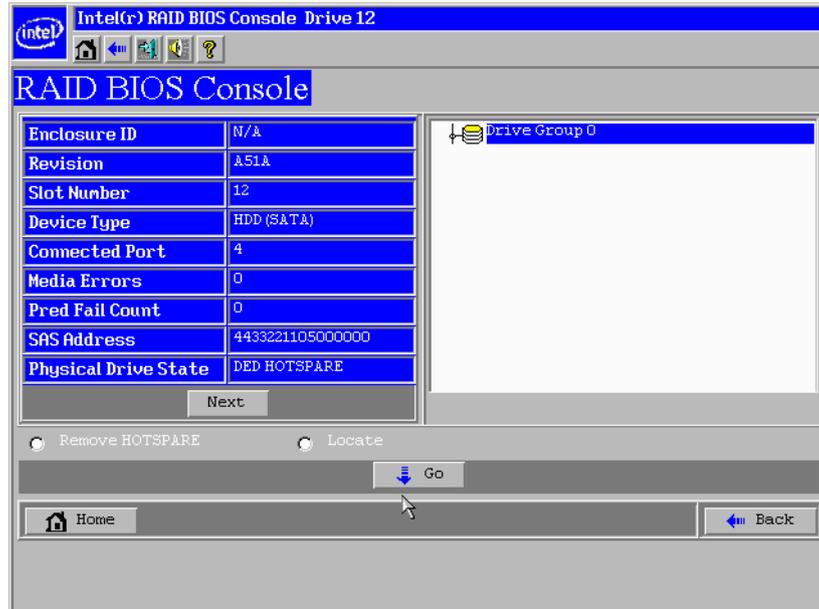


Figure 79. Intel® RAID BIOS Console 2 – Viewing Hot Spare

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

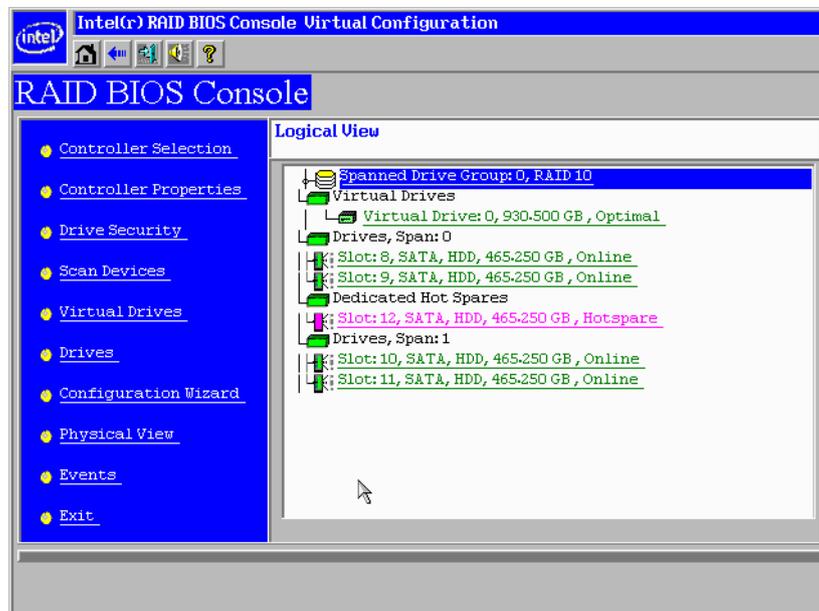


Figure 80. Intel® RAID BIOS Console 2 – Main Screen showing Hot Spare Drive

8.8 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.

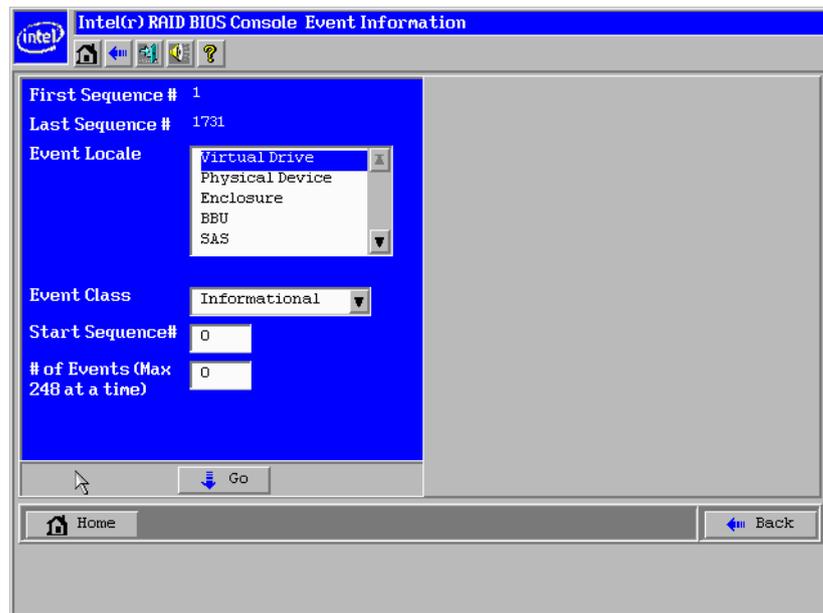


Figure 81. Intel® RAID BIOS Console 2 – Event Information Screen

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.

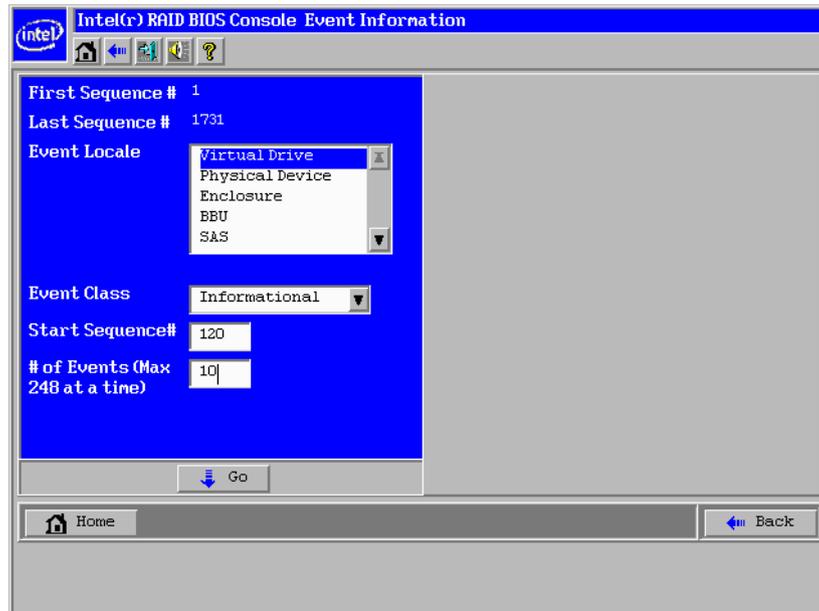


Figure 82. Intel® RAID BIOS Console 2 – Selecting Events to View

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.*

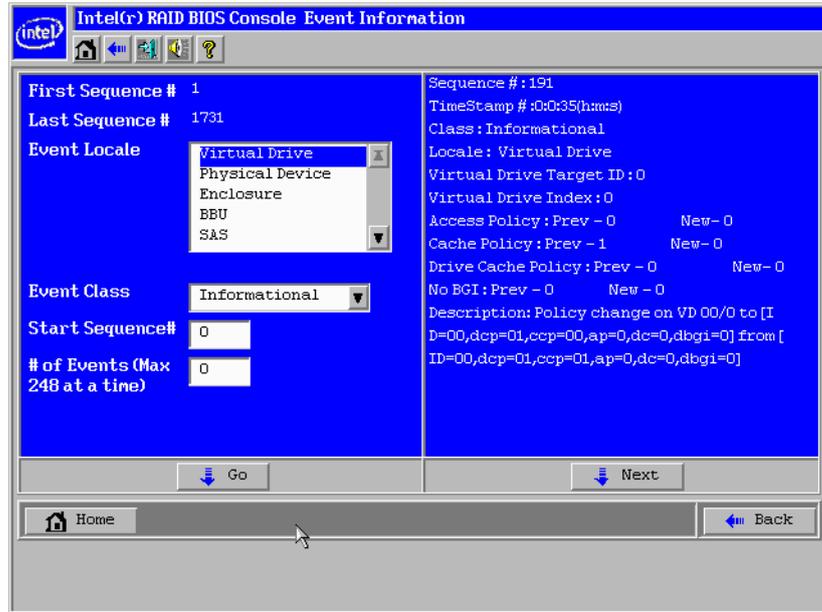


Figure 83. Intel® RAID BIOS Console 2 – Viewing an Event

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

9 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® RAID Web Console 2, the actual screen displayed could be different from this section. This section is for reference only.

9.1 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.

9.2 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.

9.3 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.

9.4 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

9.5 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 84](#).

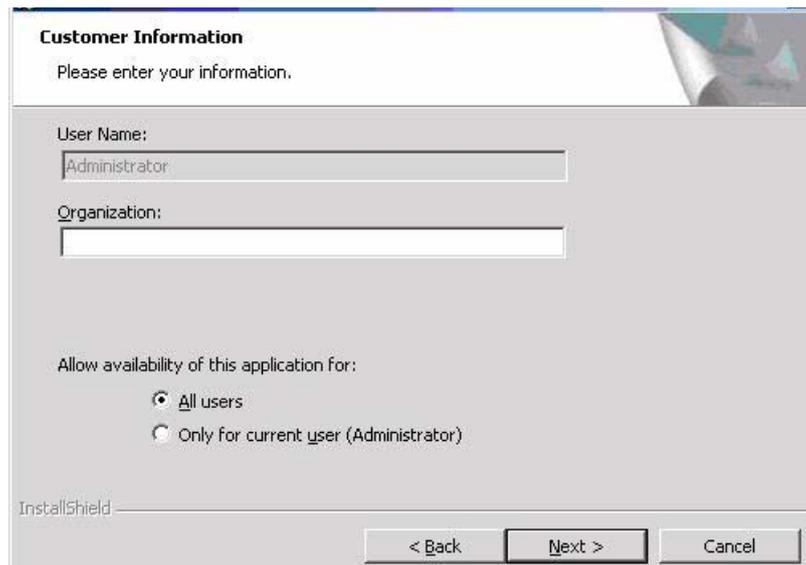


Figure 84. Intel® RAID Web Console 2 – Customer Information Screen

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 85](#).

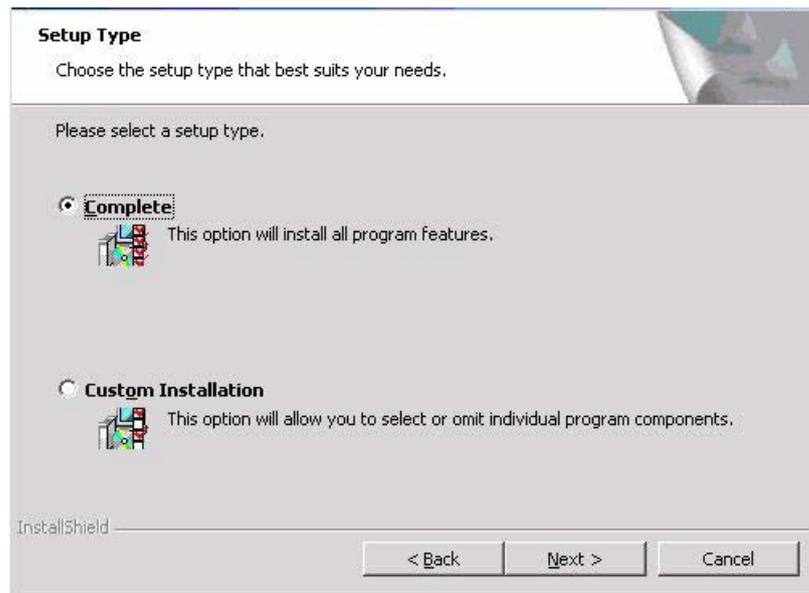


Figure 85. Setup Type Screen

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.

9.6 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.

9.7 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.

9.7.1 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.

9.7.2 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
```

9.7.3 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.

9.7.3.1 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.

9.7.3.2 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`

9.7.3.3 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.*

9.8 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**.

9.9 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.

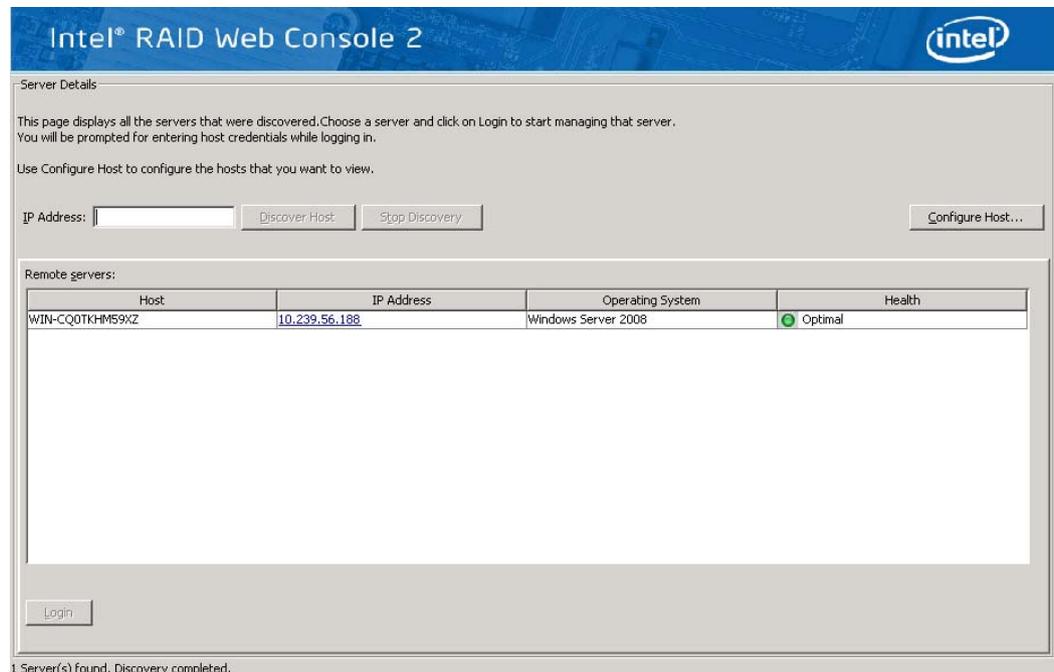


Figure 86. Intel® RAID Web Console 2 – Select Server Screen

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 87. 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.

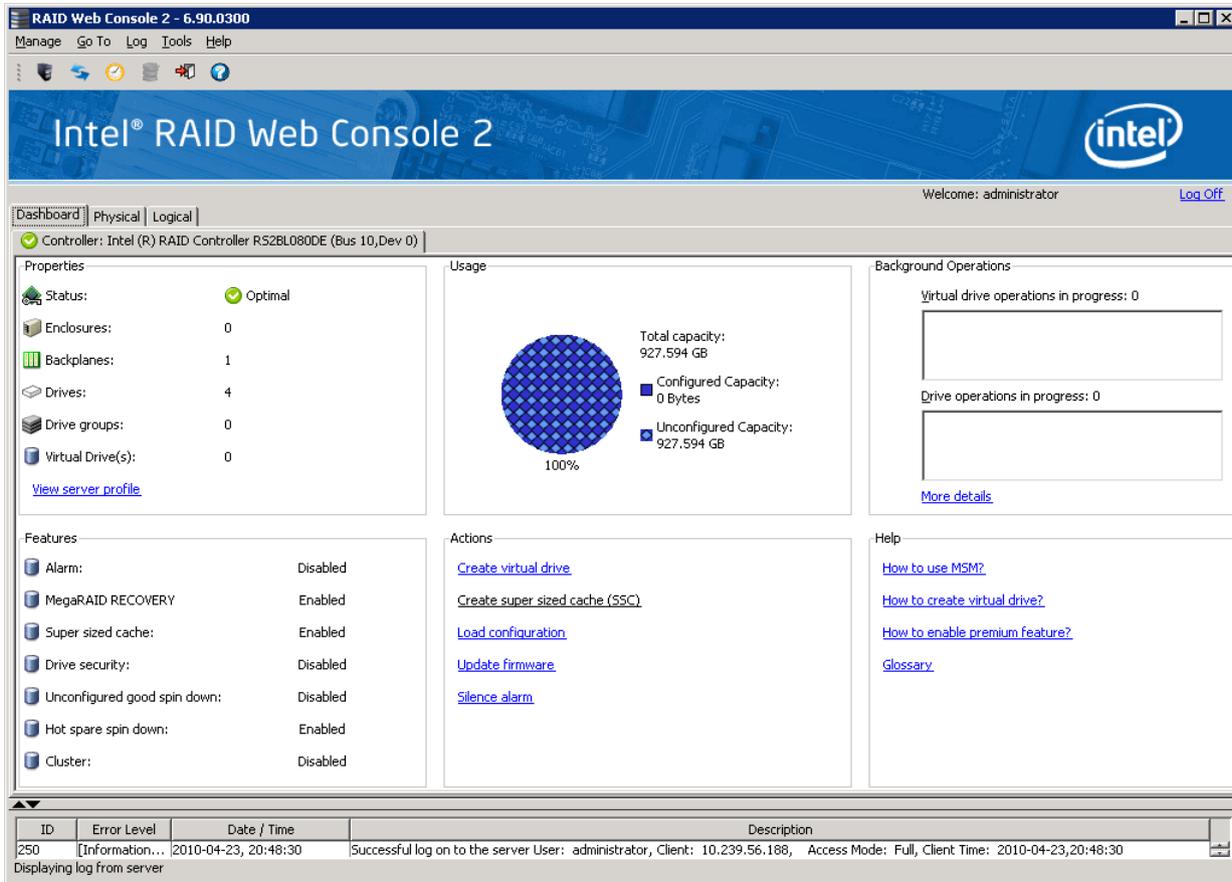


Figure 88. Intel® RAID Web console 2 dashboard

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 and so on.

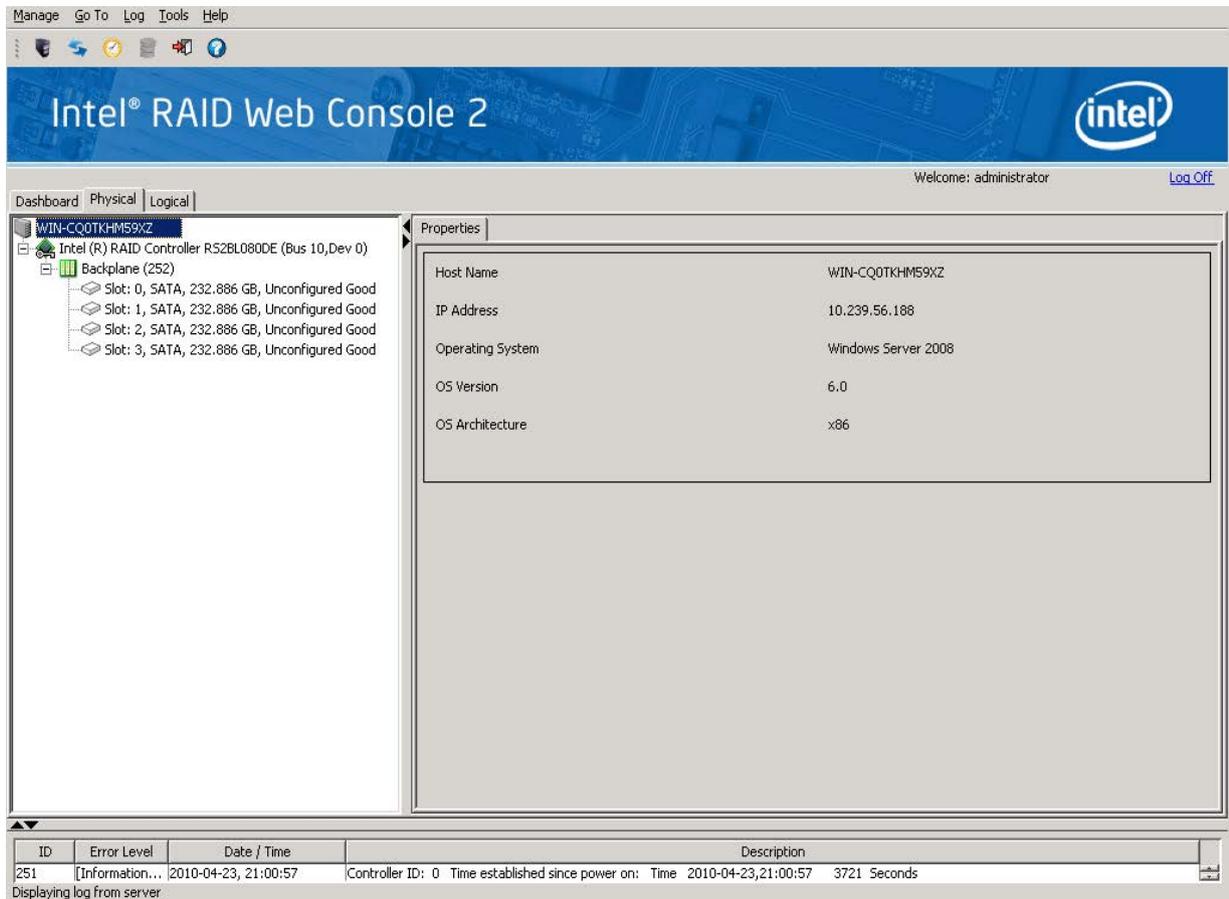


Figure 89. Intel® RAID Web Console 2 – Main Screen

The information on the screen in [Figure 89](#) shows the connected devices such as RAID controllers, physical devices, battery and so on.

The following subsections describe this screen in more detail.

9.9.1 LDAP Support

The Intel® RAID Web Console 2 application supports the discovery of remote Intel® RAID Web Console 2 servers using LDAP. To enable LDAP support, the Intel® RAID Web Console 2 servers must be registered with the LDAP server.

Note:

1. LDAP supports only Windows Active Directory LDAP Server Implementation.
2. ESXi servers are not discovered during LDAP discovery.

To register the Intel® RAID Web Console 2 servers with the LDAP server, define a new attribute on the machine on which the LDAP server is configured, and give this attribute the value MSM. This registration enables the discovery of only the Intel® RAID Web Console 2 servers that have been registered with the LDAP server.

To use LDAP support, follow these steps:

1. Double-click the Intel® RAID Web Console 2 software shortcut icon on your desktop. The Select Server dialog appears.
2. Select the **Use LDAP Login** check box, and click **Discover Host**.

All the Intel® RAID Web Console 2 servers registered with the LDAP server are displayed in the Remote servers box.

Note: If the Use LDAP Login check box is selected, the IP Address field is disabled.

3. Click on a server link to connect to the LDAP server.

Note: Based on the privileges allotted to you, the Intel® RAID Web Console 2 servers are launched with full access rights or read-only rights.

If you have selected the **Do not prompt for credentials when connecting to LDAP** check box (in the LDAP Settings tab in the Configure Host dialog), you are directly connected to the LDAP server; otherwise, the **LDAP Login** dialog appears, as shown in the figure below.



The screenshot shows a dialog box titled "LDAP Login" from the Intel RAID Web Console 2. The dialog has a blue header with the Intel logo and the text "Intel RAID Web Console 2". Below the header, there are several input fields and checkboxes. The fields are: "LDAP Server IP Address:" (disabled), "User Name:", "Password:", and "Distinguished Name:". There are checkboxes for "Use Default Port" and "Remember my Login Details". A "Port:" field is also present. At the bottom, there are "Login" and "Cancel" buttons. The text "Use your LDAP server credentials to login" is displayed above the User Name field.

Figure 90. LDAP Login

Follow these steps to enter the LDAP login details:

1. Enter the IP address of the LDAP server in the **LDAP Server IP Address** field.
2. Enter the LDAP server's user name and password in the **User Name** and **Password** fields, respectively. An example of a user name can be `username@testldap.com`.
3. Enter the name of the Domain Controller in the **Distinguished Name** field. As an example, the Domain Controller name can be `dc= TESTLDAP, dc=com`.

***Note:** The LDAP Server IP Address, User Name, Password, and Distinguished Name fields are already populated if their corresponding values have been stored in the LDAP Settings tab in the Configure Host dialog.*

4. Perform one of these actions:
 - If you want to use the default port number, select the **Use Default Port** check box. The default port number, 389, appears in the **Port** field.
 - If you do not want to use the default port number, uncheck the **Use Default Port** check box, and enter a port number in the **Port** field.
5. Select the **Remember my Login Details** check box if you want to save all the values entered in this dialog in the **LDAP Settings** tab in the **Configure Host** dialog.
6. Click **Login** to log in to the LDAP server.

9.9.2 Configuring LDAP Support Settings

To configure settings for LDAP support, follow these steps:

1. Navigate to the **Configure Host** dialog and click the **LDAP Settings** tab. The following fields appear:

RAID Web Console 2 12.01.03.01 - Configure Host

Intel® RAID Web Console 2

Configure Host **LDAP Settings**

Use LDAP login as default login mode

Do not prompt for credentials when connecting to LDAP

LDAP Server

IP Address: Port:

Distinguished Name:

Connection

User Name:

Password:

Save Settings Cancel

Figure 91. Configure Host LDAP

2. Select the **Use LDAP Login as default login** mode check box to always connect to the LDAP server.
3. Select the **Do not prompt for credentials when connecting to LDAP** check box if you do not want the LDAP Login dialog to appear when connecting to the LDAP server.
4. Enter the IP address of the LDAP server in the **IP Address** field.
5. Enter the port number in the **Port** field.
6. Enter the name of the Domain Controller in the **Distinguished Name** field.
7. Enter the user name and password for logging into the LDAP server in the **User Name** and **Password** fields, respectively.
8. Click **Save Settings** to save all the values entered in the fields in the msm.properties file.

9.9.3 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' may be 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: 

9.9.4 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® 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.*

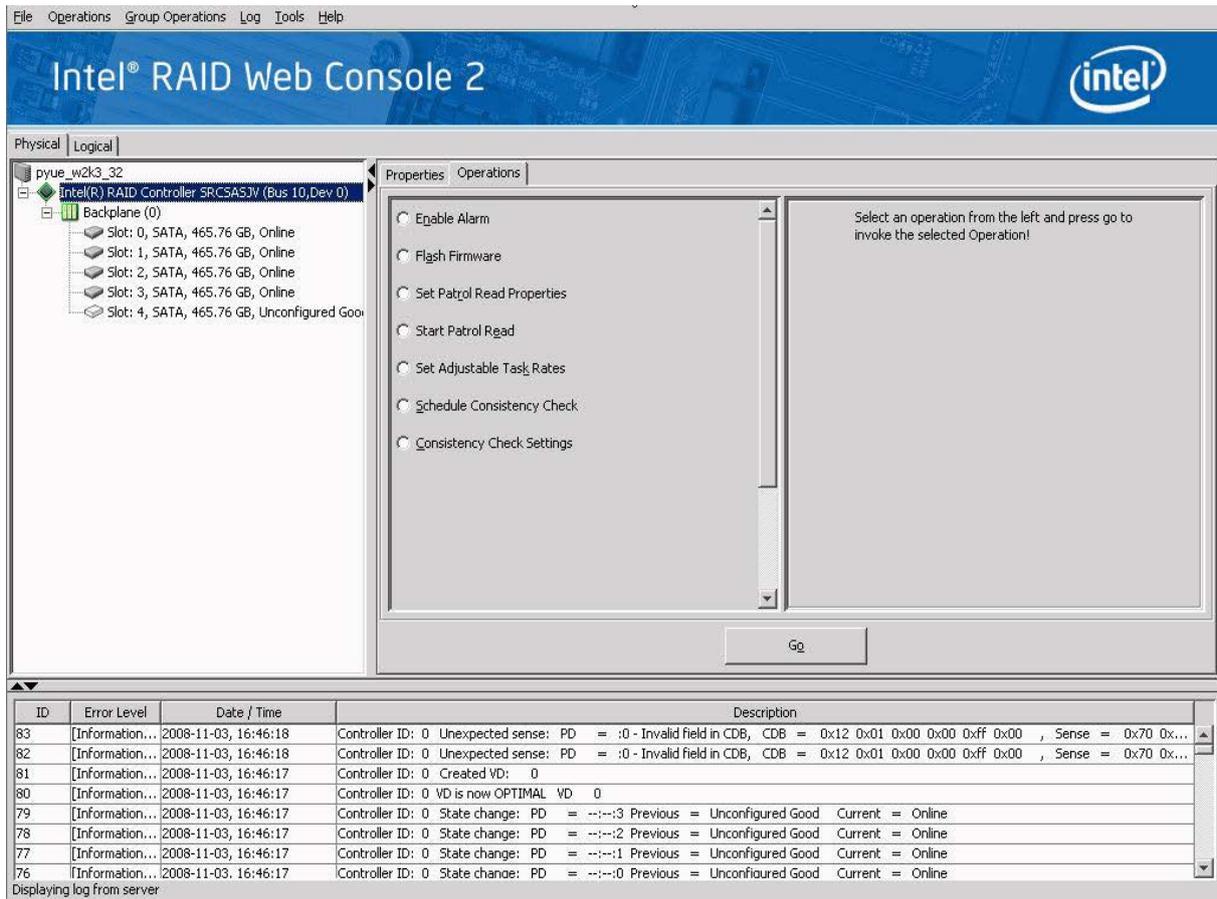


Figure 92. Intel® RAID Web Console 2 – Operations Tab

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..

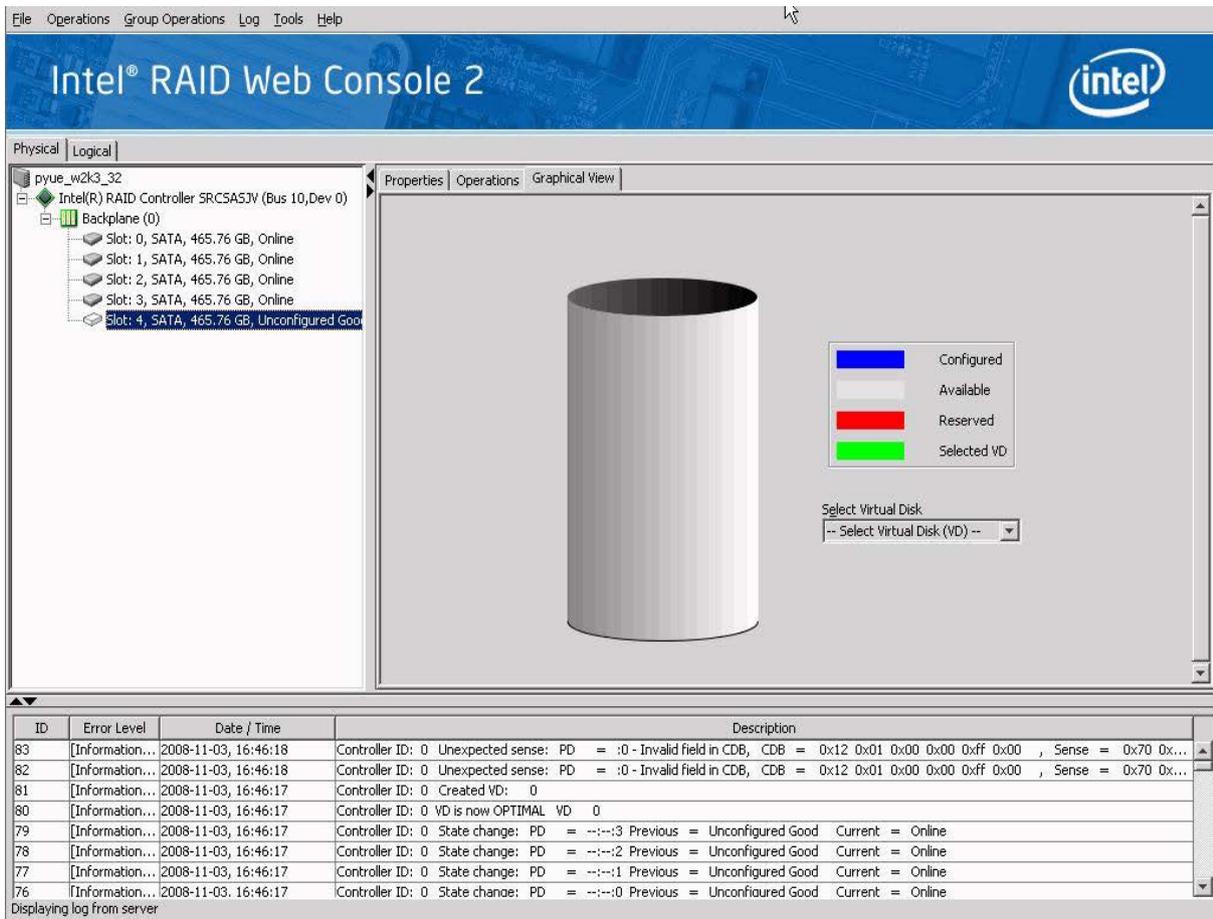


Figure 93. Intel® RAID Web Console 2 – Graphical Tab (Optional feature)

9.9.5 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 315](#)

9.9.6 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.

9.9.7 Menu Bar/Go To Menu

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

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.

9.9.8 File Menu/Log Menu

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

9.9.9 File Menu/Tool Menu

The Tools menu includes options for configuring alerts.

9.9.10 File Menu/Help Menu

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

9.10 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 169](#)
- [“Creating a Spanned Disk Group” on page 183](#)
- [“Creating Hot Spares” on page 184](#)
- [“Setting Adjustable Task Rates” on page 186](#)
- [“Adding a Drive to a Virtual Disk” on page 188](#)

- “Removing a Drive or Changing the RAID Level of a Virtual Disk” on page 191
- “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 193
- “Changing Virtual Disk Properties” on page 193
- “Deleting a Virtual Disk” on page 194
- “Managing Configurations” on page 195

9.10.1 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.*

9.11 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 174.*

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 94](#).

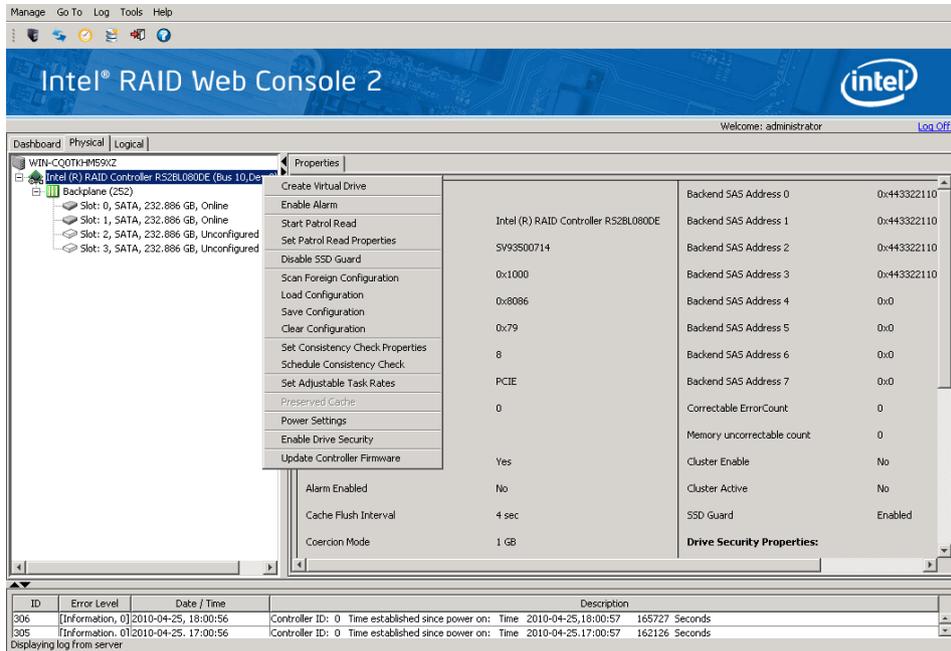


Figure 94. Virtual Drive Creation Menu

The dialog box for the configuration mode (simple or advanced) appears, as shown in [Figure 95](#).

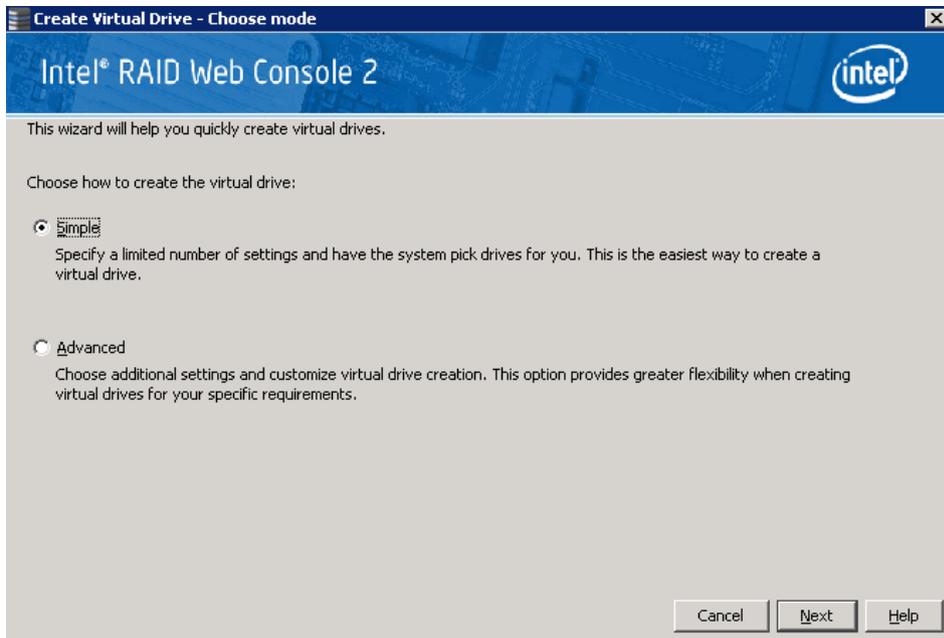


Figure 95. Virtual Drive Creation Mode

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

The Create Virtual Drive screen appears, as shown in [Figure 96](#).

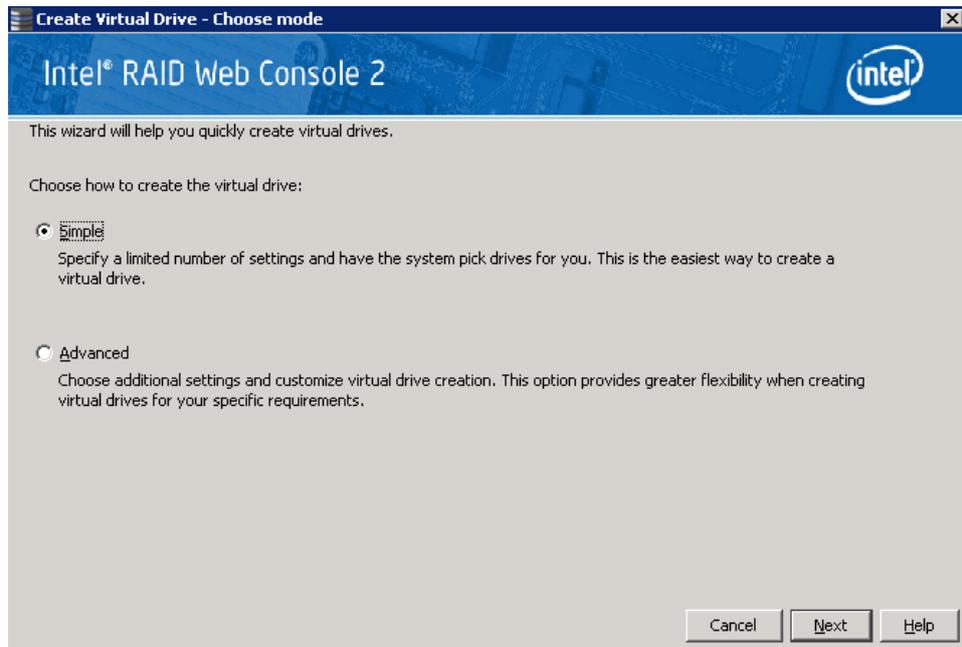


Figure 96. Create Virtual Drive Screen

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 97](#). This window shows the selections you made for simple configuration.

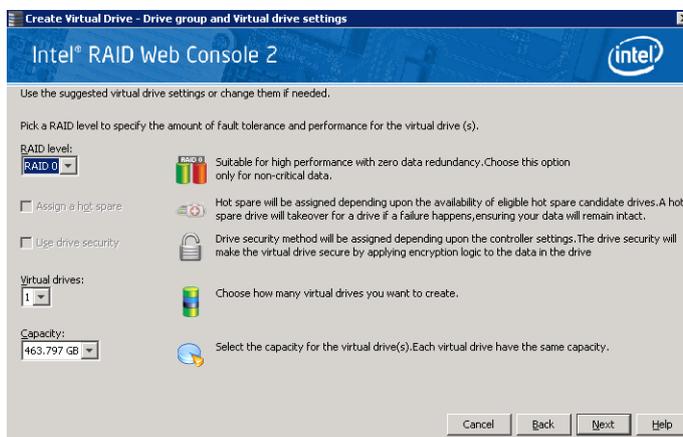


Figure 97. Create Virtual Drive – Summary Window

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 98](#). 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.

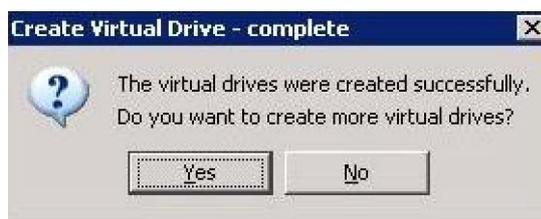


Figure 98. Option to Create Additional Virtual Drives

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 96](#). If you select **No**, the utility asks whether you want to close the wizard, as shown in [Figure 99](#).
10. If you selected **No** in step 9, select **Yes** or **No** to indicate whether to close the wizard.



Figure 99. 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.

9.12 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 100](#)

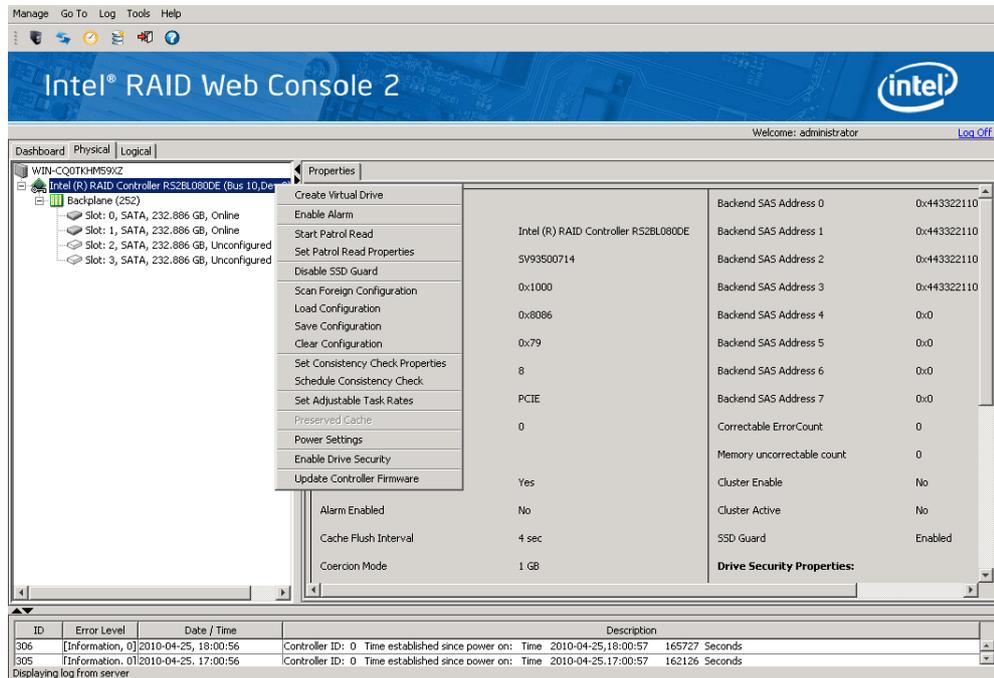


Figure 100. Virtual Drive Creation Menu

The dialog box shown in Figure 101 appears.

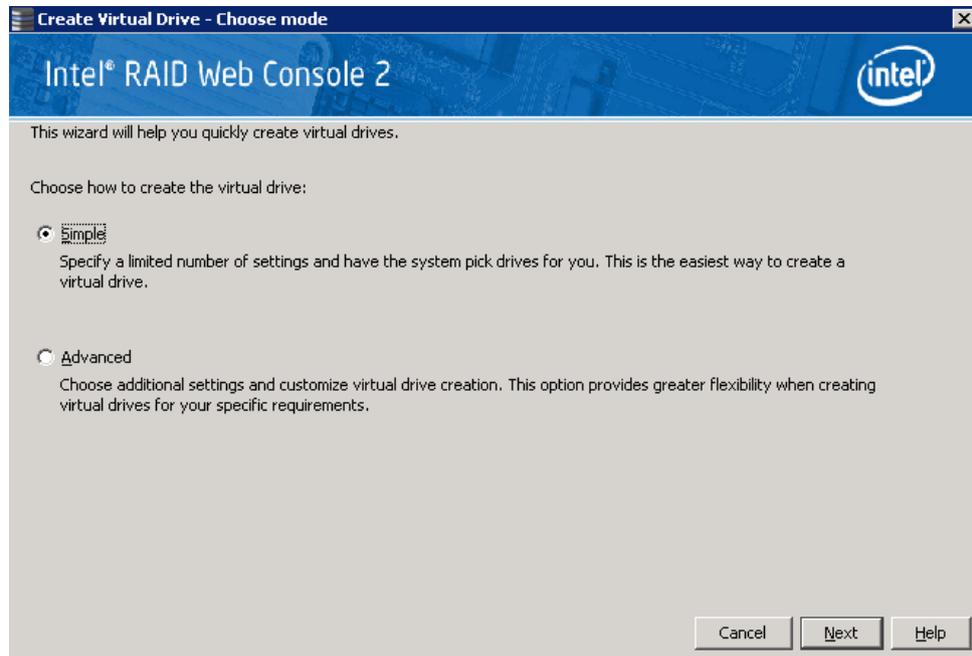


Figure 101. Virtual Drive Creation Mode

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

The Create Drive Group Settings screen appears, as shown in [Figure 102](#).

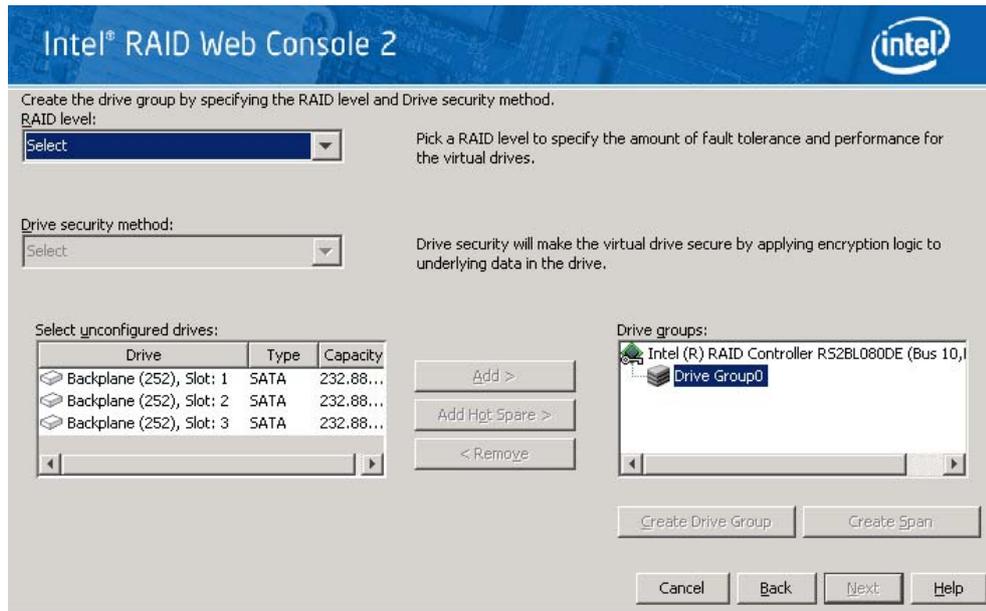


Figure 102. Create Drive Group Settings Screen

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 103](#).

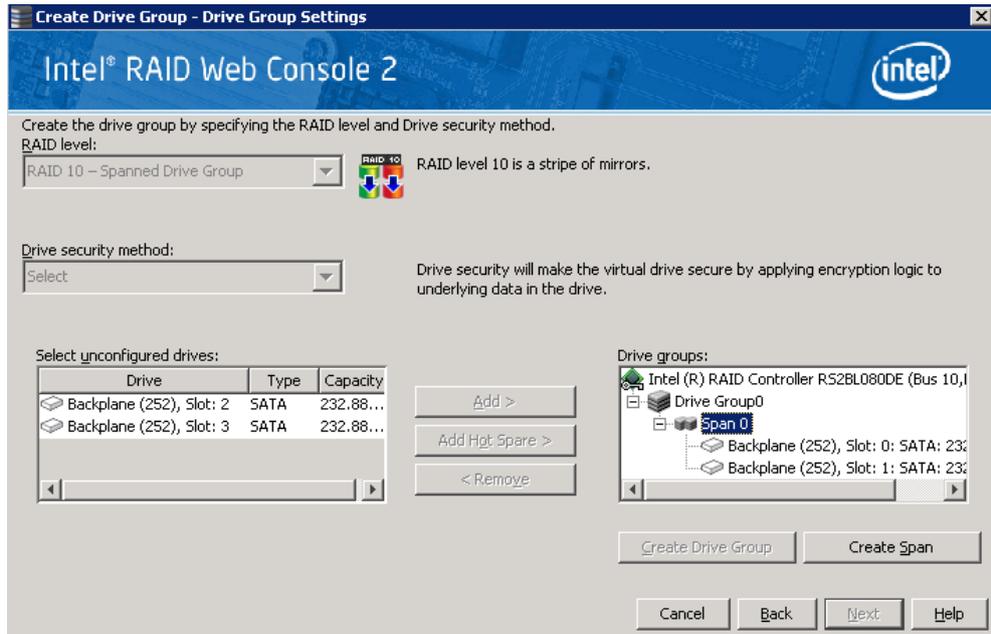


Figure 103. Span 0 of Drive Group 0

- 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 104](#)

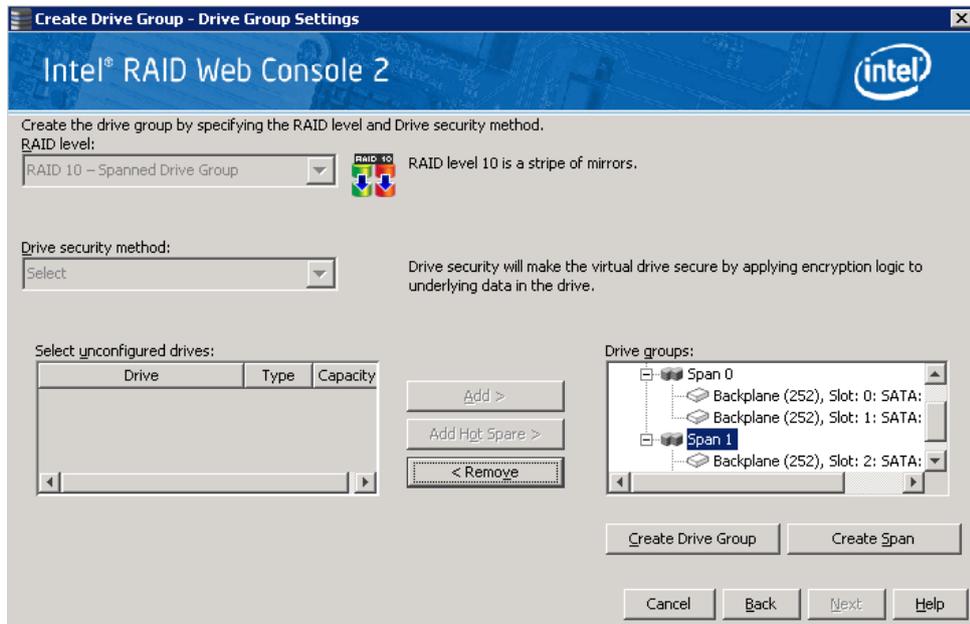


Figure 104. Span 0 and Span 1 of Drive Group 0

- 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 105](#). 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.

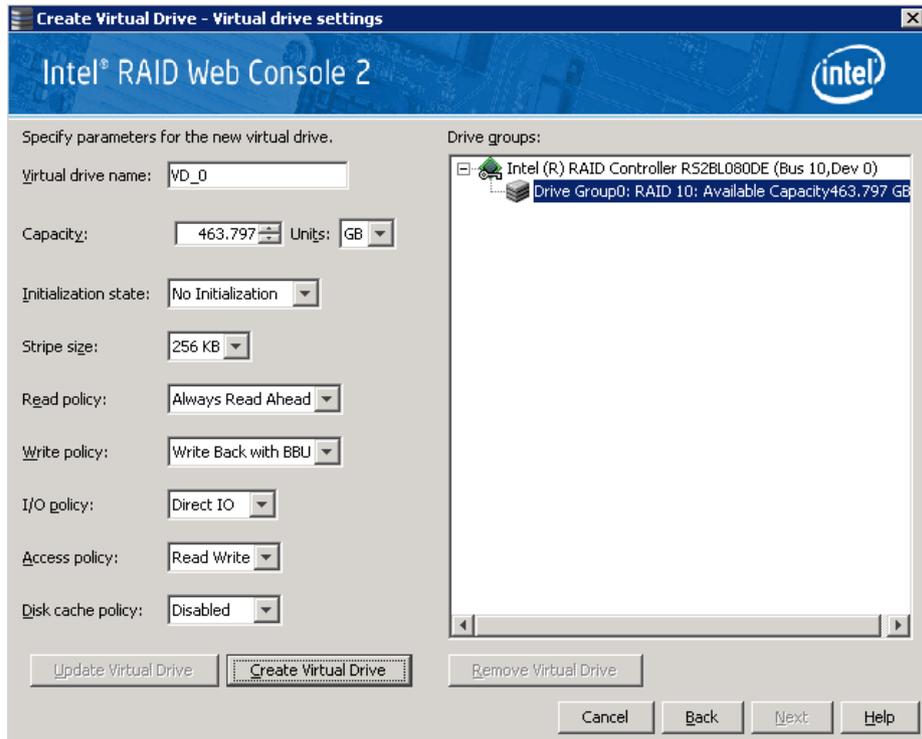


Figure 105. Virtual Drive Settings Window

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 106](#). 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.

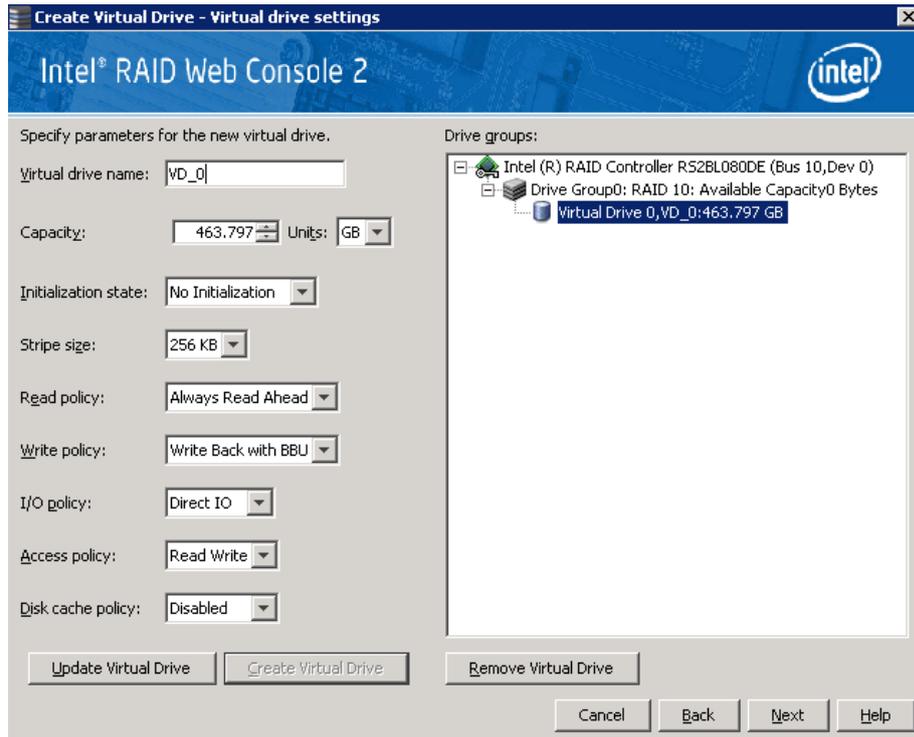


Figure 106. New Virtual Drive 0

6. Click **Next**.

The **Create Virtual Drive - Summary** window appears, as shown in [Figure 107](#). This window shows the selections you made for advanced configuration.

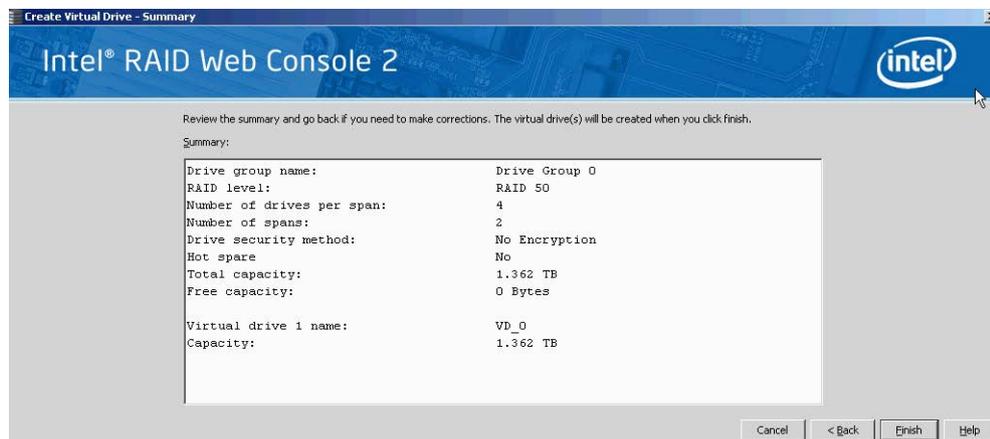


Figure 107. Create Virtual Drive Summary Window

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 108](#). 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.

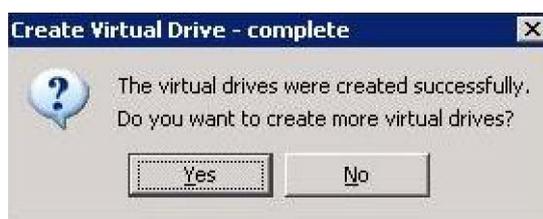


Figure 108. Option to Create Additional Virtual Drives

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 96](#). If you select **No**, the utility asks whether you want to close the wizard, as [Figure 109](#).
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 109. Option to Close the Configuration Wizard

9.13 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”](#) on page 174.

9.14 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 Hotspare**, as shown in [Figure 110](#).

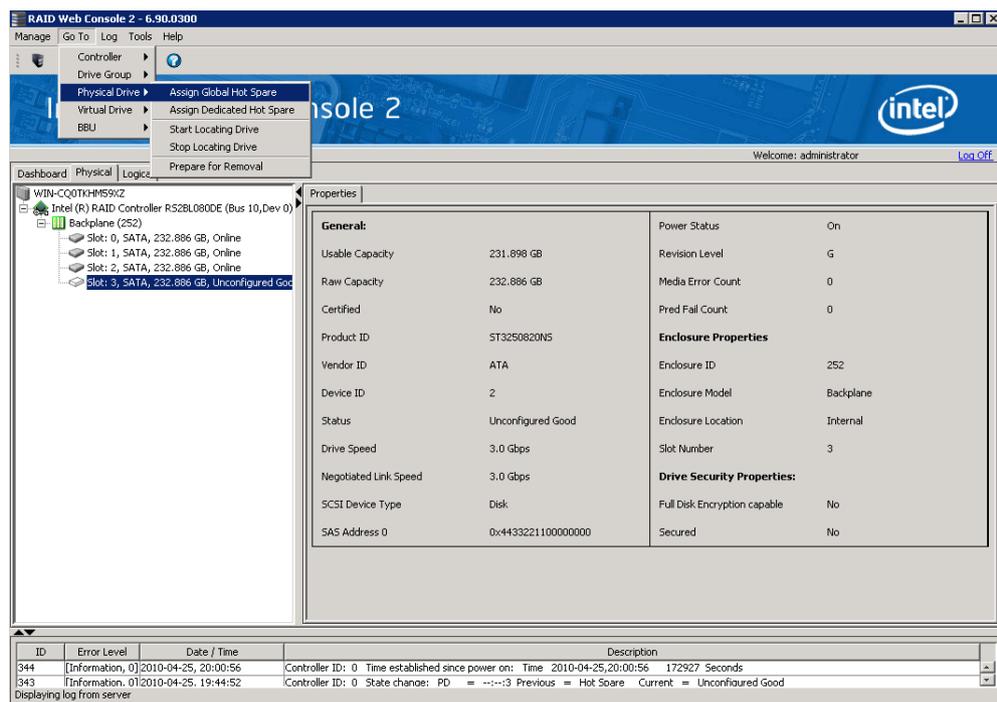


Figure 110. Assign Global Hotspare

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 111](#).

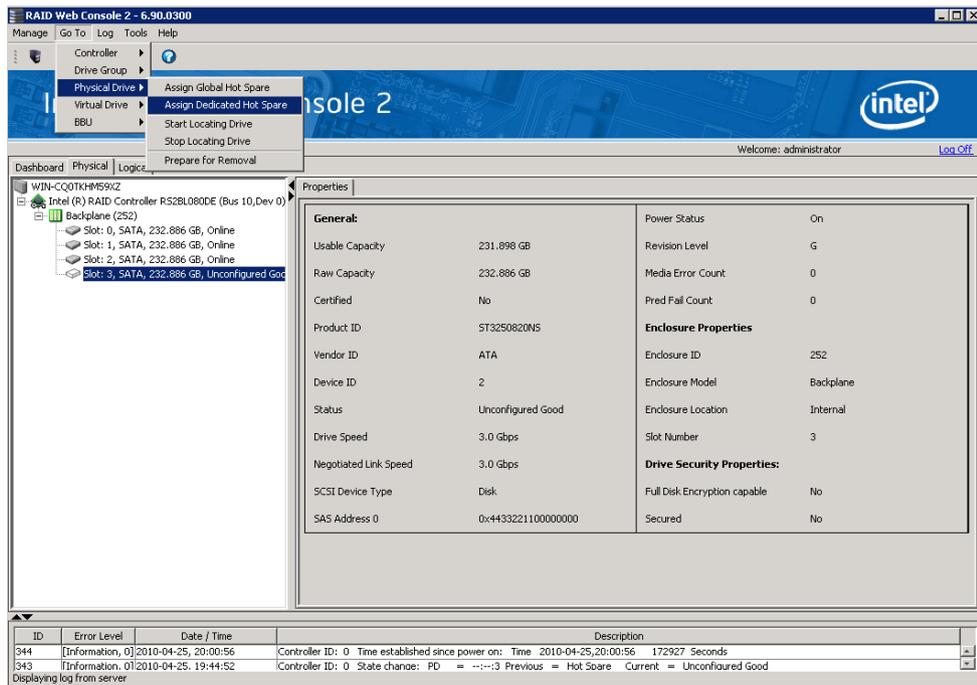


Figure 111. Assign Dedicated Hotspare

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

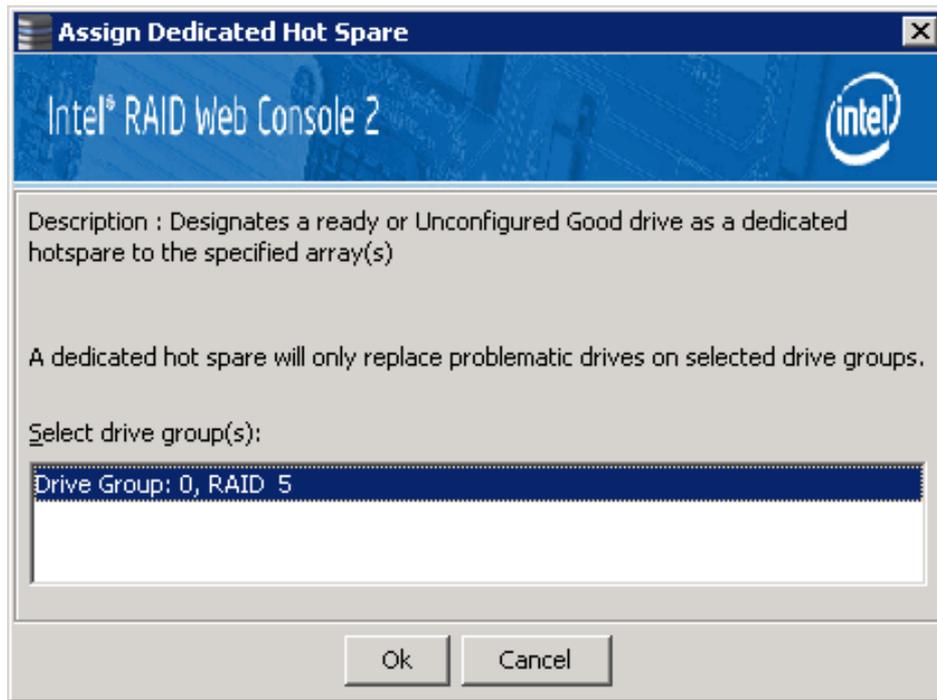


Figure 112. Select Hotspare Drive

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

9.15 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 113](#).

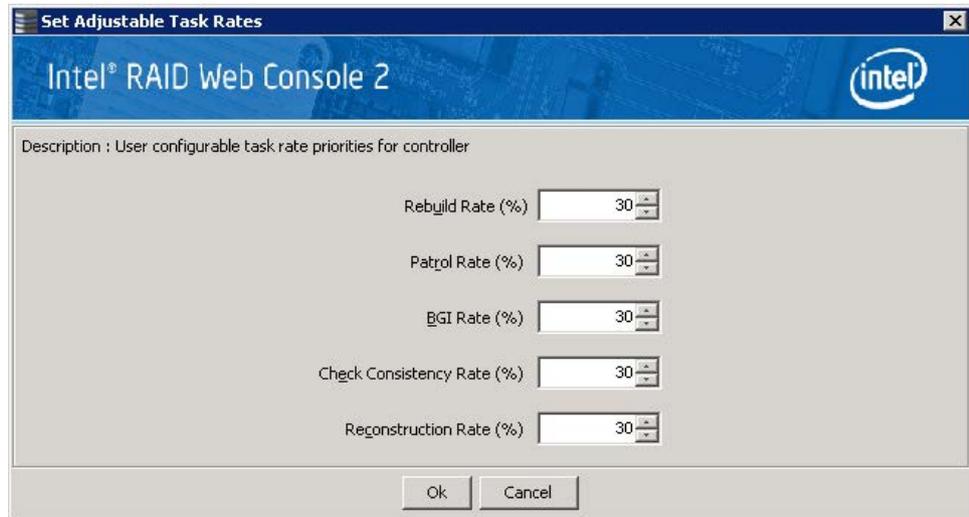


Figure 113. Set Adjustable Task Rates

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.*

9.16 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.**

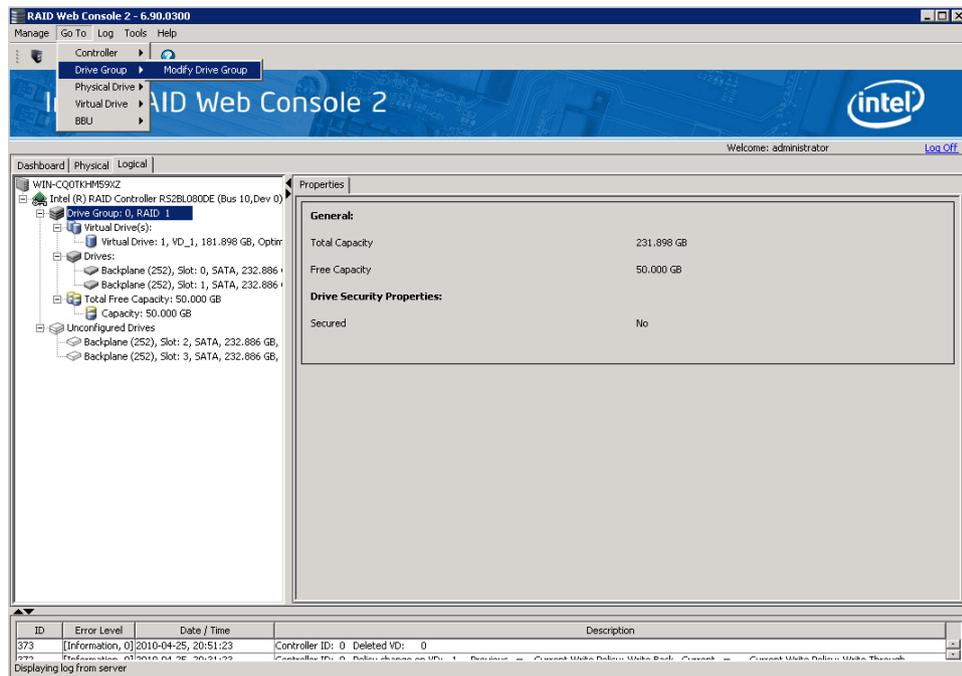


Figure 114. Starting Modify Drive Group

The Modify Drive Group window is displayed, as shown in Figure 115.

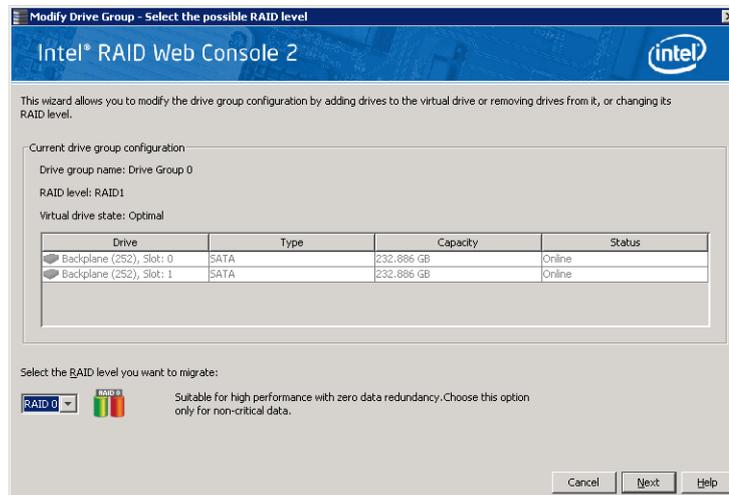


Figure 115. Select RAID level to migrate

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:

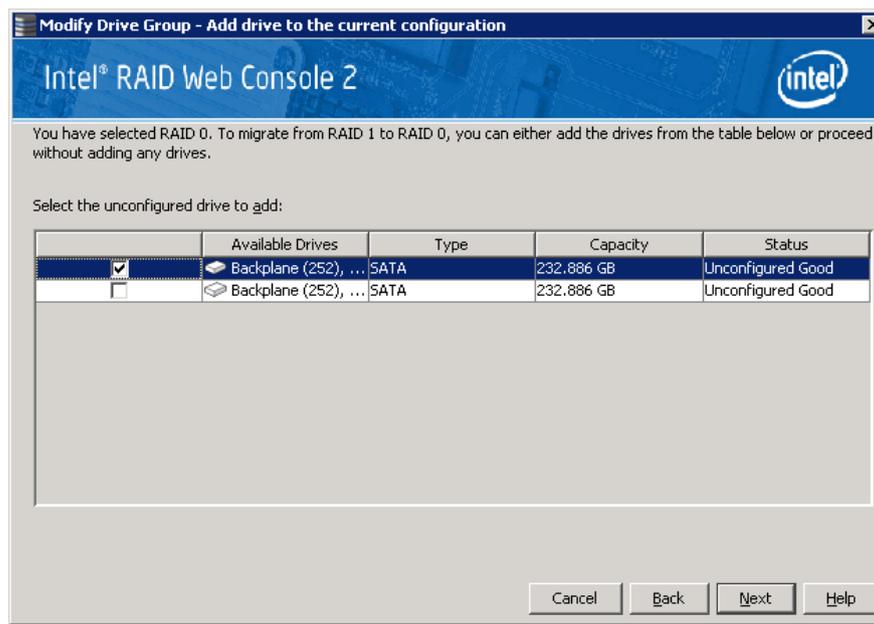


Figure 116. Selecting Drives to Add

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 117](#).

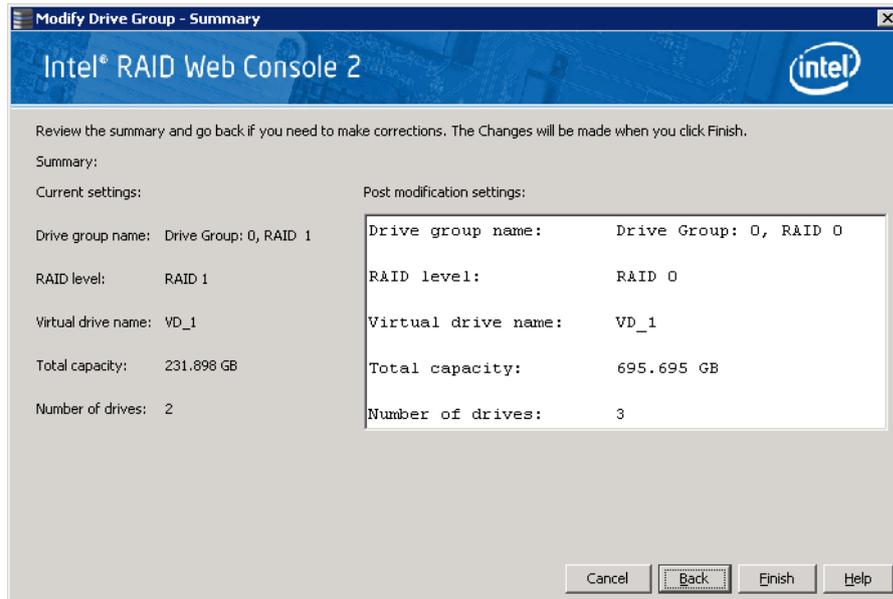


Figure 117. Changing RAID Level

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**.

9.17 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:

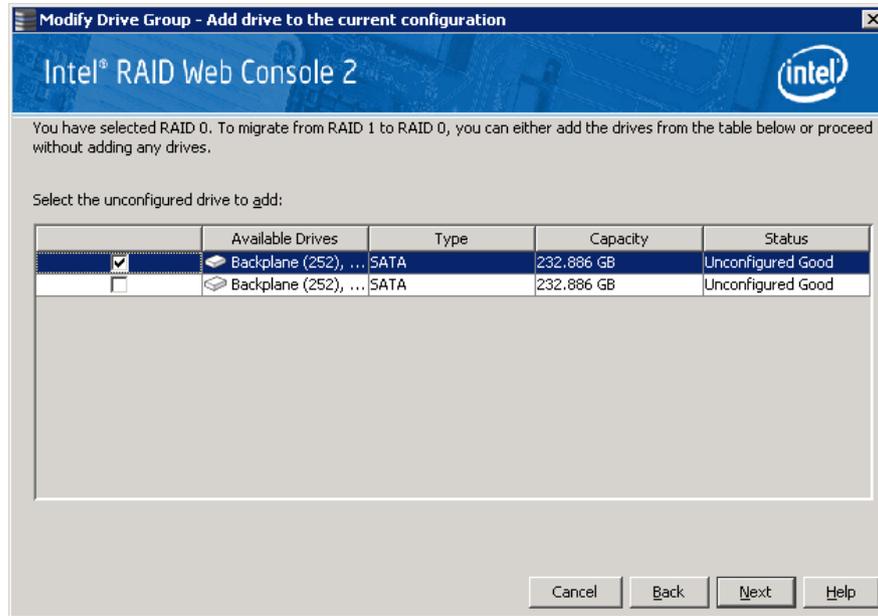


Figure 118. Selecting Drives to Remove

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 119](#).

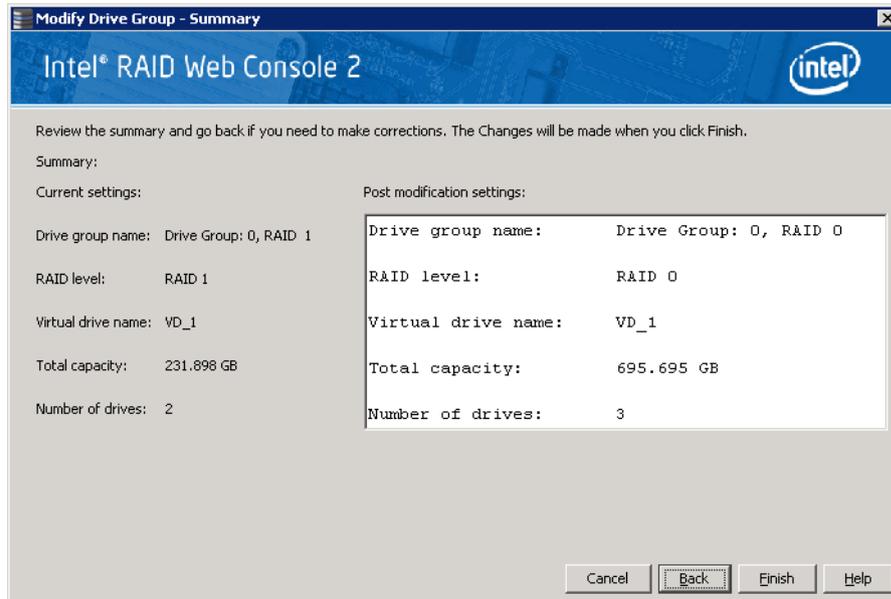


Figure 119. Changing RAID Level

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**.

9.18 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 120](#) shows the Set Virtual Disk Properties screen.

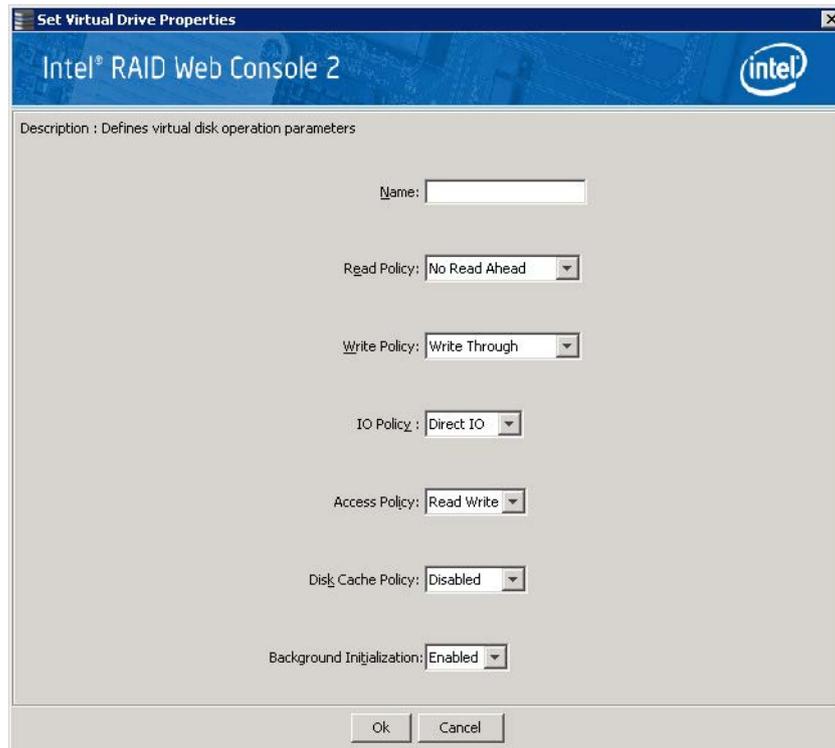


Figure 120. Set Virtual Disk Properties

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

9.19 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.

9.20 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.

9.20.1 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**.

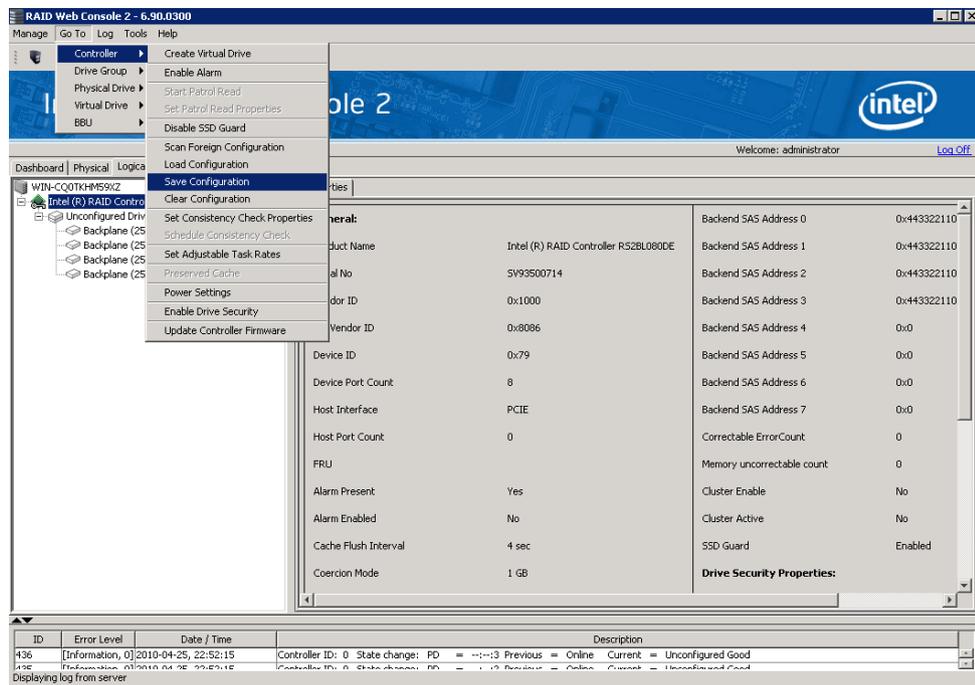


Figure 121. Save Configuration to File

The Save dialog box displays as shown in [Figure 122](#).



Figure 122. 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

9.20.2 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**.

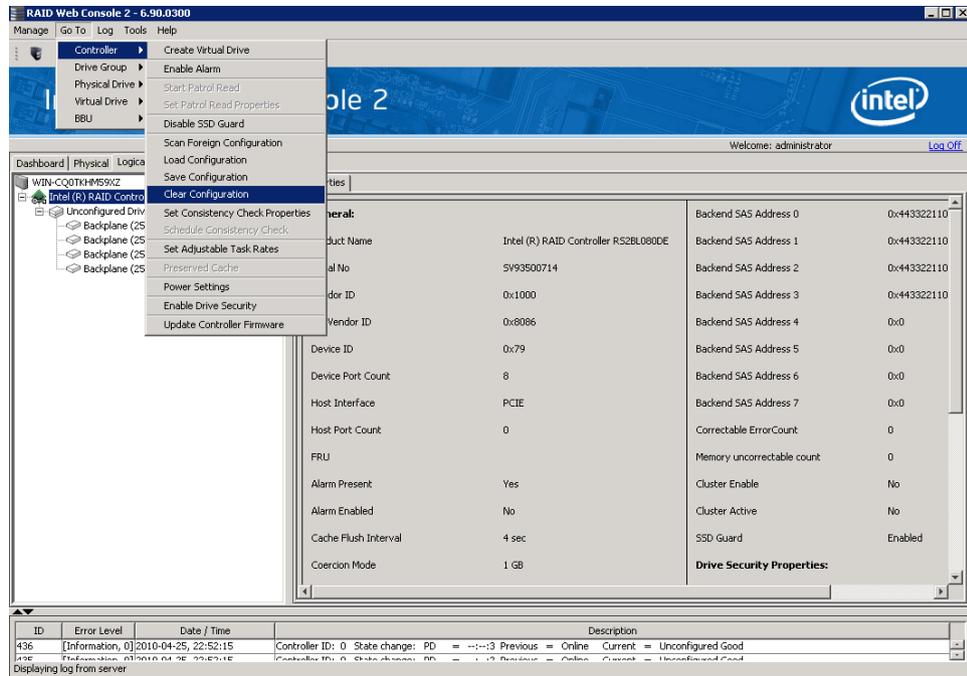


Figure 123. Clear Configuration

A Warning message displays.

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

9.20.3 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**.

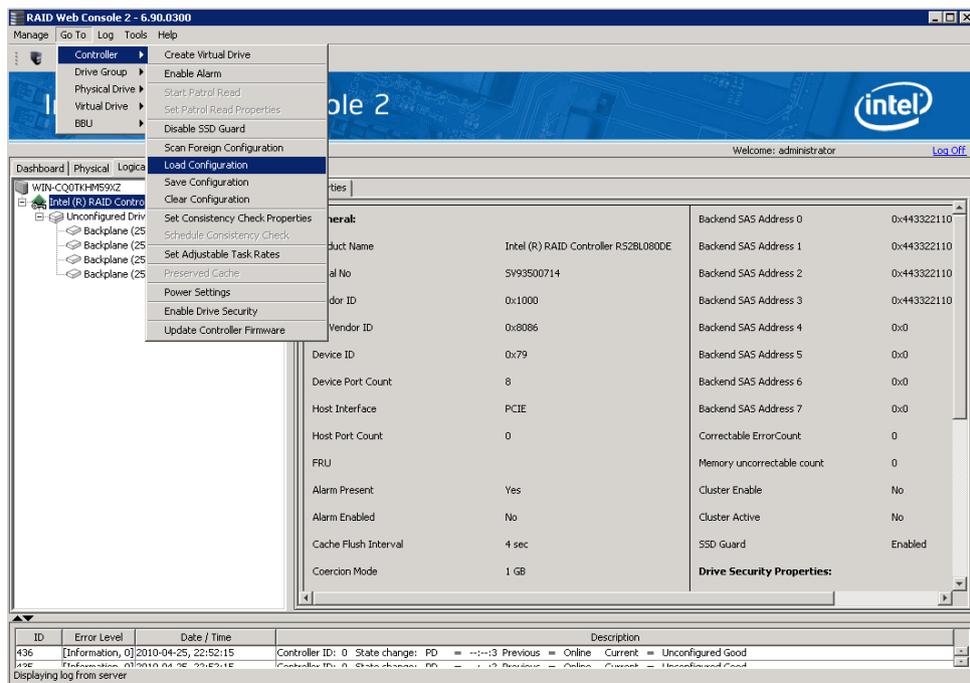


Figure 124. Add Saved Configuration

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.

9.21 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

9.21.1 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 125](#). These event messages also display in the Microsoft Windows* application log (the Event Viewer).

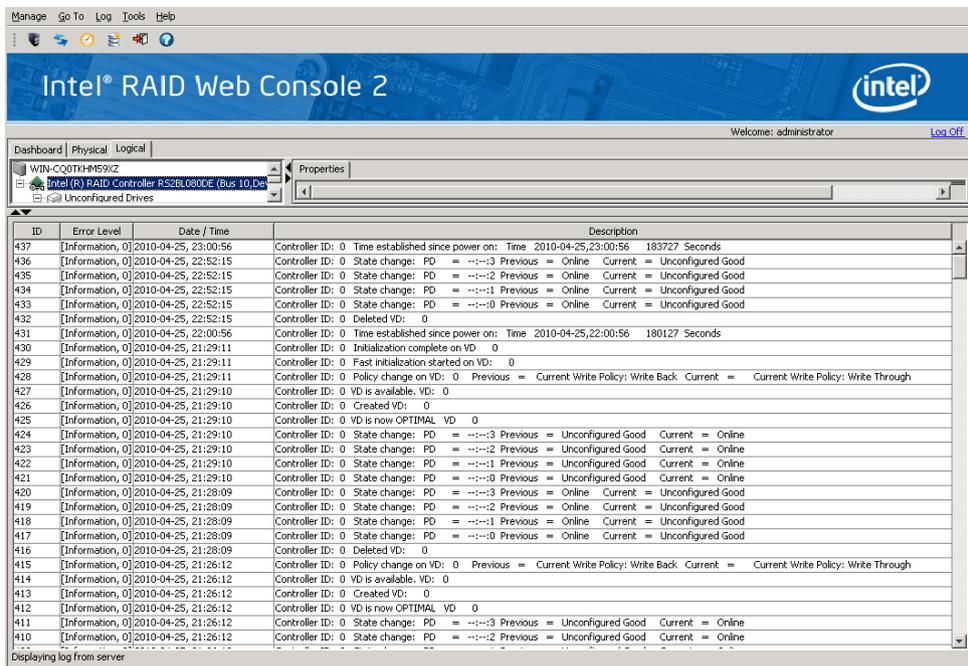


Figure 125. Event Information Window

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 315) The status bar at the bottom of the screen indicates whether the log is a system log (as in [Figure 125](#)) 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.

9.21.2 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 126](#) shows the Controller Information window on the right panel.

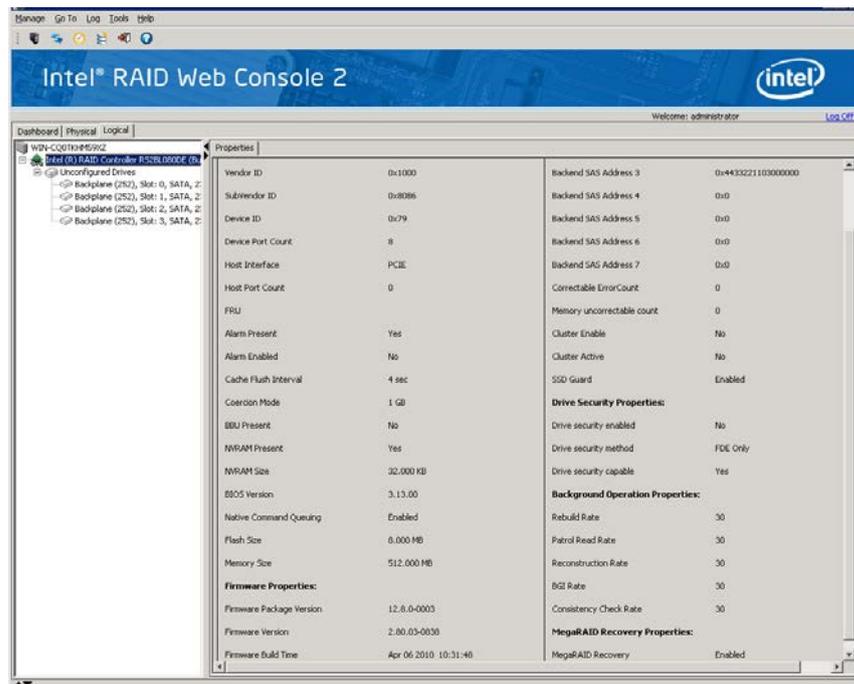


Figure 126. Controller Information

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 186.
- 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 186.

9.21.3 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 127](#) shows the Properties panel for a physical drive.

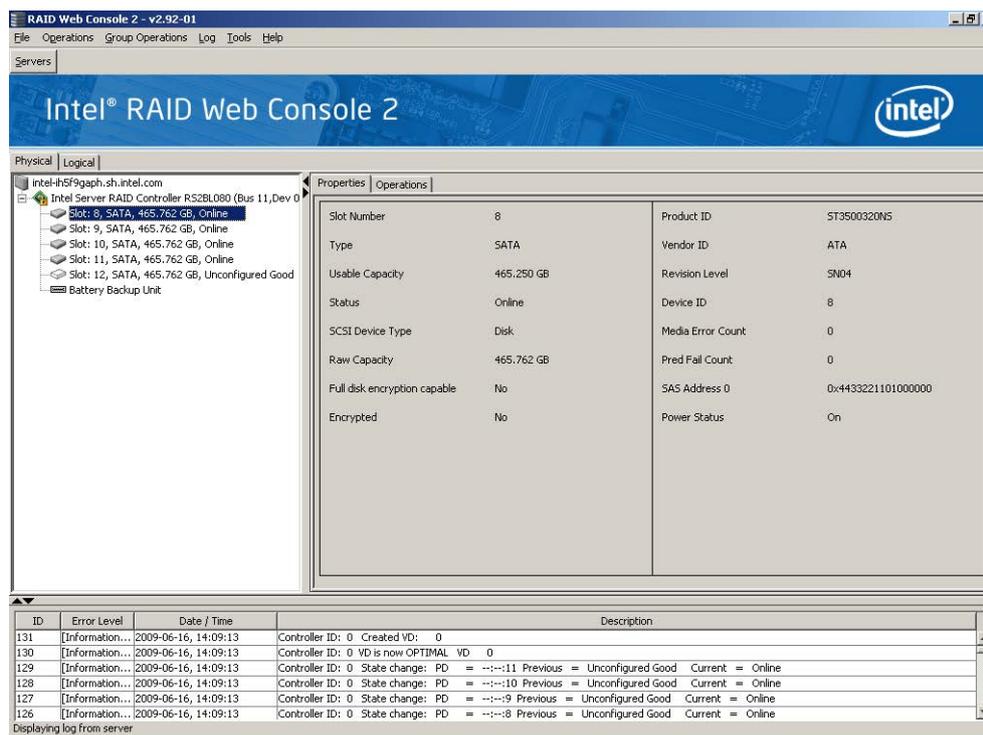


Figure 127. Physical Drive Information

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 128](#)).

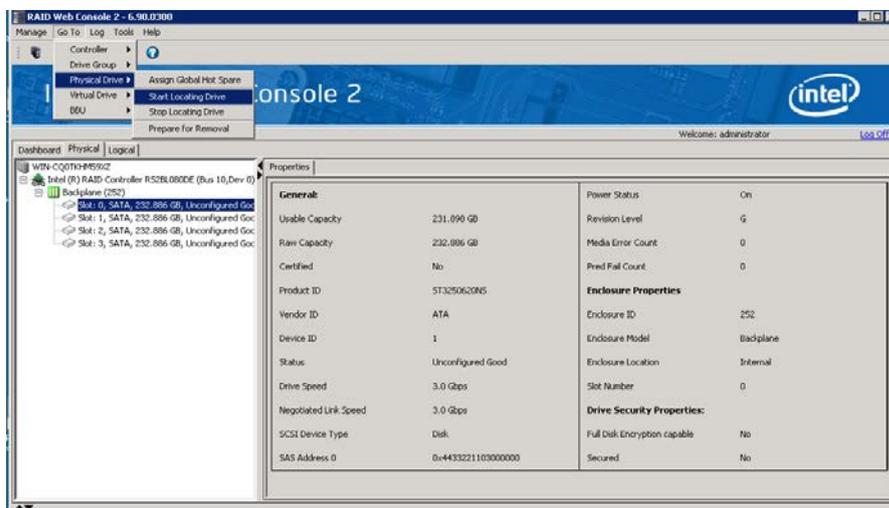


Figure 128. Locating a Physical Drive

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**.

9.21.3.1 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 129](#).

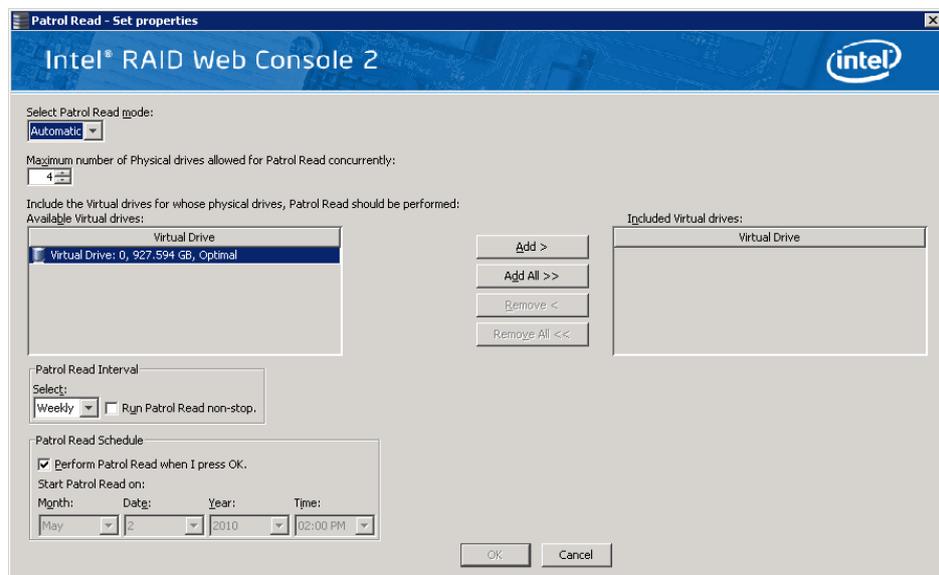


Figure 129. Patrol Read Configuration

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 129](#)).
 - **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.

9.21.4 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 130](#) 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 130](#) shows the Properties tab for a virtual disk in the right panel.

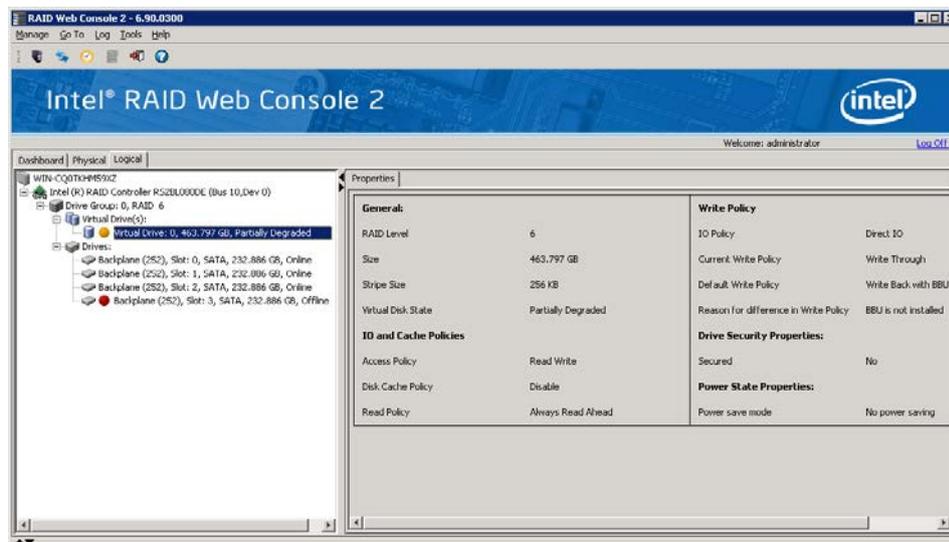


Figure 130. Virtual Drive Properties

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**.

9.21.5 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 131](#) shows it.

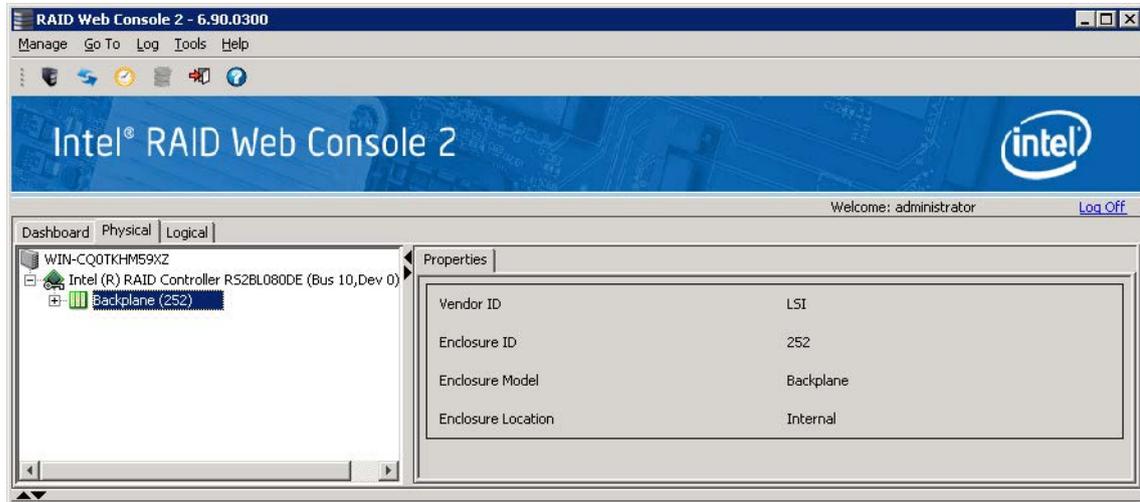


Figure 131. Enclosure Information

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

9.21.6 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 132](#).

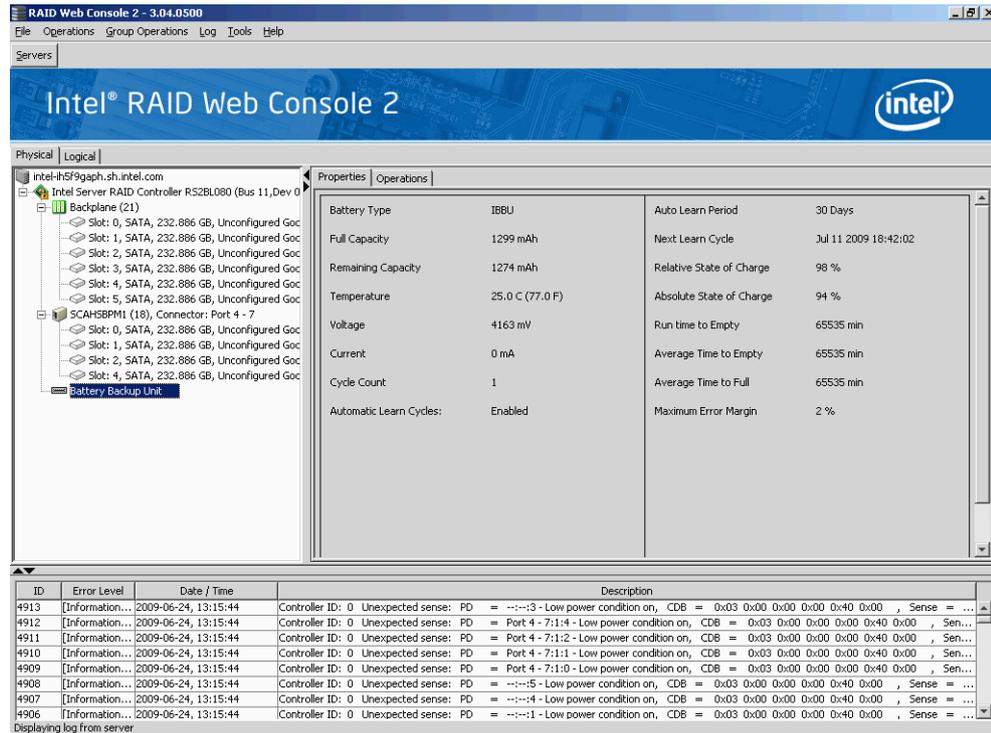


Figure 132. Battery Backup Unit Information

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

9.21.7 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.

9.21.7.1 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 133](#) (Depending on different utility version, the BBU operation also may appear in a new pop-up window).

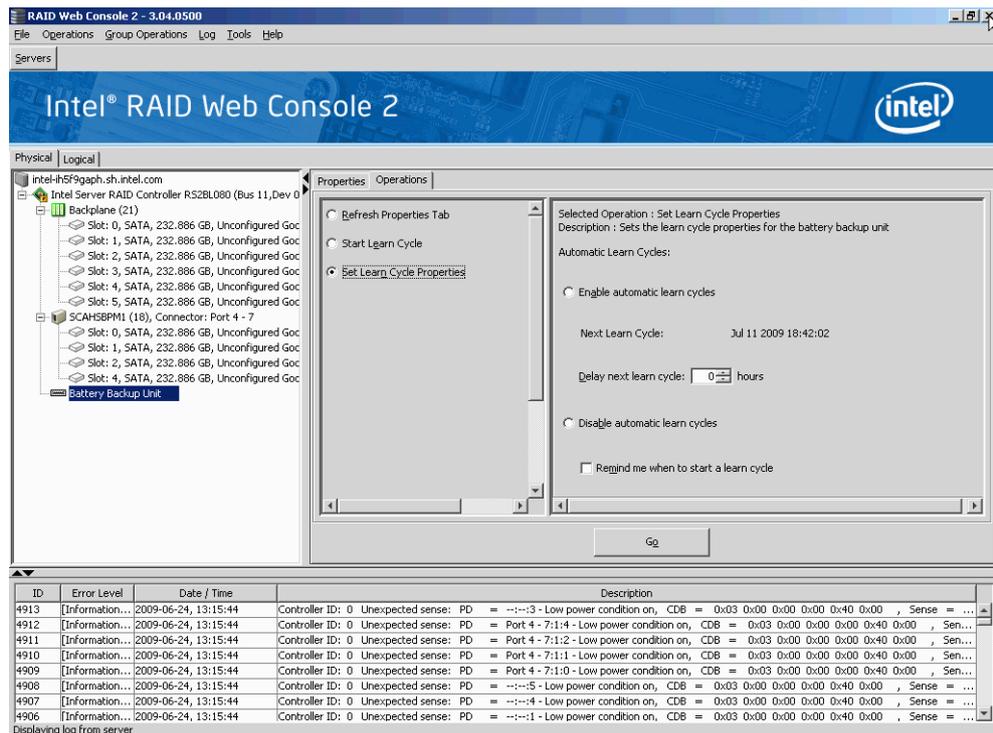


Figure 133. Battery Backup Unit Operations

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.

9.21.7.2 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 133](#) (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**.

9.21.8 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 134](#). You open this window by selecting **Manage -> Show Progress**.

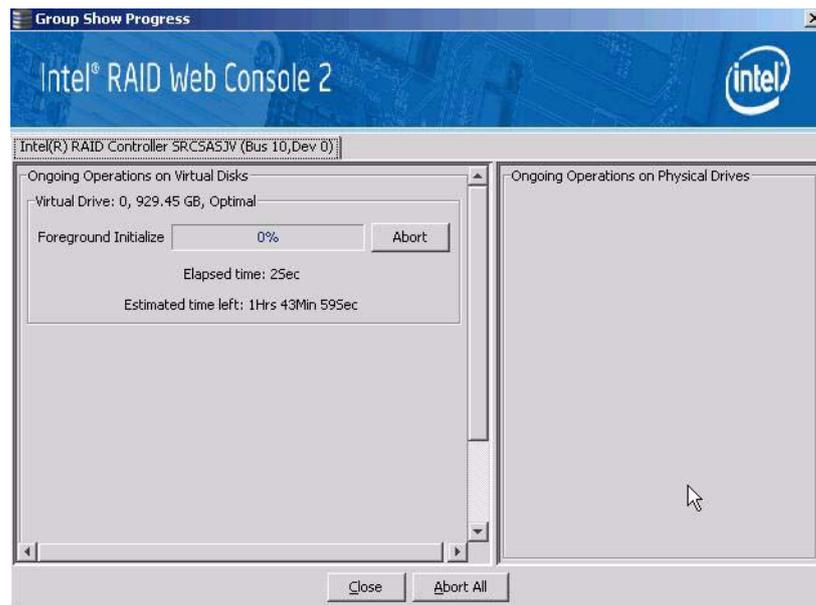


Figure 134. Group Show Progress Window

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 211).
- Rebuild (see [“Rebuilding a Drive”](#) on page 214).

- Reconstruction (see [“Adding a Drive to a Virtual Disk”](#) on page 188).
- Consistency check (see [“Running a Consistency Check”](#) on page 212).

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.*

9.22 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

9.22.1 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**.

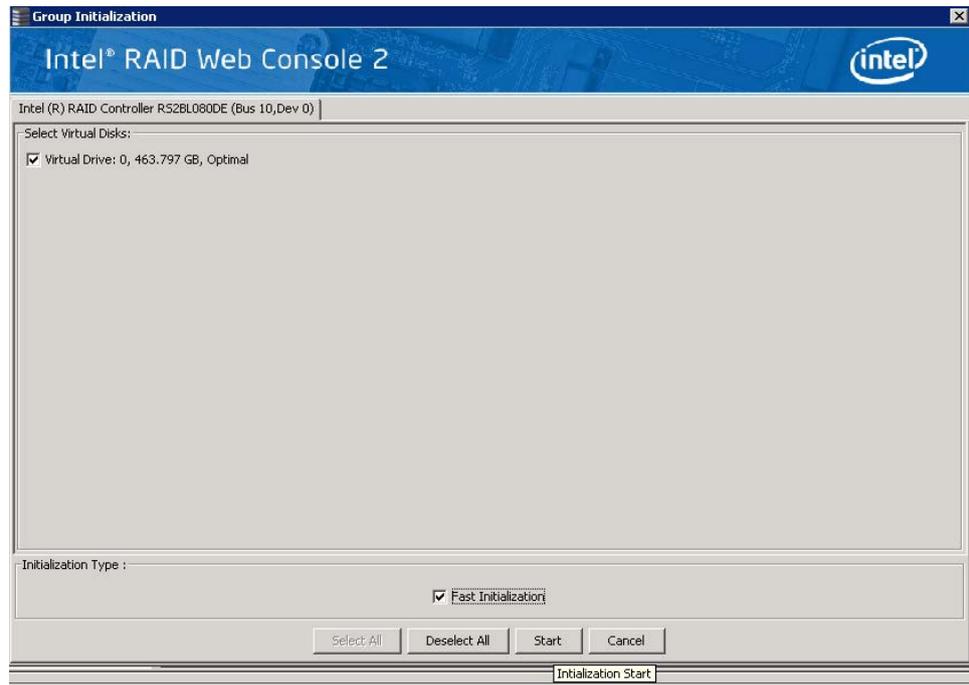


Figure 135. Selecting Initialize

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 209](#) for more information.

9.22.2 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 136](#).

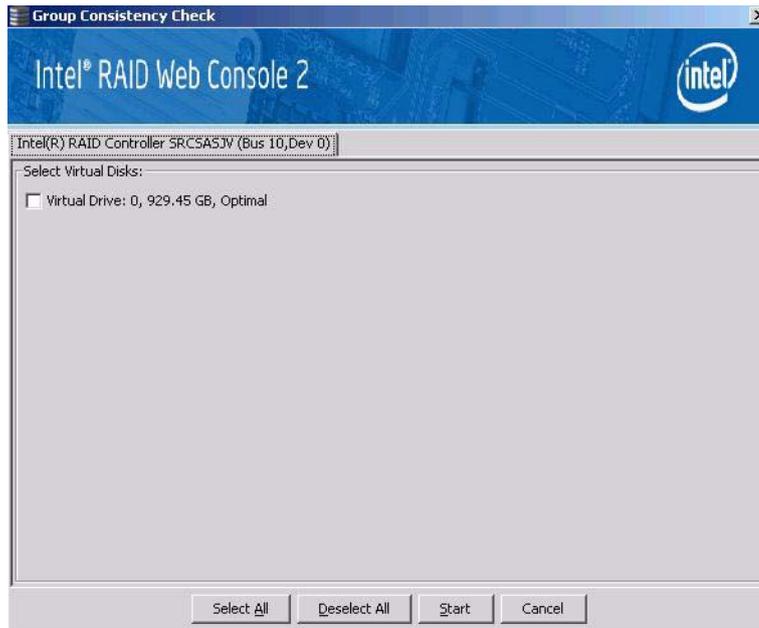


Figure 136. Group Consistency Check Window

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 209.

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**.

9.22.3 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**.

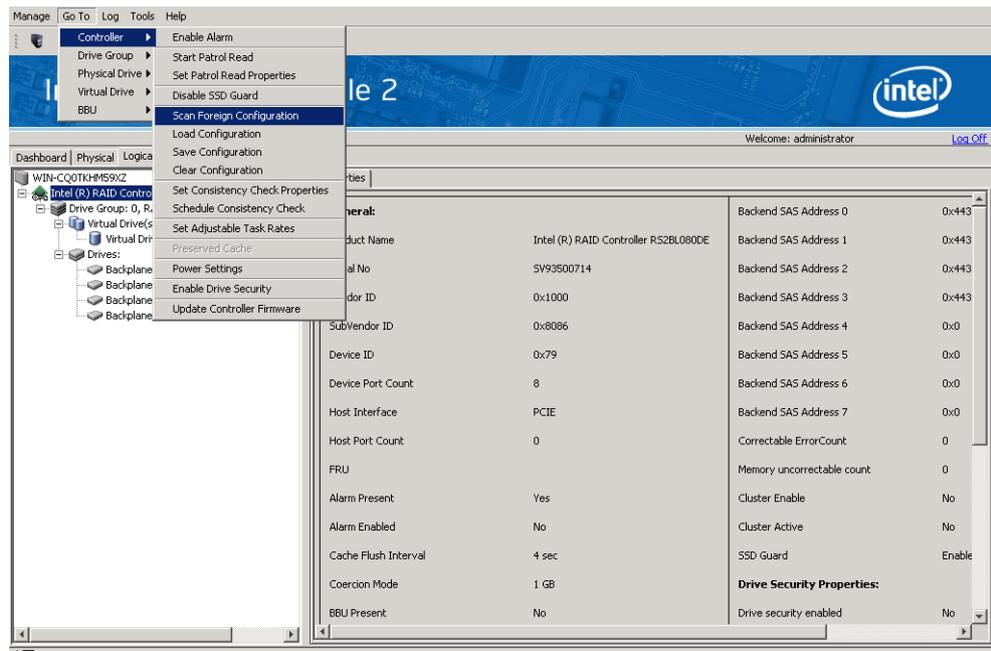


Figure 137. Scan for Foreign Configuration

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

9.22.4 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**.
2. 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.

3. Click the icon of the failed drive in the left panel.
4. Right-click and select **Prepare for Removal**.

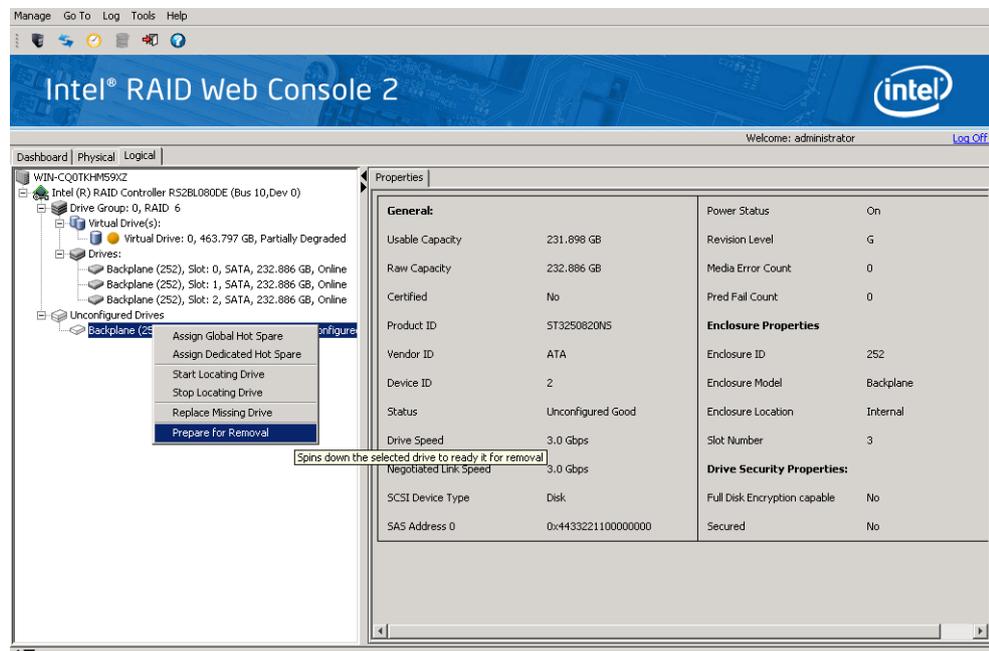


Figure 138. Preparing Drive for Removal

5. 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**.

9.22.5 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:

1. Click the icon of the drive in the left panel and right-click it.
2. Select **Make Drive Offline**.
3. 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.*

9.22.6 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.

9.23 Enabling RAID Premium Features

9.23.1 Enabling Full Disk Encryption feature

This feature is supported by Intel® RAID controller with Intel® RAID Premium Feature Key AXXRPFKDE or AXXRPFKDE2 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), or Intel® RAID Premium Feature Keys AXXRPFKSSD2, AXXRPFKDE2, and AXXRPFKSNSH2 Installation Guide (G29824-00x) 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 locked when powered down.
- It provides secure data when a drive is misplaced, failed, experiencing drive theft, experiencing Server theft “Smash and 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 follows:

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**.

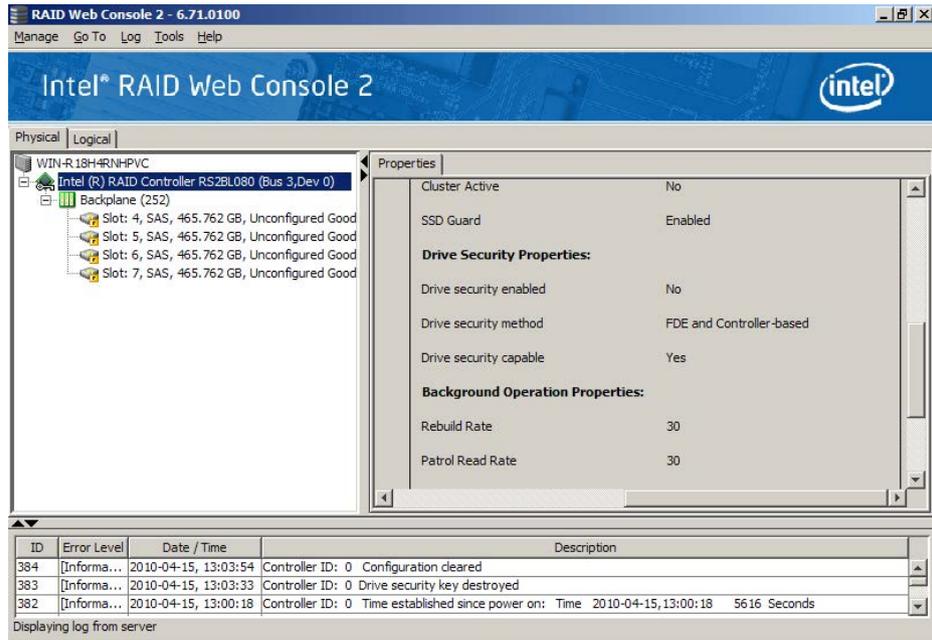


Figure 139. Check Controller Security status

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**.

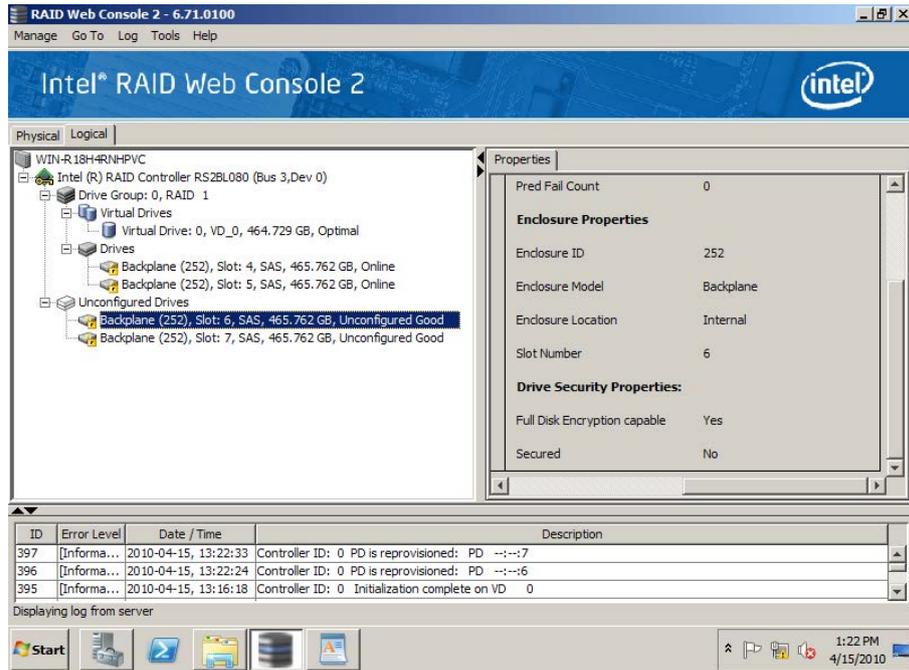


Figure 140. Check Drive Security status

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

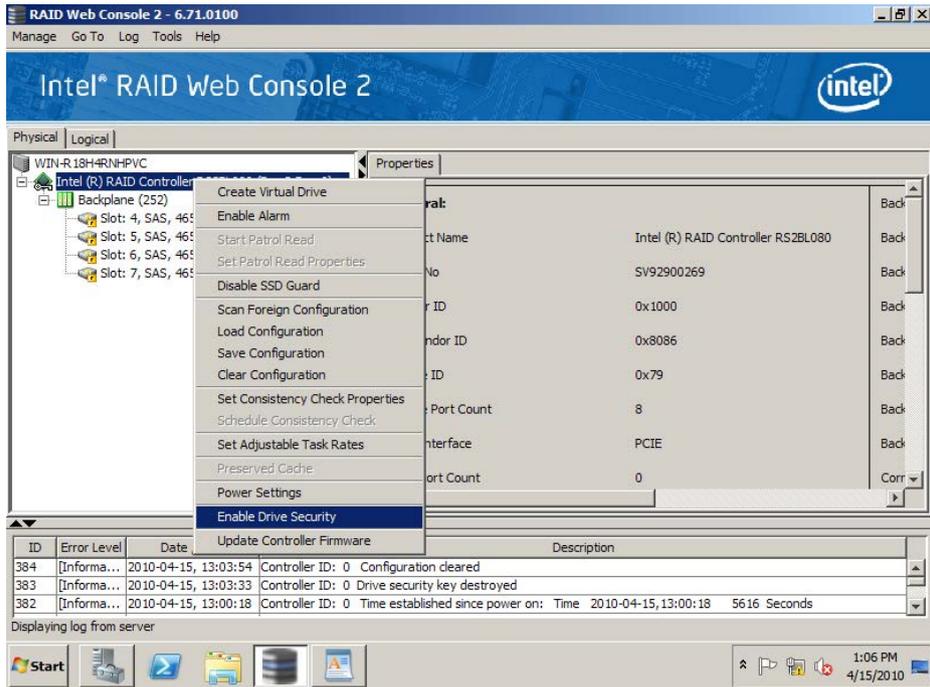


Figure 141. Enable Drive Security

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

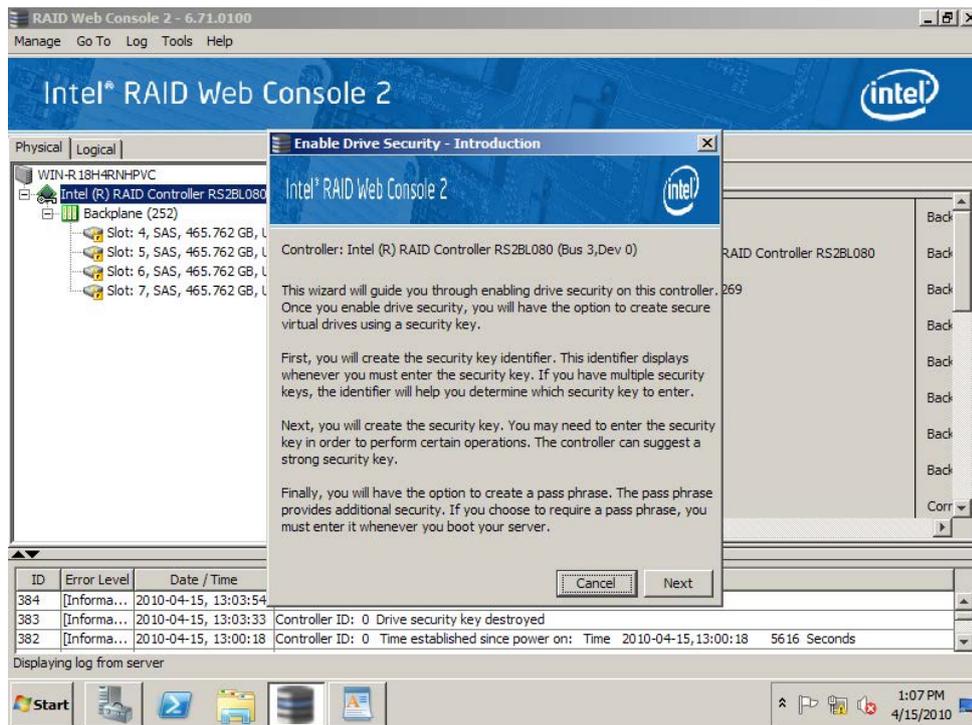


Figure 142. Start Security Wizard

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

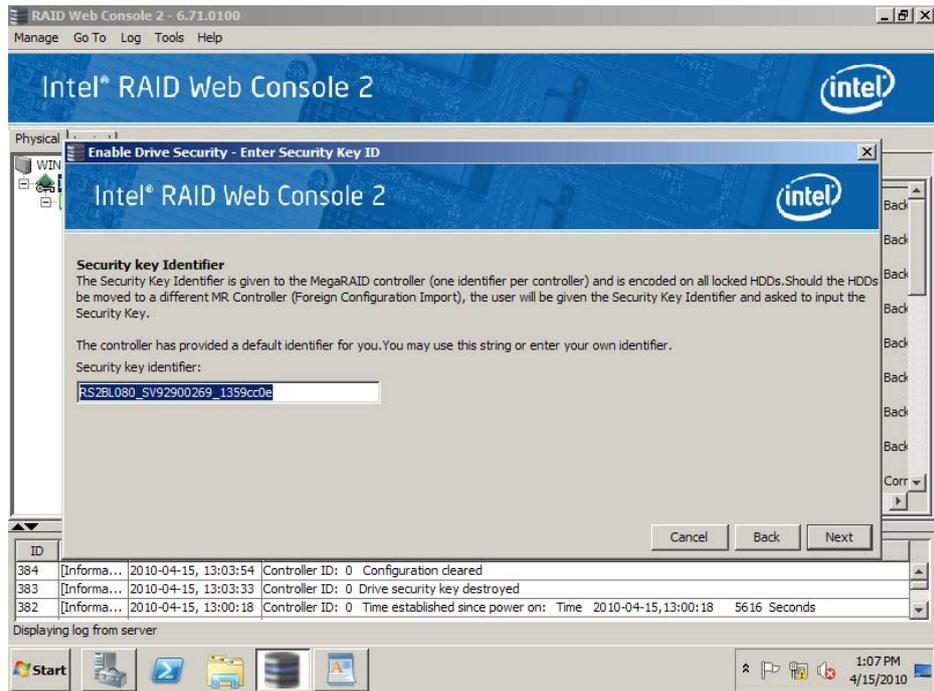


Figure 143. Enter Security Key Identifier

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

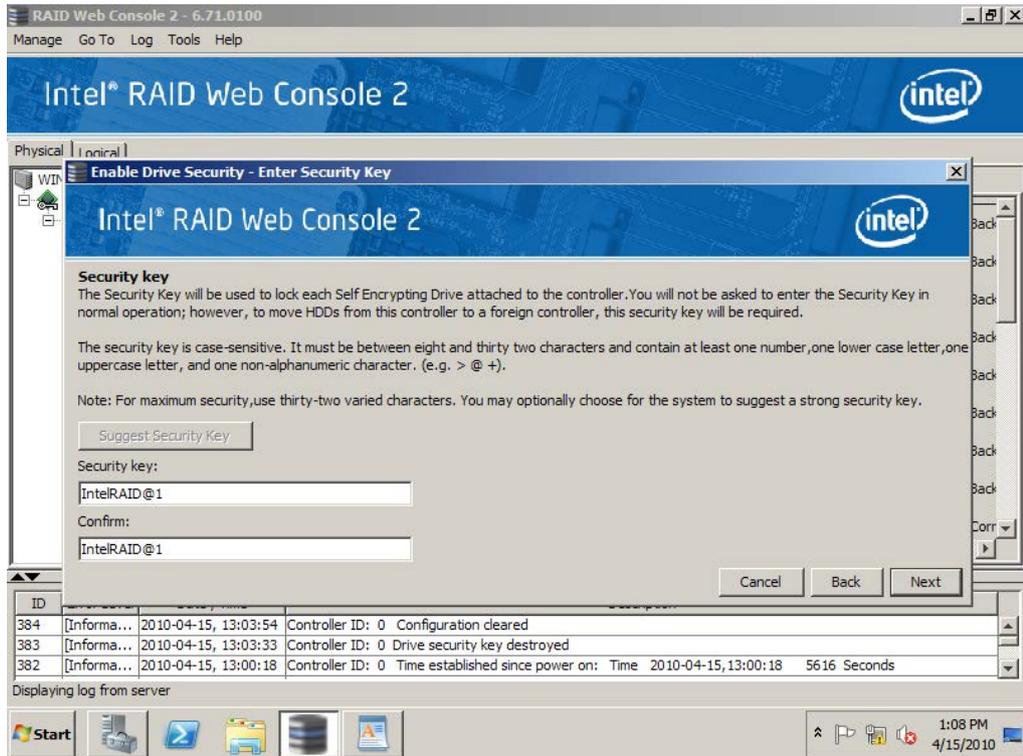


Figure 144. Enter Security Key

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

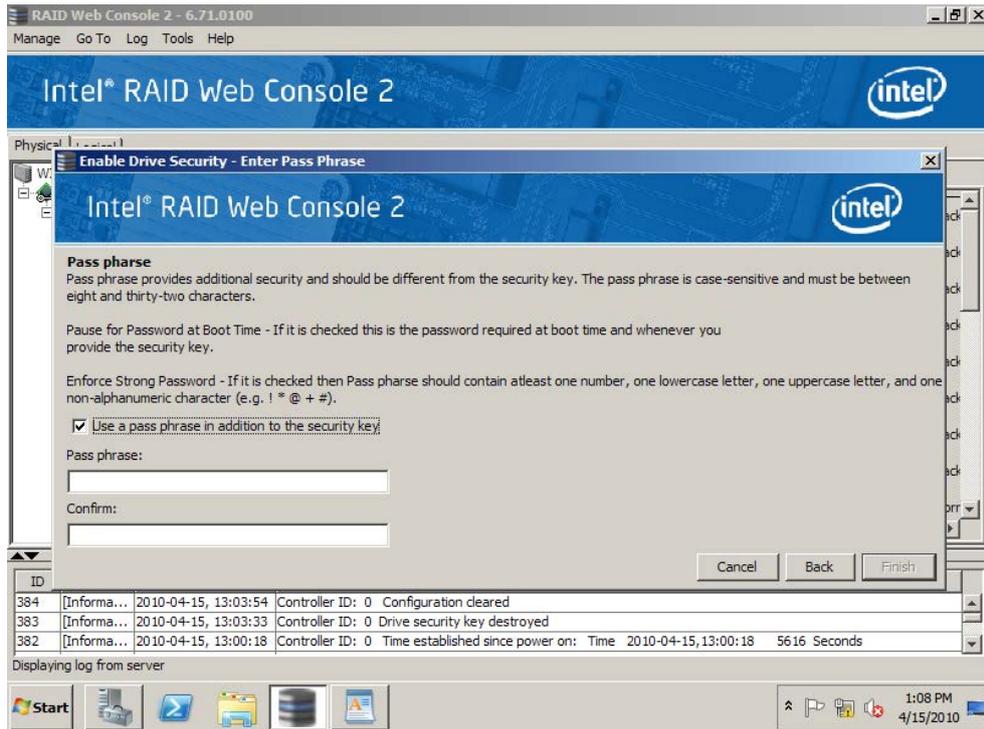


Figure 145. Enter Pass Phrase

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

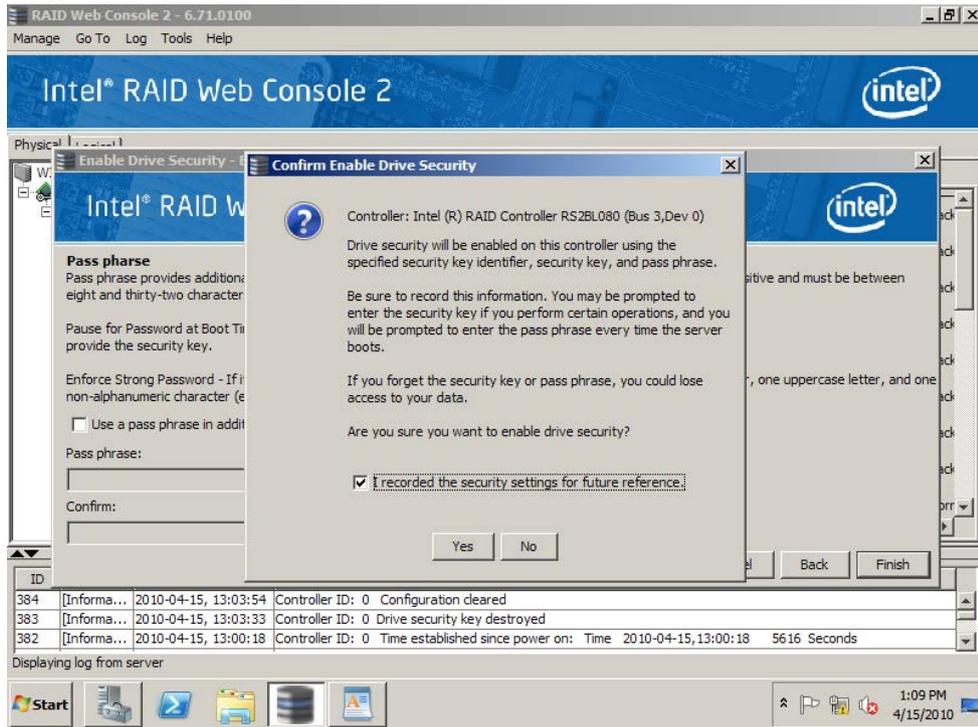


Figure 146. Confirm Enable Drive Security

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

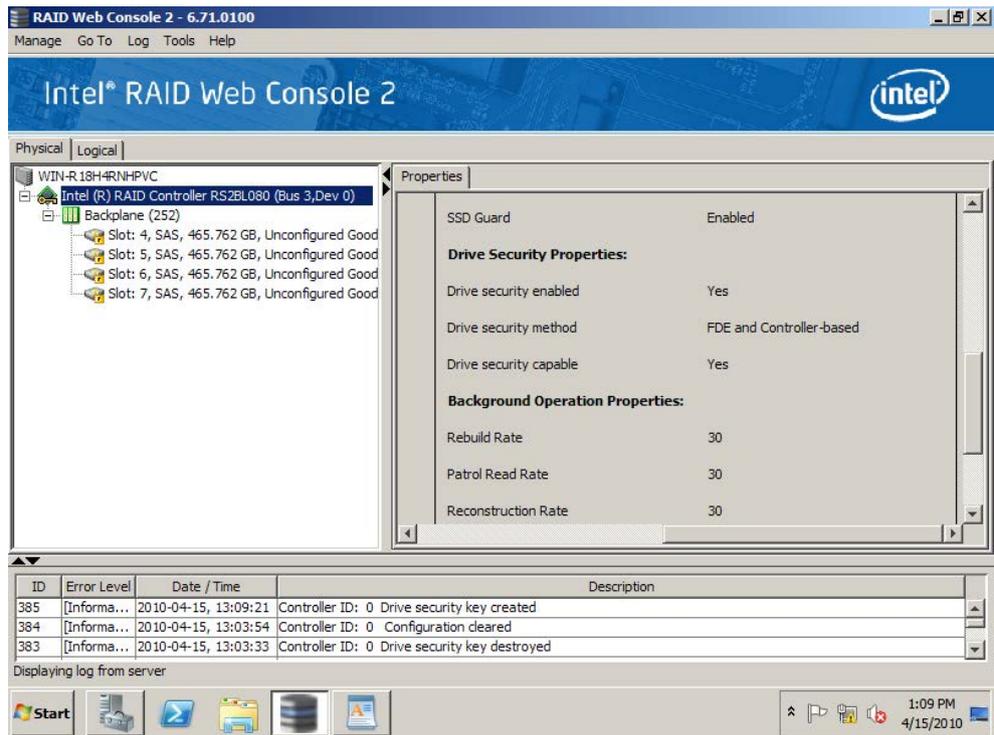


Figure 147. Check Drive Security Enabled status

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

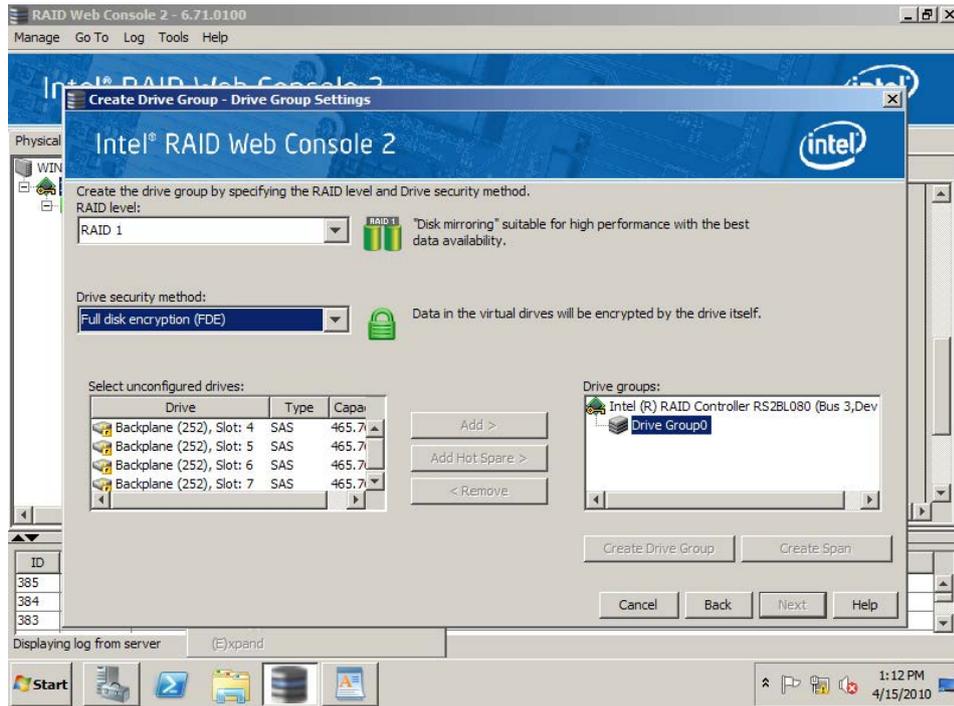


Figure 148. Select Full Disk Encryption

- 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.

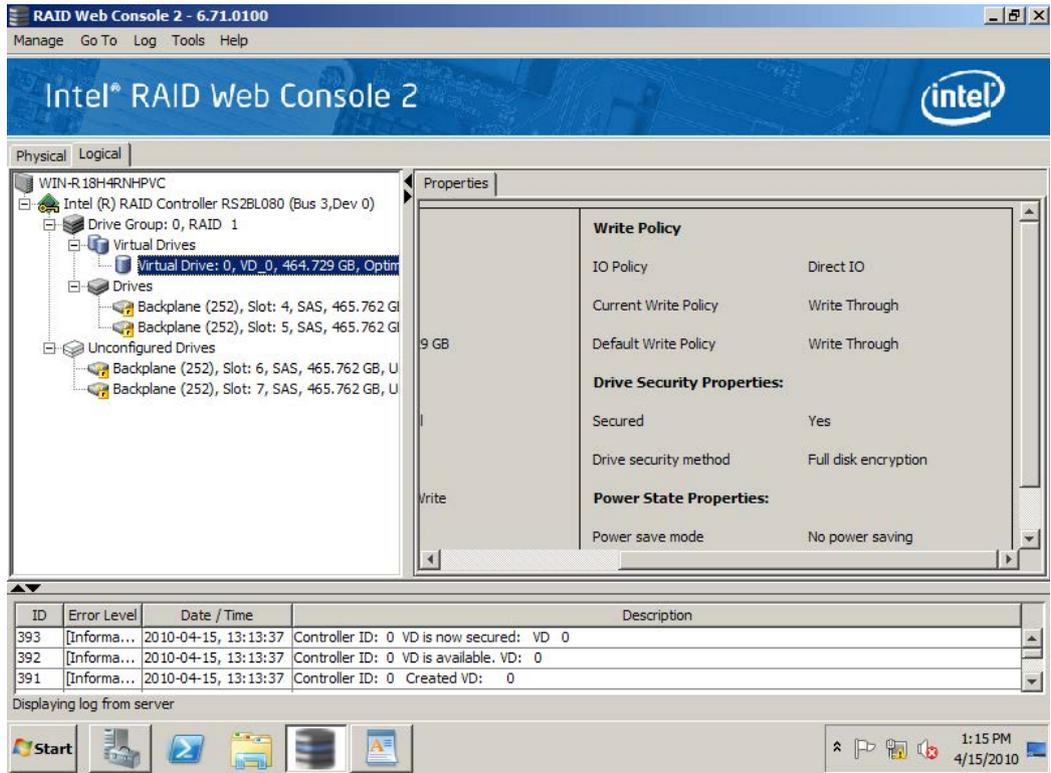


Figure 149. Create RAID Virtual Drive with FDE enabled

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**.

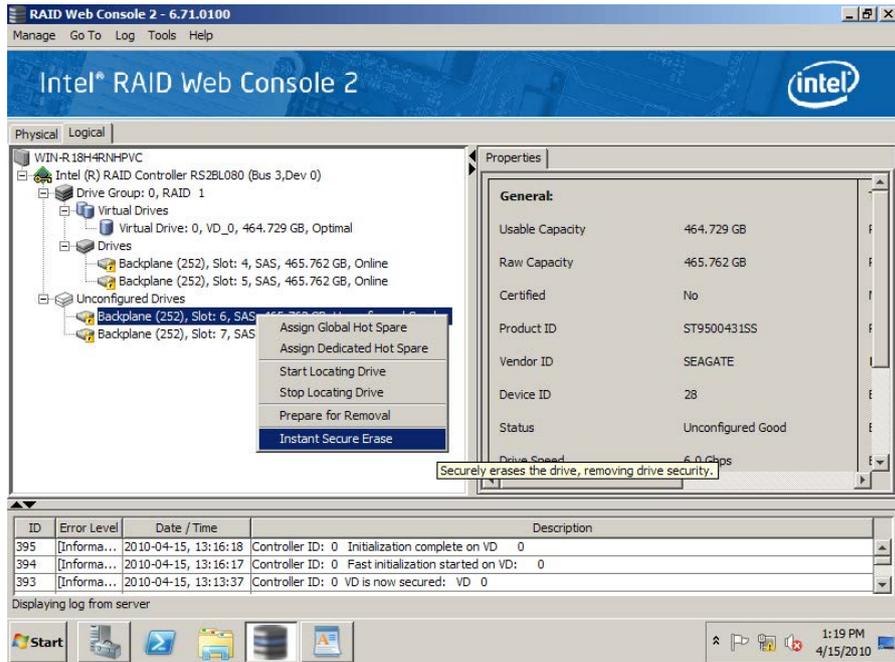


Figure 150. Instant Secure Erase

13. Accept confirmation to permanently erase disk drive.

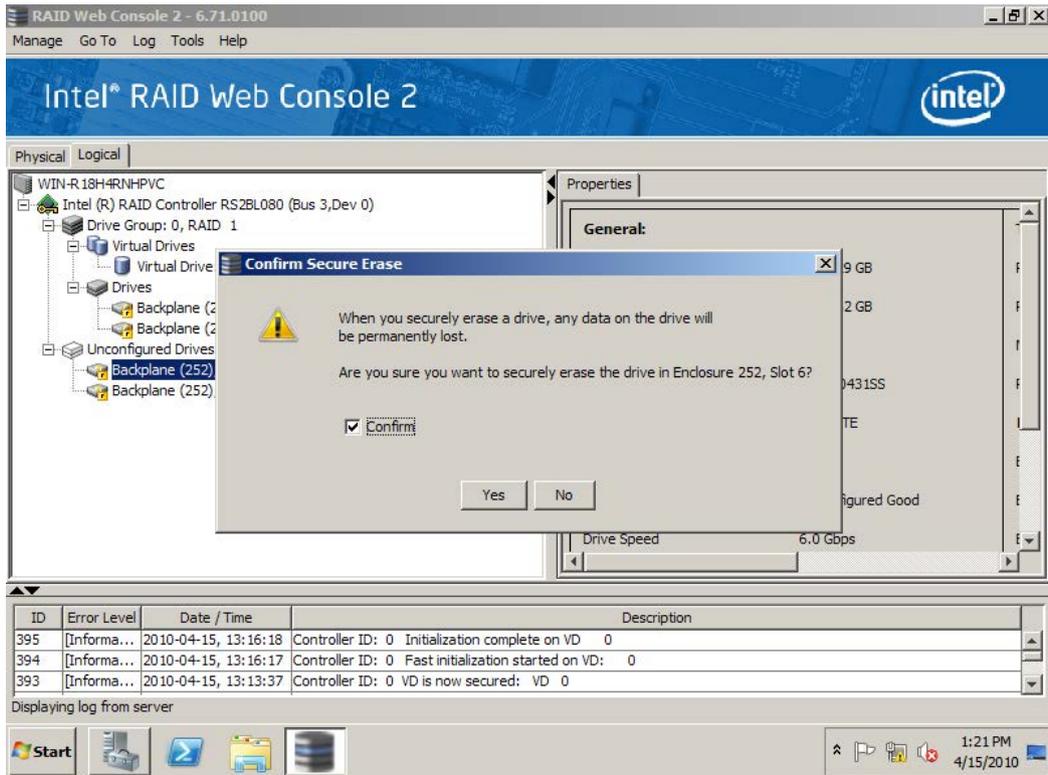


Figure 151. Confirm Secure Erase

9.23.2 Enabling Snapshot feature

This feature is supported by Intel® RAID RAID controller with Intel® RAID Premium Feature Key AXRPFKSNH or AXRPFKSNH2 installed. Refer to Intel® RAID Premium Feature Key AXRPFKSSD, AXRPFKDE, and AXRPFKSNH Installation Guide (E88588-00x), or Intel® RAID Premium Feature Keys AXRPFKSSD2, AXRPFKDE2, and AXRPFKSNH2 Installation Guide (G29824-00x) 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 follows:

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.

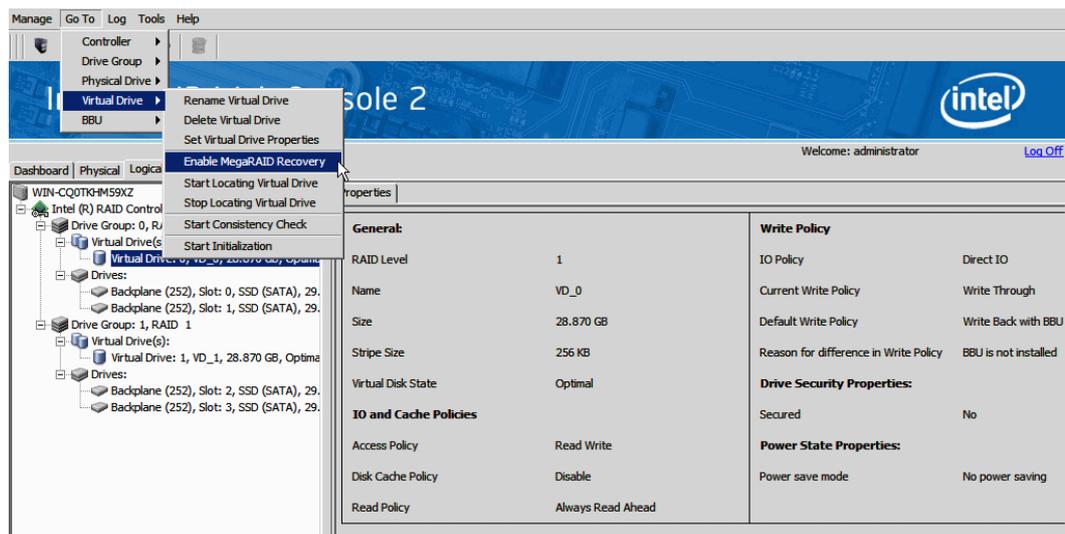


Figure 152. Enable MegaRAID Recovery

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.

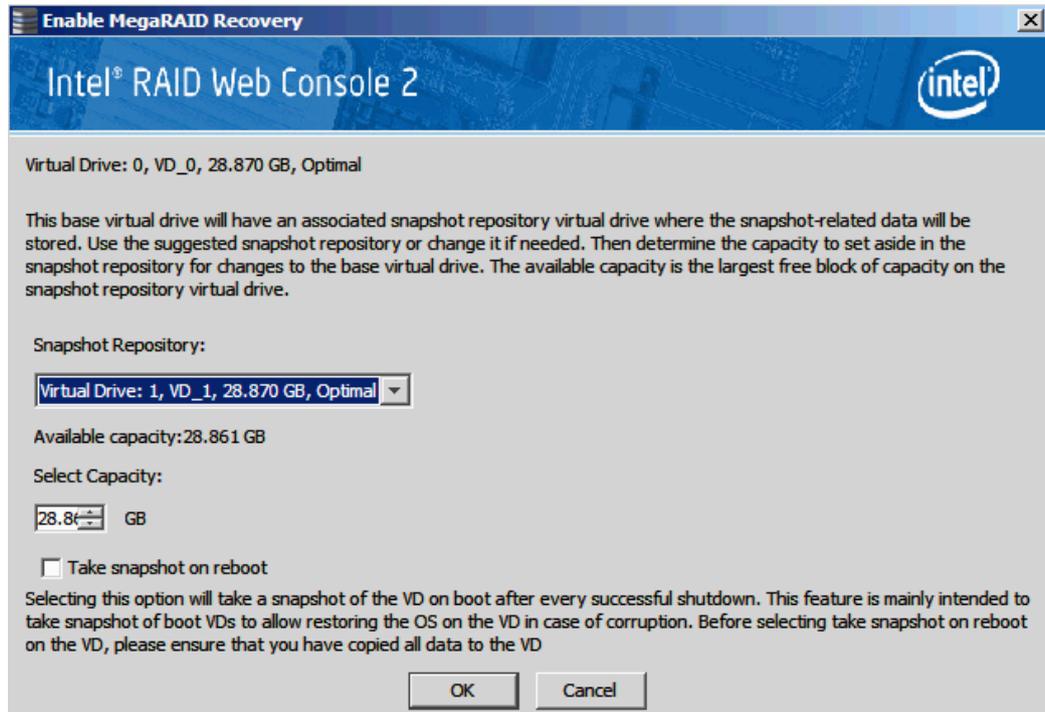


Figure 153. Enter the Capacity for Snapshot Repository

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**.

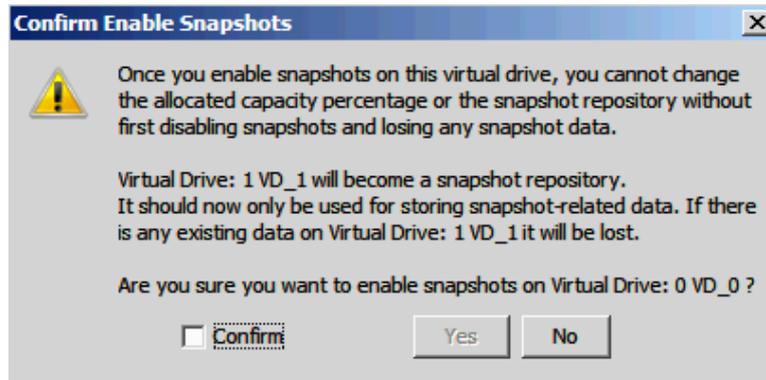


Figure 154. Confirm Enable Snapshot

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.

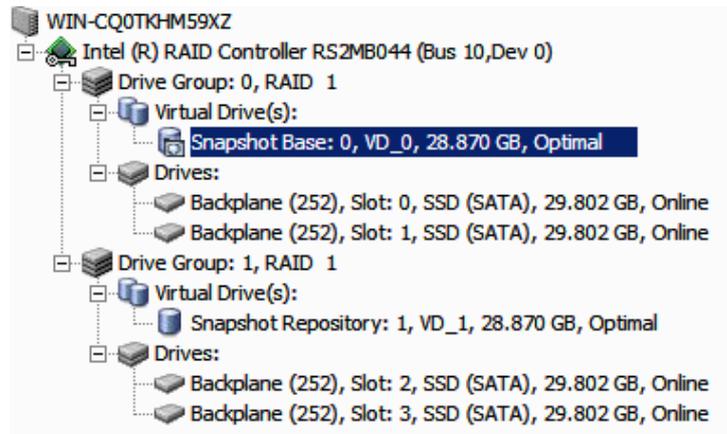


Figure 155. Snapshot Base is shown

9. 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 showed 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.

10. 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.

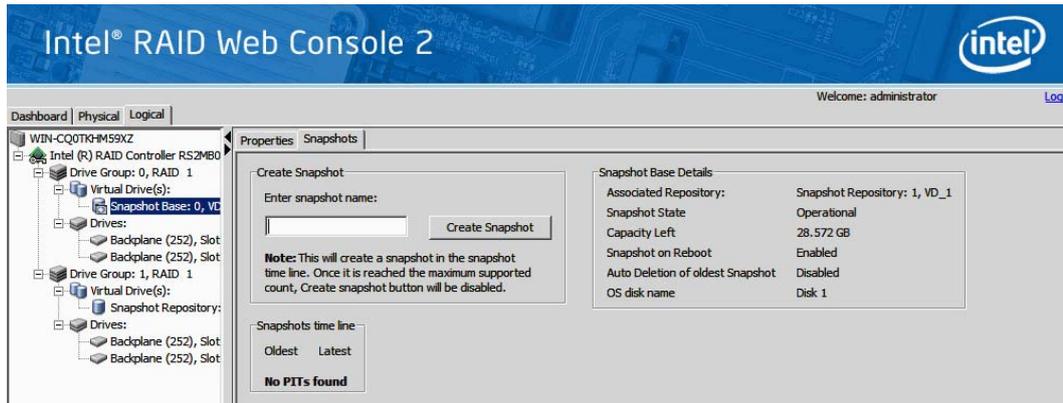


Figure 156. Enter Snapshot Name

11. Enter the snapshot name in the **Enter snapshot name** field. Click **Create Snapshot**.
12. 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.

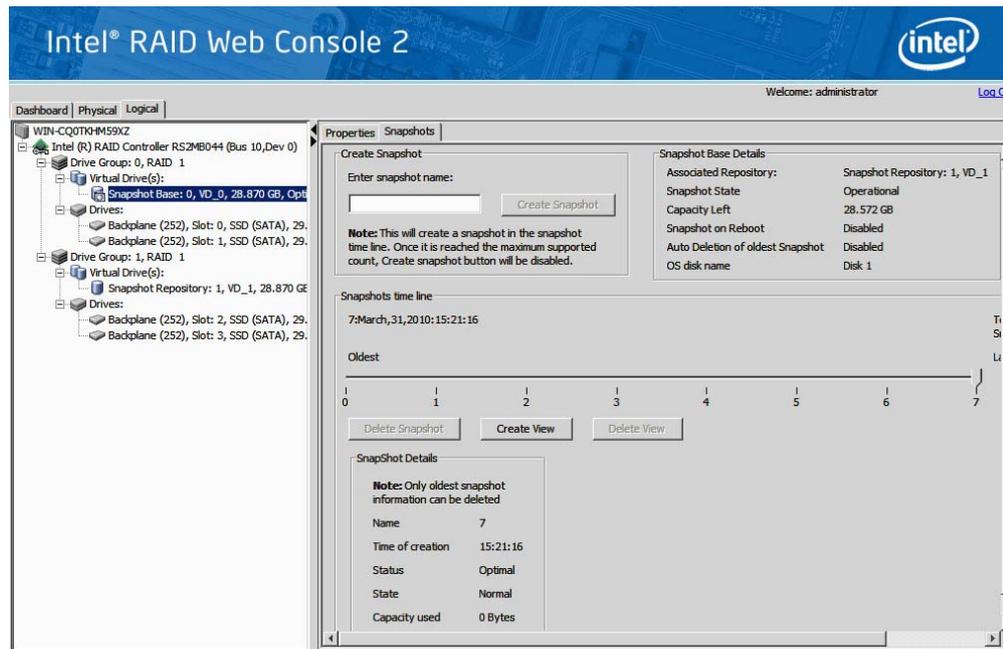


Figure 157. Create Snapshot

13. (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**.

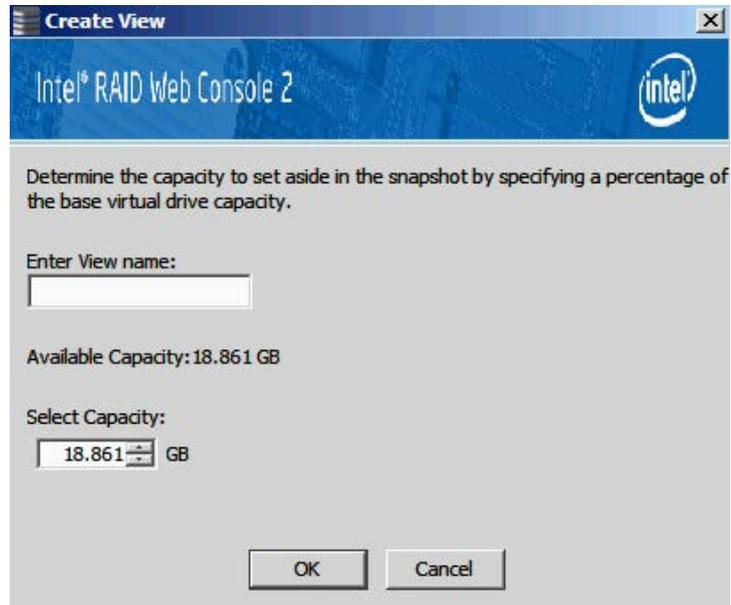


Figure 158. Create View

14. 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**.

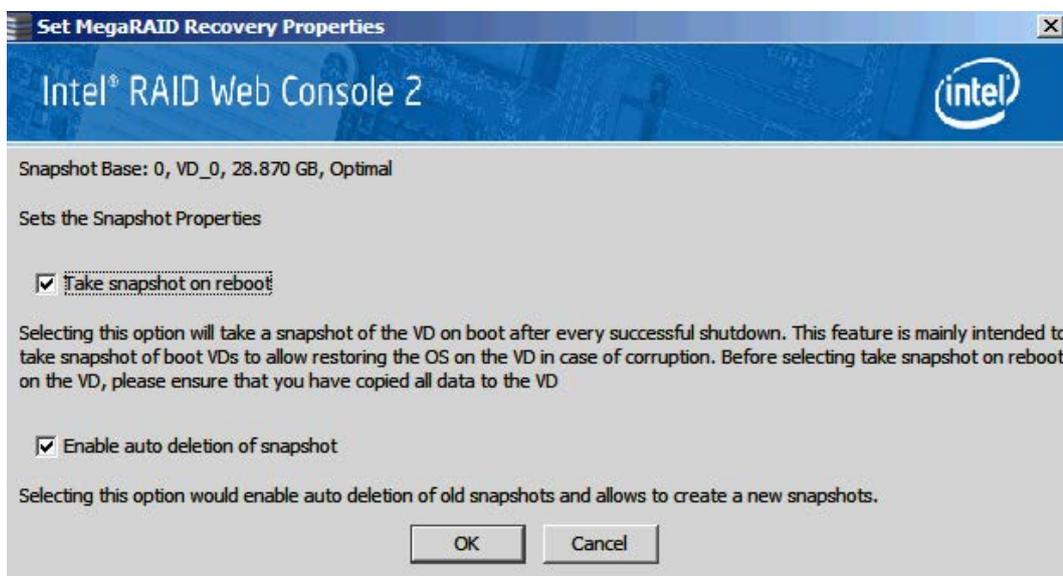


Figure 159. Set MegaRAID Recovery Properties

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.

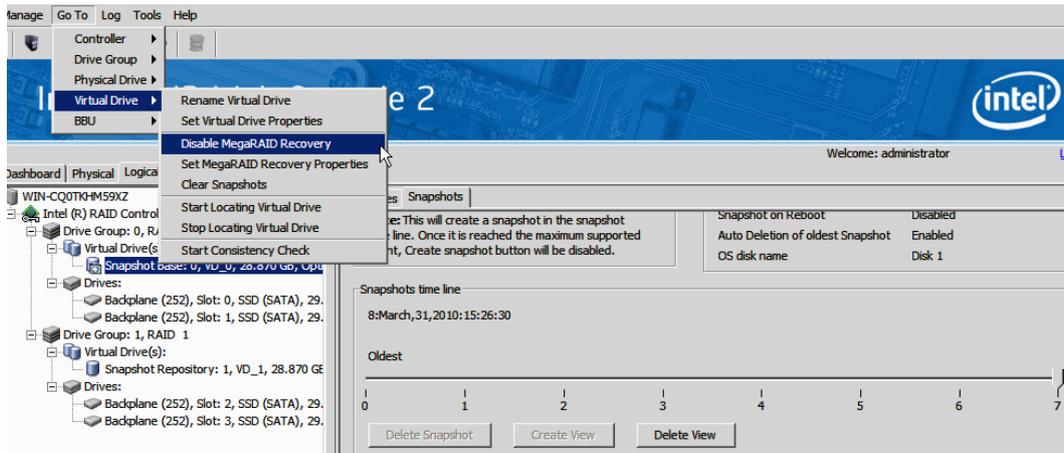


Figure 160. Disable MegaRAID Recovery

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.

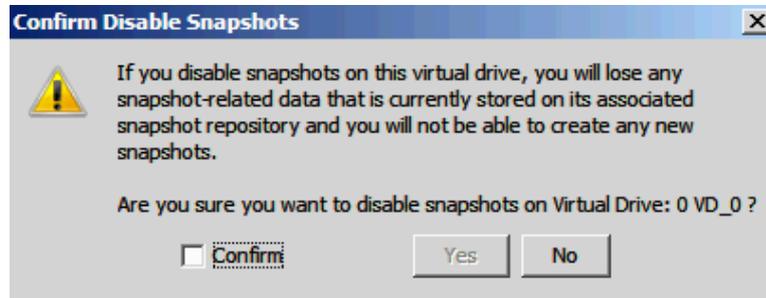


Figure 161. Confirm Disable Snapshots

23. After snapshots are successfully created (in step 12, [Figure 157](#)), 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 following screen, press **Start**.



Figure 162. Adapter Selection

- In the home page, from the **Logical View** in the right panel, click the **Snapshot Base** which has the operating system installed.

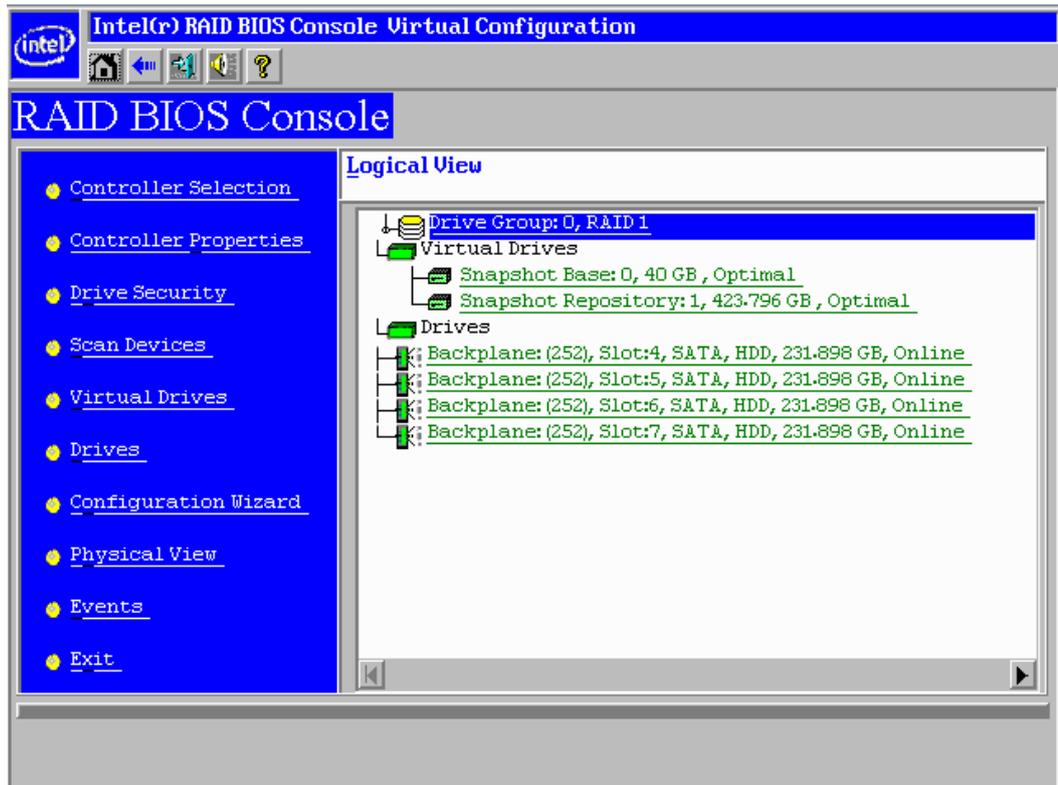


Figure 163. Selecting Snapshot Base

- In the new screen, select **Adv Opers** (which means Advanced Operations), and click **Go**.

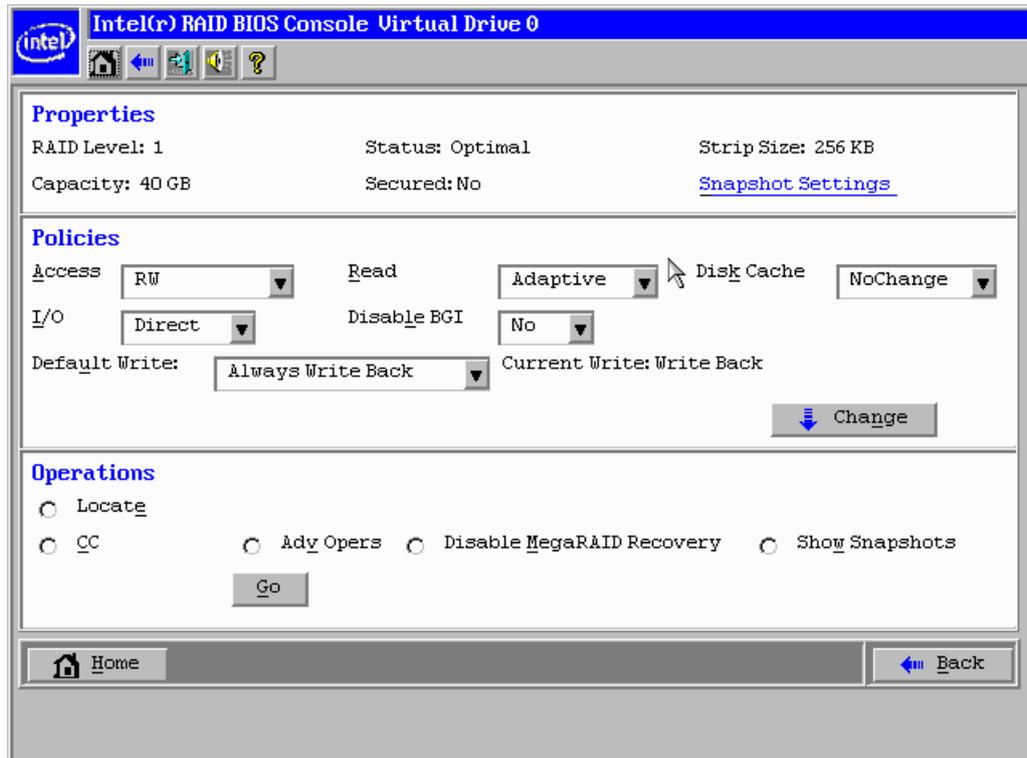


Figure 164. Selecting Advanced Operations

27. In the Advanced Operations screen, select **Rollback**.

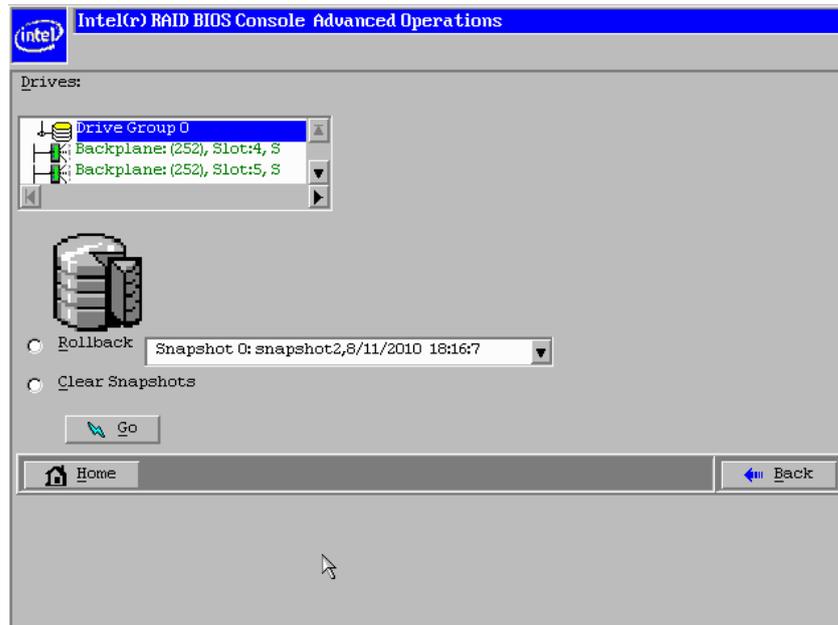


Figure 165. Selecting Rollback

28. From the drop-down menu, select a previous created snapshot, and then click **Go**.

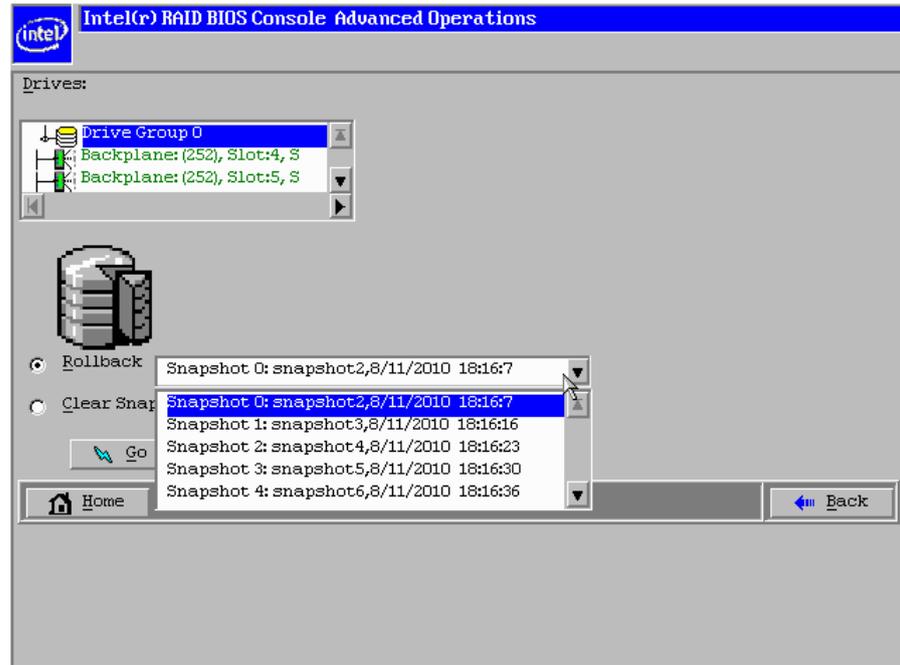


Figure 166. Selecting a snapshot

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**.

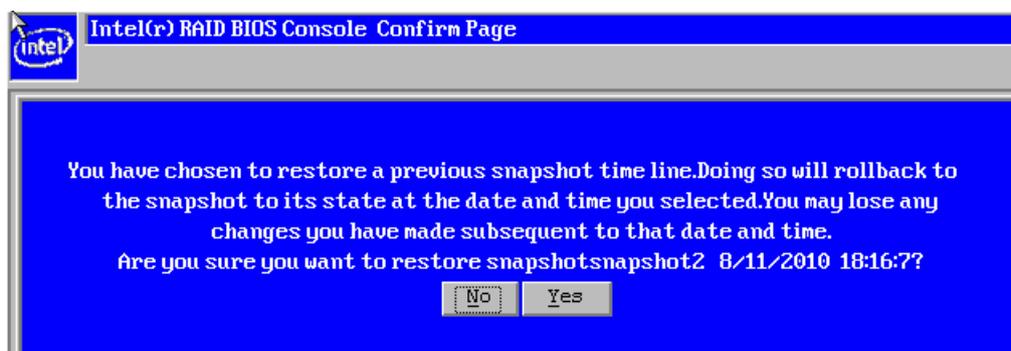


Figure 167. Confirm Page

30. The screen should go back to the display shown 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.

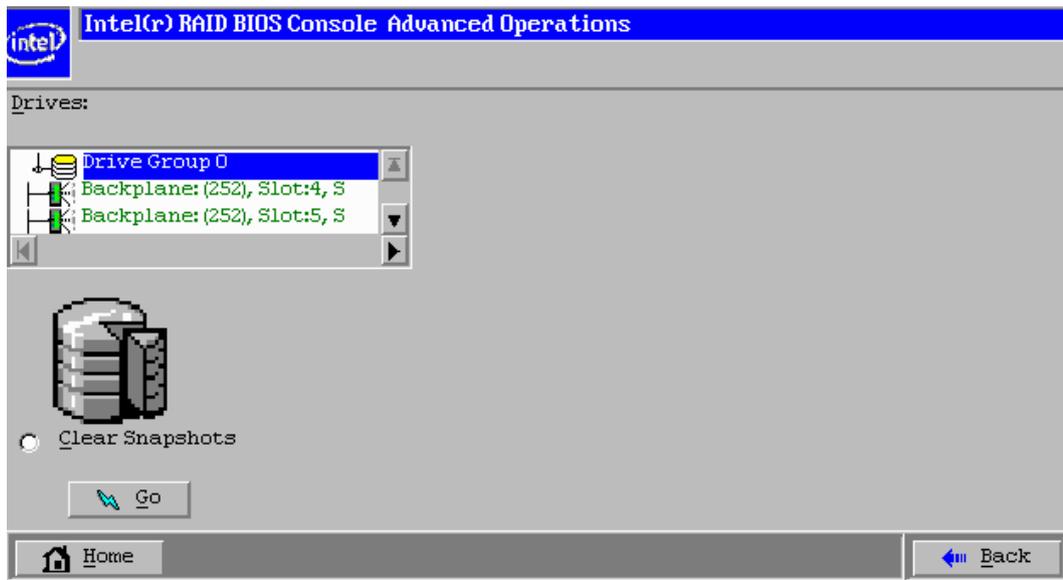


Figure 168. Rollback operation is done

9.23.3 Enabling Super Sized Cache

This feature is supported by Intel® RAID Premium Feature Key AXXRPFKSSD or AXXRPFKSSD2 installed. Refer to Intel® RAID Premium Feature Key AXXRPFKSSD, AXXRPFKDE, and AXXRPFKSNSH Installation Guide (E88588-00x), or Intel® RAID Premium Feature Keys AXXRPFKSSD2, AXXRPFKDE2, and AXXRPFKSNSH2 Installation Guide (G29824-00x) 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 follows:

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).

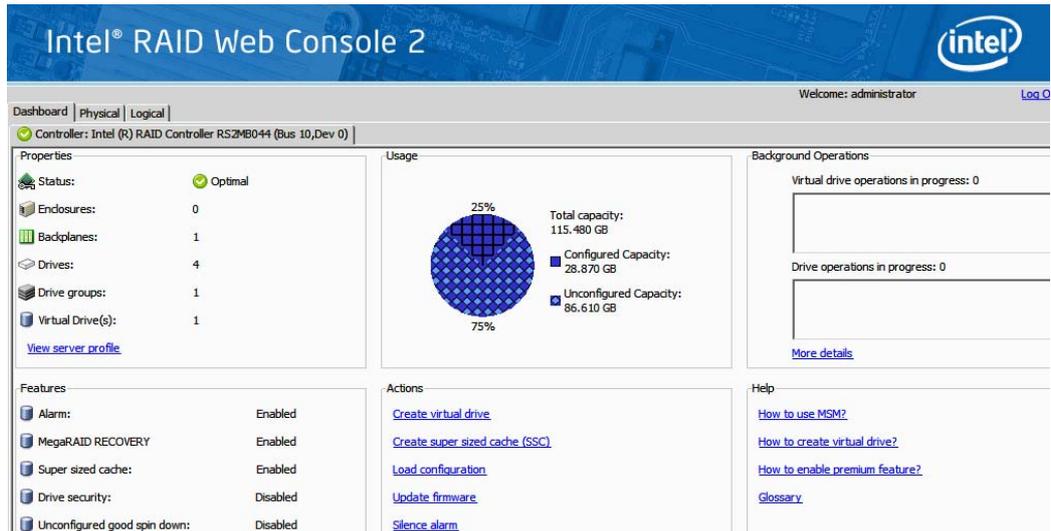


Figure 169. Create SSC from Dashboard

2. The **Create Super Sized Cache Drive group** screen appears. Select the **unconfigured SSD drive**. Click **Add >** -> **Create SSC Drive Group** -> **Next**.

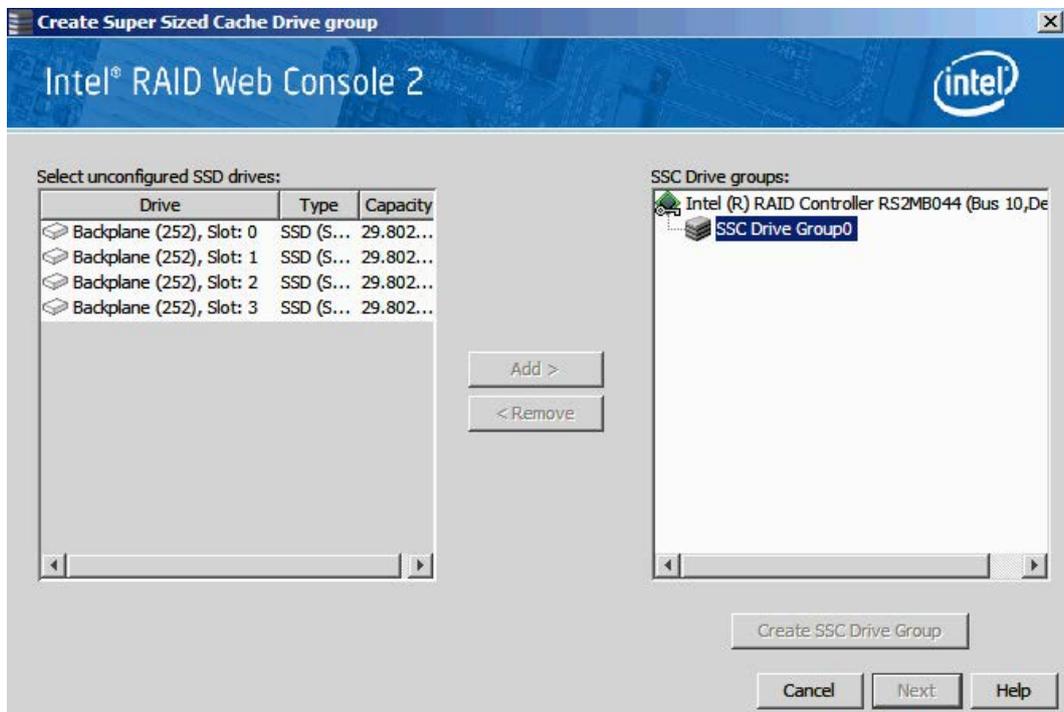


Figure 170. Create SSC Drive Group

3. The **Create Super Sized Cache Drive** screen appears. Click **Create** to create a SSC drive group. Click **Next**.

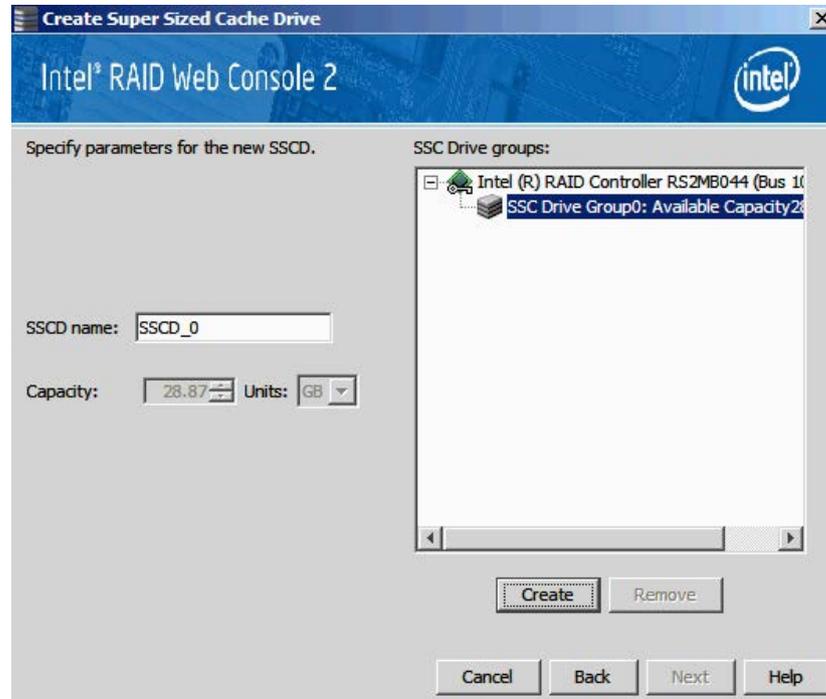


Figure 171. Create SSCD name

4. Create Super Sized Cache screen appears. Click **Finish** to confirm the information and create the Super Sized Cache Disk (SSCD).

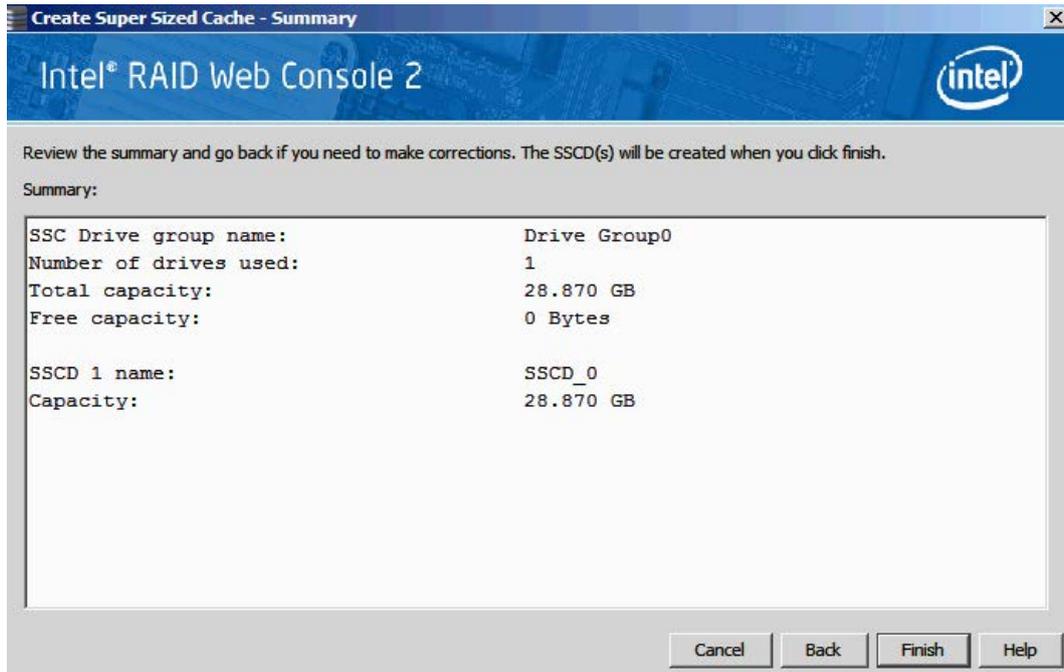


Figure 172. SSC Summary

5. You can see **Super Sized Cache Drive Group** in the logical view.



Figure 173. SSCD status shown

6. 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.

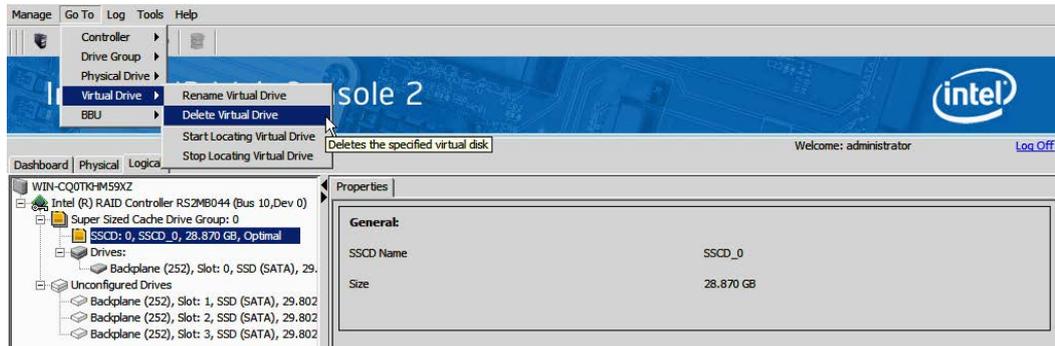


Figure 174. Delete SSCD

7. A confirmation dialog box appears. Check the box next to the **Confirm** field and click **Yes**. The SSCD is deleted.

10 StorCLI

10.1 Overview

The Storage Command Line Tool (StorCLI) is the command line management software designed for the Intel® RAID Controller product line. The StorCLI is a command line interface that is designed to be easy to use, consistent, and easy to script. This chapter provides information on installing and using the StorCLI and explains the various features of the StorCLI.

Note: The legacy commands are deprecated from this guide.

10.2 Devices Supported by the StorCLI Tool

The StorCLI tool is designed to work with the Intel® RAID Controller product line. The StorCLI tool supports the following Intel® RAID Controller products.

- Intel® RAID Controller RS3MC044
- Intel® RAID Controller RS3SC008
- Intel® RAID Controller RS3DC080
- Intel® RAID Controller RS3DC040

10.3 Installation

The Intel® RAID Controllers can be used with the following operating systems for Intel and AMD 32-bit and 64-bit x86-based motherboards:

- Microsoft Windows Server 2008* R2
- Microsoft Windows 7* (32/64 bit)
- Red Hat* Enterprise Linux 5.8 (32/64 bit)
- Red Hat* Enterprise Linux 6.1
- Red Hat* Enterprise Linux 6.2 (32/64 bit)
- SUSE* Linux Enterprise Server 11 SP2 (32/64 bit)
- SUSE* Linux Enterprise Server 10 SP4 (32/64 bit)
- Fedora Core Linux 15
- VMware* ESX 4.0

- VMware* ESX 4.1 U2
- VMware* ESXi 4.1 U2
- VMware* ESXi 5.0 U1
- Solaris*
- FreeBSD (32/64 bit)
- EFI
- Ubuntu*

Note: *The LSISAS2208 and LSISAS2108 controllers provide support for Microsoft Windows 8 and Microsoft Windows Server 2012 operating systems.*

10.3.1 Installing StorCLI on Microsoft Windows Operating Systems

The Windows StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the Intel website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.

Note: *StorCLI must be run with the administrator privileges.*

10.3.2 Installing StorCLI on Linux Operating Systems

To install StorCLI on Linux operating systems, perform the following steps:

1. Unzip the StorCLI package.
2. To install the StorCLI RPM, run the `rpm -ivh <StorCLI-x.xx-x.noarch.rpm>` command.
3. To upgrade the StorCLI RPM, run the `rpm -Uvh <StorCLI-x.xx-x.noarch.rpm>` command.

10.3.3 Installing StorCLI on Ubuntu Operating Systems

To install StorCLI on the Ubuntu operating systems, perform the following steps:

Note: *Run all the commands using the super user (sudo) login.*

1. Run the command `sudo dpkg -i storcli_1.0_all.deb` for installing the debian package.
2. Run the command `dpkg -l | grep -i storcli` for verifying if the debian package was installed successfully or not.

3. To uninstall the debian package, run the command `sudo dpkg -r storcli`.

10.3.4 Installing StorCLI on VMware Operating Systems

To install StorCLI on VMware operating systems, run the following syntax from the command line:

```
esxcli software vib install -v=<path-to-vib-package>
```

Example:

```
esxcli software vib install -  
v=/vmfs/volumes/datastore1/StorCliMN/vmware-esx-StorCli-  
1.01.04.vib
```

10.3.5 Installing StorCLI on FreeBSD Operating Systems

The FreeBSD StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the Intel website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.

10.3.6 Installing StorCLI on Microsoft EFI

The EFI StorCLI binary is provided in a binary format, and no separate installation is required.

1. Copy the binary file from the CD or from the Intel website.
2. Place the binary file in the directory from which you want to run the Storage Command Line Tool, and run the tool.

10.3.7 Installing StorCLI on Solaris Operating Systems

To install StorCLI on Solaris operating systems, run the following command:

```
pkgadd -d Storcli.pkg
```

10.4 StorCLI Command Syntax

This chapter describes the StorCLI command syntax and the valid values for each parameter in the general command syntax.

Note: To get the output in JSON format, add *J* at the end of the command syntax.

Example: storcli /cx show <property1>|<property2> J

Note: JSON format output is not supported in the EFI operating system. The EFI platform ignores the *J* when it is added at the end of the command syntax.

Note: Background operations are blocked in the EFI and HII environments and these operations are resumed in the operating system environments.

The StorCLI syntax uses the following general format:

```
<[object identifier]> <verb> <[adverb | attributes | properties] >
<[key=value]>
```

The StorCLI tool supports the object identifiers listed in the following table.

Table 17. Object Identifiers in the StorCli Command Syntax

Object Identifier	Description
No object identifier specified	If there is no object identifier, the command is a system command.
/cx	This object identifier is for controller <i>x</i> .
/cx/vx	This object identifier is for a virtual drive <i>x</i> on controller <i>x</i> .
/cx/vall	This object identifier is for all virtual drives on controller <i>x</i> .
/cx/ex	This object identifier is for an enclosure <i>x</i> on controller <i>x</i> .
/cx/eall	This object identifier is for all enclosures on controller <i>x</i> .
/cx/fx	This object identifier is for a foreign configuration <i>x</i> on controller <i>x</i> .
/cx/fall	This object identifier is for all foreign configurations on controller <i>x</i> .
/cx/ex/sx	This object identifier is for the drive is slot <i>x</i> on enclosure <i>x</i> on controller <i>x</i> .
/cx/sx	This object identifier represents the drives that are directly attached to controller <i>x</i> .
/cx/ex/sall	This object identifier is for all the drives on enclosure <i>x</i> on controller <i>x</i> .
/cx/dx	This object identifier is for the drive group <i>x</i> on enclosure <i>x</i> on controller <i>x</i> .
/cx/dall	This object identifier is for the all drive groups on enclosure <i>x</i> on controller <i>x</i> .
/cx/px	This object identifier is for a PHY operation <i>x</i> on controller <i>x</i> .

Object Identifier	Description
/cx/pall	This object identifier is for all PHY operations on controller x.
/cx/bbu	This object identifier is for a bbu x on controller x.
/cx/cv	This object identifier is for a cache vault x on controller x.

Note: *If enclosures are not used to connect physical drives to the controller, you do not specify the enclosure ID in the command.*

The StorCLI tool supports the following verbs.

Table 18. Verbs in the StorCli Command Syntax

Verb	Description
add	This verb adds virtual drives, JBODs, and so on to the object identifier.
del	This verb deletes a drive, value, or property of the object identifier.
set	This verb sets a value of the object identifier.
show	This verb shows the value and properties of the object identifier.
pause	This verb pauses an ongoing operation.
resume	This verb resumes paused operation.
suspend	This verb suspends an ongoing operation. A suspended operation cannot be resumed.
compare	This verb compares an input value with a system value.
download	This verb downloads and flashes a file to the target.
start	This verb starts an operation.
flush	This verb flushes a controller cache or a drive cache.
stop	This verb stops an operation that is in progress. A stopped process cannot be resumed.
import	This verb imports the foreign configuration into the drive.
expand	This verb expands the size of the virtual drive.
insert	This verb replaces the configured drive that is identified as missing, and starts an automatic rebuild.
flasherase	This verb erases the flash memory on the controller.
transform	This verb downgrades the firmware memory on the controller.
restart	This verb restarts the controller without a system reboot.
apply	This verb applies the activation Key to a WarpDrive.

- <[adverb | attributes | properties] > – Specifies what the verb modifies or displays.
- <[key=value]> – Specifies a value, if a value is required by the command.

10.5 Working with the Storage Command Line Tool

This section describes the commands supported by the Storage Command Line Tool.

Note: The Storage Command Line Tool is not case sensitive.

Caution: The order in which you specify the command options must be the same as in the User Guide; otherwise, the command will fail.

Note: The Storage Command Line Tool does not support the Snapshot feature.

10.5.1 System Commands

10.5.1.1 System Show Commands

The Storage Command Line Tool supports the following system show commands:

```
storcli show
storcli show all
storcli show ctrlcount
storcli show help
storcli -v
```

The detailed description for each command follows.

storcli show

This command shows a summary of controller and controller-associated information for the system. The summary includes the number of controllers, the host name, the operating system information, and the overview of existing configuration.

storcli show all

This command shows the list of controllers and controller-associated information, information about the drives that need attention, and advanced software options.

storcli show ctrlcount

This command shows the number of controllers detected in the server.

storcli show help

This command shows help for all commands at the server level.

storcli -v

This command shows the version of the Storage Command Line Tool.

10.5.2 Controller Commands

Controller commands provide information and perform actions related to the specified controller, such as the /c0 controller. The Storage Command Line Tool supports the controller commands described in this section.

10.5.2.1 Show and Set Controller Properties Commands

Table 19. Controller Commands Quick Reference Table

Command	Value Range	Description
<code>show <properties></code>	See Table 20.	Shows specific controller properties.
<code>set <properties></code>	See Table 20.	Sets controller properties.
<code>show</code>	all: Shows all properties of the virtual drive. freespace: Shows the freespace in the controller. See "Controller Show Commands" on page 260.	Shows physical drive information.

This section provides command information to show and set controller properties.

Note: *You cannot set multiple properties with a single command.*

storcli /cx show <property>

This command shows the current value of the specified property on the specified controller.

General example output:

```
Status Code = 0
Status = Success
Description = None
Controller: 0
Property_name = Property_value
```

You can show the following properties using the `storcli /cx show <property1> | <property2>` command.

```
storcli /cx show abortconerror
storcli /cx show activityforlocate
storcli /cx show alarm
```

```
storcli /cx show backplane
storcli /cx show batterywarning
storcli /cx show bgirate
storcli /cx show bootwithpinnedcache
storcli /cx show cachebypass
storcli /cx show cacheflushint
storcli /cx show ccrate
storcli /cx show clusterenable
storcli /cx show coercion
storcli /cx show consistencycheck|cc
storcli /cx show copyback
storcli /cx show directpdmapping
storcli /cx show dimmerswitch|ds
storcli /cx show eccbucketleakrate
storcli /cx show eccbucketsize
storcli /cx show enableeeghsp
storcli /cx show enableesmarter
storcli /cx show enableeug
storcli /cx show exposeencldevice
storcli /cx show jbod
storcli /cx show loadbalancemode
storcli /cx show maintainpdfailhistory
storcli /cx show migraterate
storcli /cx show ncq
storcli /cx show patrolread|pr
storcli /cx show perfmode
storcli /cx show pi
storcli /cx show preventpiimport
storcli /cx show prcorrectunconfiguredareas
storcli /cx show prrate
storcli /cx show rebuildrate
storcli /cx show rehostinfo
storcli /cx show restorehotspare
storcli /cx show safeid
storcli /cx show smartpollinterval
storcli /cx show spinupdelay
storcli /cx show spinupdrivecount
storcli /cx show time
storcli /cx show usefdeonlyencrypt
storcli /cx show badblocks
```

storcli /cx set <property> = <value>

General example output:

Status Code = 0
Status = Success
Description = None

Controller 0, new Property_name = Property_value

The following commands are examples of the properties that can be set using the storcli /cx set <property>=<value> command:

```
storcli /cx set abortcconerror=<on|off>
storcli /cx set termlog[=on|off|offthisboot]
storcli /cx set activityforlocate=<on|off>
storcli /cx set alarm=<on|off|silence>
storcli /cx set backplane=<value>
storcli /cx set batterywarning=<on|off>
storcli /cx set bgirate=<value>
storcli /cx set bootwithpinnedcache=<on|off>
storcli /cx set cachebypass=<on|off>
storcli /cx set cacheflushinterval=<value>
storcli /cx set ccrate=<value>
storcli /cx set coercion=<value>
storcli /cx set consistencycheck|cc=[off|seq|conc][delay=value]
[starttime=yyyy/mm/dd hh] [excludevd=x-y,z]
storcli /cx set clusterenable=<value>
storcli /cx set copyback=<on|off> type=<smartssd|smarthdd|all>
storcli /cx set directpdmapping=<on|off>
storcli /cx set eccbucketleakrate=<value>
storcli /cx set eccbucketsize=<value>
storcli /cx set enableeeghsp=<on|off>
storcli /cx set enableesmarter=<value>
storcli /cx set enableeug=<on|off>
storcli /cx set exposeencldevice=<on|off>
storcli /cx set dimmerswitch|ds=<on|off type=1|2|3|4>
storcli /cx set foreignautoimport=<on|off>
storcli /cx set jbod=<on|off>
storcli /cx set loadbalancemode=<value>
storcli /cx set maintainpdfailhistory=<on|off>
storcli /cx set migraterate=<value>
storcli /cx set ncq=<on|off>
storcli /cx set patrolread|pr {=on mode=<auto|manual>}|{off}
storcli /cx set perfmode=<value>
storcli /cx set pi=<on|off>
storcli /cx set preventpiimport=<on|off>
storcli /cx set prcorrectunconfiguredareas=<on|off>
storcli /cx set prrate=<value>
storcli /cx set rebuildrate=<value>
```

```

storcli /cx set restorehot spare=<on|off>
storcli /cx set smartpollinterval=<value>
storcli /cx set spinupdelay=<value>
storcli /cx set spinupdrivecount=<value>
storcli /cx set stoponerror=<on|off>
storcli /cx set usefdeonlyencrypt=<on|off>
storcli /cx set time=yyyymmdd hh:mm:ss/systemtime
storcli /cx set usefdeonlyencrypt=<on|off>

```

The following table lists and describes the properties for the show and set commands.

Table 20. Properties for Show and Set Commands

Property Name	Set Command Range	Description
abortcconerror	on off	Aborts consistency check when it detects an inconsistency.
activityforlocate	on off	Enables/disables drive activity, drive activity locates function for systems without SGPIO/SES capabilities.
alarm	on off silence silence: Silences the alarm.	Enables/disables alarm on critical errors.
backplane	0: Use autodetect logic of backplanes, such as SGPIO and I ² C SEP using GPIO pins. 1: Disable autodetect SGPIO. 2: Disable I ² C SEP autodetect. 3: Disable both the autodetects.	Configures enclosure detection on a non-SES/expander backplane.
batterywarning	on off	Enables/disables battery warnings.
bgirate	0 to 100	Sets background initialization rate in percentage.
cacheflushint	0 to 255, default value 4	Sets cache flush interval in seconds.
ccrate	0 to 100	Sets consistency check rate in percentage.
coercion	0: No coercion 1: 128 MB 2: 1 GB	Sets drive capacity in coercion mode.
consistencycheck	See "Consistency Check" on page 263.	See "Consistency Check" on page 263.

Property Name	Set Command Range	Description
copyback	on off type = smartssd smarthdd all smartssd: Copy back enabled for SSD drives. smarthdd: Copy back enabled for HDD drives. all: Copy back enabled for both ssd drives and HDD drives. Example: storcli /cx set copyback=on type=all	Enables/disables copy back for drive types.
directpdmapping	on off	Enables/disables direct physical drive mapping. When enclosures are used, this feature is disabled; otherwise it should be enabled.
eccbucketleakrate	0 to 65535	Sets leak rate of the single-bit bucket in minutes (one entry removed per leak-rate).
eccbucketsize	0 to 255	Sets size of ECC single-bit-error bucket (logs event when full).
enableeeghsp	on off	Enables/disables the commissioning of otherwise incompatible global hot spare drives as Emergency Hot Spare (EHSP) drives.
enableesmarter	on off	Enables/disables the commissioning of Emergency Hot Spare (EHSP) drives for Predictive Failure (PFA) events.
enableeug	on off	Enables/disables the commissioning of Unconfigured Good drives as Emergency Hot Spare (EHSP) drives.
exposeencldevice	on off	Enables/disables device drivers to expose enclosure devices; for example, expanders, SEPs.
dimmerswitch ds	See “Dimmer Switch Commands” on page 293.	See “Dimmer Switch Commands” on page 293.
foreignautoimport	on off	Imports foreign configuration automatically, at boot.
jbod	on off	Enables/disables JBOD mode; by default, drives become system drives. Not supported by all controllers.

Property Name	Set Command Range	Description
loadbalancemode	on off	Enables/disables automatic load balancing between SAS phys or ports in a wide port configuration.
maintainpdfailhistory	on off	Maintains the physical drive fail history.
migraterate	0 to 100	Sets data migration rate in percentage.
patrolread pr	See "Patrol Read" on page 261.	See "Patrol Read" on page 261.
perfmode	0: Tuned to provide best IOPS, currently applicable to non-FastPath 1: Tuned to provide least latency, currently applicable to non-FastPath	Performance tuning setting for the controller.
pi	on off	Enables/disables data protection on the controller.
preventpiimport	on off	Enables/disables import data protection drives on the controller.
prcorrectunconfiguredareas	on off	Correct media errors during PR by writing 0s to unconfigured areas of the disk.
prrate	0 to 100	Sets patrol read rate of the virtual drives in percentage.
rebuildrate	0 to 100	Sets rebuild rate of the drive in percentage.
reconrate	0 to 100	Sets reconstruction rate for a drive in percentage.
restorehotspare	on off	Becomes a hot spare on insertion of a failed drive.
smartpollinterval	0 to 65535	Set time for polling of SMART errors in seconds.
spinupdrivecount	0 to 255	Sets number of drives that are spun up at a time.
spinupdelay	0 to 255	Sets spin-up delay between a group of drives or a set of drives, in seconds.
stoponerror	on off	Stops the Intel® RAID BIOS during POST, if any errors are encountered.
time	Valid time in <code>yyymmdd hh:mm:ss</code> format or <code>systemtime</code>	Sets the controller time to your input value or the system time (local time in 24-hour format).

Property Name	Set Command Range	Description
usefdeonlyencrypt	on off	Enables/disables FDE drive-based encryption.

10.5.2.2 Controller Show Commands

The Storage Command Line Tool supports the following show commands:

```
storcli /cx show
storcli /cx show all
storcli /cx show freespace
```

The detailed description for each command follows.

storcli /cx show

This command shows the summary of the controller information. The summary includes basic controller information, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

Input example:

```
storcli /c1 show
```

storcli /cx show all

This command shows all controller information, which includes basic controller information, bus information, controller status, advanced software options, controller policies, controller defaults, controller capabilities, scheduled tasks, miscellaneous properties, foreign configurations, drive groups, virtual drives, physical drives, enclosures, and BBU information.

Input example:

```
storcli /c0 show all
```

Note: *The PCI information displayed as a part of storcli /cx show and storcli /cx show all commands is not applicable for the FreeBSD operating system. Hence, the PCI information fields are displayed as N/A.*

storcli /cx show freespace

This command shows the usable free space in the controller.

Input example:

```
storcli /c0 show freespace
```

10.5.2.3 Controller Background Tasks Operation Commands

10.5.2.3.1 Rebuild Rate

```
storcli /cx set rebuildrate=<value>
storcli /cx show rebuildrate
```

The detailed description for each command follows.

```
storcli /cx set rebuildrate=<value>
```

This command sets the rebuild task rate of the specified controller. The input value is in percentage.

Input example:

```
storcli /c0 set rebuildrate=30
```

Note: *A high rebuild rate slows down I/O processing.*

```
storcli /cx show rebuildrate
```

This command shows the current rebuild task rate of the specified controller in percentage.

Input example:

```
storcli /c0 show rebuildrate
```

10.5.2.3.2 Patrol Read

The Storage Command Line Tool supports the following patrol read commands:

```
storcli /cx resume patrolread
storcli /cx set patrolread ={{on mode=<auto|manual>}}|{off}}
storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>]
[maxconcurrentpd=<value>] [includessds=<on|off>]
[uncfgareas=<on|off>]
storcli /cx set patrolread delay=<value>
storcli /cx show patrolread
storcli /cx start patrolread
storcli /cx stop patrolread
storcli /cx suspend patrolread
```

Note: *A patrol read operation is scheduled for all the physical drives of the controller.*

The detailed description for each command follows.

```
storcli /cx resume patrolread
```

This command resumes a suspended patrol read operation.

Input example:

```
storcli /c0 resume patrolread
```

```
storcli /cx set patrolread {=on mode=<auto|manual>}|{off}
```

This command turns the patrol read scheduling on and sets the mode of the patrol read to automatic or manual.

Input example:

```
storcli /c0 set patrolread=on mode=manual
```

```
storcli /cx set patrolread [starttime=<yyyy/mm/dd hh>]  
[maxconcurrentpd=<value>] [includessds=<on|off>]  
[uncfgareas=on|off]
```

This command schedules a patrol read operation. You can use the following options for patrol read command.

Table 21. Set Patrolread Input Options

Option	Value Range	Description
starttime	A valid date and hour in 24-hour format	Sets the start time in yyyy/mm/dd hh format.
maxconcurrentpd	Valid number of physical drives present	Sets the number of physical drives that can be patrol read at a single time.
includessds	—	Include SSDs in the patrol read.
uncfgareas	—	Include the areas not configured in the patrol read.

Note: Controller time is taken as a reference for scheduling a patrol read operation.

Input example:

```
storcli /c0 set patrolread=on starttime=2012/02/21 00
```

```
storcli /cx set patrolread [delay=<value>]
```

This command delays the scheduled patrol read in hours.

Input example:

```
storcli /c0 set patrolread delay=30
```

```
storcli /cx show patrolRead
```

This command shows the progress on the current patrol read in percentage.

Input example:

```
storcli /c0 show patrolread
```

```
storcli /cx start patrolread
```

This command starts the patrol read operation. This command starts a patrol read immediately.

Input example:

```
storcli /c0 start patrolread
```

```
storcli /cx stop patrolread
```

This command stops a running patrol read operation.

Input example:

```
storcli /c0 stop patrolread
```

Note: *You cannot resume a stopped patrol read.*

```
storcli /cx suspend patrolread
```

This command pauses a running patrol read operation.

Input example:

```
storcli /c0 suspend patrolread
```

Note: *You can run this command only when a patrol read operation is running on the controller.*

10.5.2.3.3 Consistency Check

The Storage Command Line Tool supports the following commands to schedule, perform, and view the status of a consistency check (CC) operation:

```
storcli /cx set consistencycheck|cc=[off|seq|conc][delay=value]
starttime=yyyy/mm/dd hh [excludevd=x-y,z]
storcli /cx show cc
storcli /cx show ccrate
```

The detailed description for each command follows.

```
storcli /cx set consistencycheck|cc=[off|seq|conc][delay=value]
starttime=yyyy/mm/dd hh [excludevd=x-y,z]
```

This command schedules a consistency check (CC) operation. You can use the following options with the consistency check command.

Table 22. Set CC Input Options

Option	Value Range	Description
cc	seq: Sequential mode. conc: Concurrent mode. off: Turns off the consistency check.	Sets CC to either sequential mode, or concurrent mode, or turns off the CC. Note: <i>The concurrent mode slows I/O processing.</i>

Option	Value Range	Description
delay	-1 and any integer value.	Delays a scheduled consistency check. The value is in hours. A value of 0 makes the CC runs continuously with no delay (in a loop). Note: Only scheduled consistency checks can be delayed.
starttime	A valid date and hour in 24-hour format.	Start time of a consistency check is in yyyy/mm/dd hh format.
excludevd	The range should be less than the number of virtual drives.	Excludes virtual drives from the consistency checks. To exclude particular virtual drives, you can provide list of virtual drive names (Vx,Vy ... format) or the range of virtual drives that you want to exclude from a consistency check (Vx-Vy format). If this option is not specified in the command, no virtual drives are excluded.

Input example:

```
storcli /c0 set CC=on starttime=2012/02/21 00 excludevd v0-v3
```

```
storcli /cx show cc
```

This command shows the consistency check schedule properties for a controller.

Input example:

```
storcli /c0 show cc
```

```
storcli /cx show ccrate
```

This command checks the status of a consistency check operation. The CC rate appears in percentage.

Input example:

```
storcli /c0 show ccrate
```

Note: A high CC rate slows I/O processing.

10.5.2.4 Premium Feature Key Commands

The Storage Command Line Tool supports the following commands for premium feature keys:

```
storcli /cx set advancedsoftwareoptions(aso) key=<value> [preview]
storcli /cx aso
[transfertovault][rehostcomplete][deactivatetrialkey]
storcli /cx show safeid
```

The detailed description for the command follows.

```
storcli /cx set advancedsoftwareoptions(aso) key=<value> [preview]
```

This command activates advanced software options (ASO) for a controller. You can use the following options with the advanced software options command.

Table 23. Set Advanced Software Options Input Options

Option	Value Range	Description
key	40 alpha numeric characters.	Key to activate ASO on the controller. Note: After they are activated, ASOs cannot be removed from the controller.
deactivatetrialsecuritykey	—	Deactivates the trial key applied on the specified controller.
rehostcomplete	—	Enables rehosting on the specified controller.
transfertovault	—	Transfers the ASO key to the vault and disables the ASO.

Input example:

```
storcli /c0 set Aso key=LSI0000
```

```
storcli /cx show safeid
```

This command shows the Safe ID of the specified controller.

Input example:

```
storcli /c0 show safeid
```

10.5.2.5 Controller Security Commands

The Storage Command Line Tool supports the following controller security commands:

```
storcli /cx show securitykey keyid
storcli /cx compare securitykey=ssssss
storcli /cx delete securitykey
storcli /cx set securitykey keyid=kkkk
storcli /cx set securitykey=sssss [passphrase=sssss][keyid=sssss]
storcli /cx set securitykey=sssss oldsecuritykey=sssss
[passphrase=sssss][keyid=sssss]
```

The detailed description for each command follows.

```
storcli /cx show securitykey keyid
```

This command shows the security key on the controller.

Input example:

```
storcli /c0 show securityKey keyid
```

```
storcli /cx compare securitykey=ssssss
```

This command compares and verifies the security key of the controller.

```
storcli /cx delete securitykey
```

This command deletes the security key of the controller.

Input example:

```
storcli /c0 delete securitykey
```

```
storcli /cx set securitykey keyId=kkkk
```

This command sets the key ID for the controller. The key ID is unique for every controller.

```
storcli /cx set securitykey=sssss [passphrase=sssss][keyid=sssss]
```

This command sets the security key for the controller. You can use the following options with the set security key command.

Table 24. Set Security Key Input Options

Option	Value Range	Description
passphrase	Should have a combination of numbers, upper case letters, lower case letters, and special characters. Minimum of 8 characters and maximum of 32 characters.	String that is linked to the controller and is used in the next bootup to encrypt the lock key. If the passphrase is not set, the controller generates it by default.
keyid	—	Unique ID set for different controllers to help you specify a passphrase to a specific controller.

Input example:

```
storcli /c0 set securitykey=Lsi@12345 passphrase=Lsi@123456  
keyid=1
```

```
storcli /cx set securitykey=sssss oldsecuritykey=sssss  
[passphrase=sssss][keyid=sssss]
```

This command changes the security key for the controller.

Input example:

```
storcli /c0 set securitykey=Lsi@12345 oldsecuritykey=pass123  
passphrase=Lsi@123456 keyid=1
```

10.5.2.6 Flashing Controller Firmware Command

Note: The Flashing Controller Firmware command is not supported in Embedded Intel® RAID.

The following command flash the controller firmware:

```
storcli /cx download file=filepath [fwtype=<value>] [nosigchk]
[noverchk] [resetnow]
```

This command flashes the firmware to the specified adapter from the given file location (filepath is the absolute file path). You can use the following options when you flash the firmware.

Table 25. Flashing Controller Firmware Input Options

Option	Value Range	Description
nosigchk	—	The application flashes the firmware even if the check word on the file does not match the required check word for the controller. Note: You can damage the controller if a corrupted image is flashed using this option.
noverchk	—	The application flashes the controller firmware without checking the version of the firmware image.
fwtype	0: Application 1: TMMC	The firmware type to be downloaded. The application downloads the firmware for the controller. The TMMC downloads the firmware for the TMMC battery only. Default is 0 (application).
resetnow		Invokes online firmware update on the controller; you do not need to reboot the controller to make the update effective.

10.5.2.7 Controller Cache Command

The following command flushes the controller cache:

```
storcli /cx flush|flushcache
```

This command flushes the controller cache.

Input example:

```
storcli /c0 flushcache
```

10.5.3 Drive Commands

This section describes the drive commands, which provide information and perform actions related to physical drives. The following table describes frequently used virtual drive commands.

Table 26. Physical Drives Commands Quick Reference Table

Command	Value Range	Description
set	missing: Sets the drive status as missing. good: Sets the drive status to unconfigured good. offline: Sets the drive status to offline. online: Sets the drive status to online.	Sets physical drive properties.
show	all: Shows all properties of the physical drive. See “Drive Show Commands” on page 268.	Shows virtual drive information.

10.5.3.1 Drive Show Commands

The Storage Command Line Tool supports the following drive show commands:

```
storcli /cx[/ex]/sx show
storcli /cx[/eall]/sall show
storcli /cx[/ex]/sx|sall show all
```

Note: *If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command. If no enclosures are used, you must specify the controller ID and slot ID.*

The detailed description for each command follows.

storcli /cx[/ex]/sx show

This command shows the summary of the physical drive for a specified slot in the controller.

Input example:

```
storcli /c0/e0/s4,5 show
```

storcli /cx[/eall]/sall show

This command shows the summary information for all the enclosures and physical drives connected to the controller.

Input example:

```
storcli /c0/eall/sall show
```

```
storcli /cx[/ex]/sx|sall show all
```

This command shows all information of a physical drive for the specified slot in the controller. If you use the `all` option, the command shows information for all slots on the controller. `x` stands for a number, a list of numbers, a range of numbers, or all numbers.

Input examples:

```
storcli /c0/e3/s0-3 show all
storcli /c0/e35/sall show all
```

Note: *The storcli /cx/sx show all command shows tape drives information.*

10.5.3.2 Missing Drives Commands

The Storage Command Line Tool supports the following commands to mark and replace missing physical drives:

```
storcli /cx[/ex]/sx insert dg=A array=B row=C
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set offline
storcli /cx/dall
```

The detailed description for each command follows.

```
storcli /cx[/ex]/sx insert dg=A array=B row=C
```

This command replaces the configured drive that is identified as missing, and then starts an automatic rebuild.

Input example:

```
storcli /c0/e25/s3 insert dg=0 array=2 row=1
```

```
storcli /cx[/ex]/sx set missing
```

This command marks a drive as missing.

Input example:

```
storcli /c0/s4 set missing
```

```
storcli /cx/dall
```

This command is used to find the missing drives.

```
storcli /cx[/ex]/sx set offline
```

This command marks the drive in an array as offline.

Note: *To set a drive that is part of an array as missing, first set it as offline. After the drive is set to offline, you can then set the drive to missing.*

10.5.3.3 Set Drive State Commands

The Storage Command Line Tool supports the following commands to set the status of physical drives:

```
storcli /cx[/ex]/sx set jbod
storcli /cx[/ex]/sx set good [force]
storcli /cx[/ex]/sx set offline
storcli /cx[/ex]/sx set online
storcli /cx[/ex]/sx set missing
storcli /cx[/ex]/sx set bootdrive=<on|off>
```

The detailed description for each command follows.

storcli /cx[/ex]/sx set jbod

This command sets the drive state to JBOD.

Input example:

```
storcli /c1/e56/s3 set jbod
```

storcli /cx[/ex]/sx set good [force]

This command changes the drive state to unconfigured good. If the drive has the operating system in it, use the force option.

Input example:

```
storcli /c1/e56/s3 set good
```

storcli /cx[/ex]/sx set offline

This command changes the drive state to offline.

Input example:

```
storcli /c1/e56/s3 set offline
```

storcli /cx[/ex]/sx set online

This command changes the drive state to online.

Input example:

```
storcli /c1/e56/s3 set online
```

storcli /cx[/ex]/sx set missing

This command marks a drive as missing.

Input example:

```
storcli /c1/e56/s3 set missing
```

storcli /cx[/ex]/sx set bootmode=<on|off>

This command sets or unsets a physical drive as a boot drive.

Input example:

```
storcli /c1/e56/s3 set bootmode=on
```

10.5.3.4 Drive Initialization Commands

When you initialize drives, all the data from the drives is cleared. The Storage Command Line Tool supports the following commands to initialize drives:

```
storcli /cx[/ex]/sx show initialization
storcli /cx[/ex]/sx start initialization
storcli /cx[/ex]/sx stop initialization
```

The detailed description for each command follows.

storcli /cx[/ex]/sx show initialization

This command shows the current progress of the initialization progress in percentage.

Input example:

```
storcli /c0/e31/s4 show initialization
```

storcli /cx[/ex]/sx start initialization

This command starts the initialization process on a drive.

Input example:

```
storcli /c0/e31/s4 start initialization
```

storcli /cx[/ex]/sx stop initialization

This command stops an initialization process running on the specified drive. A stopped initialization process cannot be resumed.

Input example:

```
storcli /c0/e56/s1 stop initialization
```

10.5.3.5 Drive Firmware Download Command

The Storage Command Line Tool supports the following command to download drive firmware:

storcli /cx[/ex]/sx download src=*filepath* [*satabridge*]

This command flashes the firmware with the specified file. The satabridge option lets you download the SATA bridge firmware in online mode.

Input example:

```
storcli /c0/e56/s1 download src=c:\file1.bin
```

10.5.3.6 Locate Drives Commands

The Storage Command Line Tool supports the following commands to locate a drive and activate the physical disk activity LED:

```
storcli /cx[/ex]/sx start locate
storcli /cx[/ex]/sx stop locate
```

The detailed description for each command follows.

storcli /cx[/ex]/sx start locate

This command locates a drive and activates the drive's LED.

Input example:

```
storcli /c0/e56/s1 start locate
```

storcli /cx[/ex]/sx stop locate

This command stops a locate operation and deactivates the drive's LED.

Input example:

```
storcli /c0/e56/s1 stop locate
```

10.5.3.7 Prepare to Remove Drives Commands

The Storage CLI supports the following commands to prepare the physical drive for removal:

```
storcli /cx[/ex]/sx spindown
storcli /cx[/ex]/sx spinup
```

The detailed description for each command follows.

storcli /cx[/ex]/sx spindown

This command spins down an unconfigured drive and prepares it for removal. The drive state is unaffiliated and it is marked offline.

Input example:

```
storcli /cx/e34/s4 spindown
```

storcli /cx[/ex]/sx spinup

This command spins up a spun-down drive and the drive state is unconfigured good.

Input example:

```
storcli /cx/e34/s4 spinup
```

10.5.3.8 Drive Security Command

The Storage Command Line Tool supports the following drive security command:

```
storcli /cx[/ex]/sx show securitykey keyid
```

```
storcli /cx[/ex]/sx show securitykey keyid
```

This command shows the security key for secured physical drives.

Input example:

```
storcli /c0/e252/s1 show SecurityKey keyid
```

10.5.3.9 Drive Secure Erase Commands

The Storage Command Line Tool supports the following drive erase commands:

```
storcli /cx[/ex]/sx secureerase [force]
```

```
storcli /cx[/ex]/sx show erase
```

```
storcli /cx[/ex]/sx start erase [simple|normal|thorough]  
[erasepatternA=<value1>] [erasepatternB=<value2>]
```

```
storcli /cx[/ex]/sx stop erase
```

The detailed description for each command follows.

```
storcli /cx[/ex]/sx secureerase [force]
```

This command erases the drive's security configuration and securely erases data on a drive. You can use the `force` option as a confirmation to erase the data on the drive and the security information.

Input example:

```
storcli /c0/e25/s1 secureerase
```

Note: *This command deletes data on the drive and the security configuration and this data is no longer accessible. This command is used for SED drives only.*

```
storcli /cx[/ex]/sx show erase
```

This command provides the status of erase operation on non-SEDs.

Input example:

```
storcli /c0/e25/s1 show erase
```

```
storcli /cx[/ex]/sx start erase [simple|normal|thorough]  
[erasepatternA=<val1>] [erasepatternB=<val2>]
```

This command securely erases non-SED drives. The drive is written with erase patterns to ensure that the data is securely erased. You can use the following options with the start erase command.

Table 27. Drive Erase Command Options

Option	Value Range	Description
erase	simple: Single pass, single pattern write normal: Three pass, three pattern write thorough: Nine pass, repeats the normal write three times	Secure erase type.
erasepatternA	8-bit value	Erase pattern A to overwrite the data.
erasepatternB	8-bit value	Erase pattern B to overwrite the data.

Input example:

```
storcli /c0/e25/s1 start erase thorough erasepatternA=10010011
erasepatternB=11110000
```

10.5.3.10 Rebuild Drives Commands

The following commands rebuild drives in the Storage Command Line Tool:

```
storcli /cx[/ex]/sx pause rebuild
storcli /cx[/ex]/sx resume rebuild
storcli /cx[/ex]/sx show rebuild
storcli /cx[/ex]/sx start rebuild
storcli /cx[/ex]/sx stop rebuild
```

Note: *If enclosures are used to connect physical drives to the controller, specify the enclosure ID in the command.*

The detailed description for each command follows.

storcli /cx[/ex]/sx pause rebuild

This command pauses an ongoing rebuild process. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 pause rebuild
```

storcli /cx[/ex]/sx resume rebuild

This command resumes a paused rebuild process. You can run this command only when a paused rebuild process for the drive exists.

Input example:

```
storcli /c0/s4 resume rebuild
```

```
storcli /cx[/ex]/sx show rebuild
```

This command shows the progress of the rebuild process in percentage.

Input example:

```
storcli /c0/s5 show rebuild
```

```
storcli /cx[/ex]/sx start rebuild
```

This command starts a rebuild operation for a drive.

Input example:

```
storcli /c0/s4 start rebuild
```

```
storcli /cx[/ex]/sx stop rebuild
```

This command stops a rebuild operation. You can run this command only for a drive that is currently rebuilt.

Input example:

```
storcli /c0/s4 stop rebuild
```

10.5.3.11 Drive Copyback Commands

The Storage Command Line Tool supports the following commands for drive copyback:

```
storcli /cx[/ex]/sx pause copyback
```

```
storcli /cx[/ex]/sx resume copyback
```

```
storcli /cx[/ex]/sx show copyback
```

```
storcli /cx[/ex]/sx start copyback target=eid:sid
```

```
storcli /cx[/ex]/sx stop copyback
```

The detailed description for each command follows.

Note: *In the copyback commands, cx[/ex]/sx indicates the source drive and eid:sid indicates the target drive.*

Note: *When a copyback operation is enabled, the alarm continues to beep even after a rebuild is complete; the alarm stops beeping only when the copyback operation is completed.*

```
storcli /cx[/ex]/sx pause copyback
```

This command pauses a copyback operation. You can run this command only when there is a copyback operation running.

Input example:

```
storcli /c0/e25/s4 pause copyback
```

```
storcli /cx[/ex]/sx resume copyback
```

This command resumes a paused copyback operation. You can run this command only when there is a paused copyback process for the drive.

Input example:

```
storcli /c0/e25/s4 resume copyback
```

```
storcli /cx[/ex]/sx show copyback
```

This command shows the progress of the copyback operation in percentage.

Input example:

```
storcli /c0/e25/s4 show copyback
```

```
storcli /cx[/ex]/sx start copyback target=eid:sid
```

This command starts a copyback operation for a drive.

Input example:

```
storcli /c0/e25/s4 start copyback target=25:8
```

```
storcli /cx[/ex]/sx stop copyback
```

This command stops a copyback operation. You can run this command only on drives that have the copyback operation running.

Input example:

```
storcli /c0/e25/s4 stop copyback
```

Note: *A stopped rebuild process cannot be resumed.*

10.5.3.12 Hot Spare Drive Commands

The following commands create and delete hot spare drives:

```
storcli /cx[/ex]/sx add hotsparedrive  
{dgs=<n|0,1,2...>}[enclaffinity][nonrevertible]  
storcli /cx[/ex]/sx delete hotsparedrive
```

Note: *If enclosures are used to connect the physical drives to the controller, specify the enclosure ID in the command.*

The detailed description for each command follows.

```
storcli /cx[/ex]/sx add hotsparedrive [{dgs=<n|0,1,2...>}]  
[enclaffinity][nonrevertible]
```

This command creates a hot spare drive. You can use the following options to create a hot spare drive.

Table 28. Add Hotsparedrive Input Options

Option	Value Range	Description
dgs	Valid drive group number	Specifies the drive group to which the hot spare drive is dedicated.

Option	Value Range	Description
enclaffinity	Valid enclosure number	Specifies the enclosure with which the hot spare is associated. If this option is specified, affinity is set; if it is not specified, there is no affinity. Affinity cannot be removed after it is set for a hot spare drive.
nonrevertible	—	Sets the drive as a nonrevertible hot spare.

Input example:

```
storcli /c0/e3/s4,5 add hotsparedrive
```

This command sets the drives /c0/e3/s4,5 as Global Hot spare.

Input example:

```
storcli /c0/e3/s6,8 add hotsparedrive dgs=0,1
```

This command sets /c0/e3/s6,8 as Dedicated Hot spare for disk groups 0,1.

```
storcli /cx/[ex]/sx delete hotsparedrive
```

This command deletes a hot spare drive.

Input example:

```
storcli /c0/e3/s4,5 delete hotsparedrive
```

10.5.4 Virtual Drive Commands

The Storage Command Line Tool supports the following virtual drive commands. The following table describes frequently used virtual drive commands.

Table 29. Virtual Drives Commands Quick Reference Table

Command	Value Range	Description
add	See Table 30 and Table 31 .	Creates virtual drives.
delete	cc or cachecade: Deletes CacheCade* virtual drives. force: Deletes the virtual drive where operating system is present.	Deletes a virtual drive.
set	See Table 30 , Table 31 , and “ Change Virtual Properties Commands ” on page 283 .	Sets virtual drive properties.
show	all: Shows all properties of the virtual drive. cc: Shows properties of CacheCade virtual drives. See “ Virtual Drive Show Commands ” on page 282 .	Shows virtual drive information.

10.5.4.1 Add Virtual Drives Commands

The Storage Command Line Tool supports the following commands to add virtual drives:

```
storcli /cx add vd
type=raid[0|1|5|6|00|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,...|all]
[name=<VDNAME1>,...] drives=e:s|e:s-x,y|e:s-x,y,z
[PDperArray=x][SED] [pdcache=on|off|default][pi]
[DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|wb]
[nora|ra] [direct|cached] [CachedBadBBU|NoCachedBadBBU][cachevd]
[Strip=<8|16|32|64|128|256|1024>] [AfterVd=X] [Spares =
[e:]s|[e:]s-x|[e:]s-x,y] [force][ExclusiveAccess]
```

```
storcli /cx add vd each type=raid0 [name=<VDNAME1>,...]
[drives=e:s|e:s-x|e:s-x,y] [SED] [pdcache=on|off|default][pi]
[DimmerSwitch(ds)=default|automatic(auto)|
none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|wb]
[nora|ra] [direct|cached]
[CachedBadBBU|NoCachedBadBBU][Strip=<8|16|32|64|128|256|1024>][Ex
clusiveAccess]
```

```
storcli /cx add VD cachecade|nytrocache Type = [raid|r][0,1,10,
R1EC] drives = [e:]s|[e:]s-x|[e:]s-x,y [WT|WB] [assignvds = 0,1,2
[BOOTVOLSIZE=x]
```

This command creates a RAID configuration. You can use the following options to create the RAID volume.

Note: * indicates default values.

The detailed description for each command follows.

```
storcli /cx add vd
type=raid[0|1|5|6|00|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,...|*all]
[name=<VDNAME1>,...] drives=e:s|e:s-x|e:s-x,y;e:s-x,y,z
[PDperArray=x][SED] [pdcache=on|off|*default][pi]
[DimmerSwitch(ds)=default|automatic(auto)|
*none|maximum(max)|MaximumWithoutCaching(maxnocache)][cachevd][Ex
clusiveAccess|SharedAccess*]** [wt|*wb] [nora|*ra]
[*direct|cached] [CachedBadBBU|*NoCachedBadBBU]
[Strip=<8|16|32|64|128|256|1024>] [AfterVd=X] [Spares =
[e:]s|[e:]s-x|[e:]s-x,y] [force]
```

Table 30. Add RAID Configuration Input Options

Option	Value Range	Description
type	RAID [0 1 5 6 00 10 50 60]	Sets the RAID type of the configuration.
size	Maximum size based on the physical drives and RAID level	Sets the size of each virtual drive. The default value is for the capacity of all referenced disks.

Option	Value Range	Description
name	15 characters of length	Specifies the drive name for each virtual drive.
drives	Valid enclosure number and valid slot numbers for the enclosure	In <code>e:s e:s-x e:s-x,y</code> : <code>e</code> specifies the enclosure ID. <code>s</code> represents the slot in the enclosure. <code>e:s-x</code> is the range convention used to represent slots <code>s</code> to <code>x</code> in the enclosure <code>e</code> .
pdperarray	1-16	Specifies the number of physical drives per array. The default value is automatically chosen.
sed	—	Creates security-enabled drives.
pdcache	on off default	Enables or disables PD cache.
pi	—	Enables protection information.
dimmerswitch	default: Logical device uses controller default power-saving policy. automatic (auto): Logical device power savings are managed by firmware. none: No power-saving policy. maximum (max): Logical device uses maximum power savings. MaximumWithoutCaching (maxnocache): Logical device does not cache write to maximize power savings.	Specifies the power-saving policy. Sets to default automatically.
direct cached	cached: Cached I/O direct: Direct I/O	Sets the logical drive cache policy. Direct I/O is the default.
wt wb	wt: Write through wb: Write back	Enables write through. Write back is the default.
nora ra	ra: Read ahead nora: No read ahead	Disables read ahead. Enabled is the default.
cachedbadbbu nocachedbadbbu	cachedbadbbu: Enables bad BBU caching. nocachedbadbbu: Disables bad BBU caching.	Enables caching when BBU is not functioning. Disabled is the default.
cachevd	—	Enables SSD caching on the created virtual drive.

Option	Value Range	Description
strip	8, 16, 32, 64, 128, 256, 512, 1024	Sets the strip size for the RAID configuration.
aftervd	Valid virtual drive number	Creates the VD in the adjacent free slot next to the specified VD.
spares	Number of spare physical drives present	Specifies the physical drives that are to be assigned to a disk group for spares.
force	—	Forces a security-capable physical drive to be added to a drive group without security.

Input example:

```
storcli /c0 add vd type=raid10 size=2gb,3gb,4gb
names=tmp1,tmp2,tmp3 drives=252:2-3,5,7 pdperarray=2
```

```
storcli /cx add vd cc|cachecade type=raid[0,1,10]
drives=[e:]s|[e:]s-x|[e:]s-x,y [[wt|*wb] ] [assignvds=0,1,2]
```

This command creates CacheCade virtual drives and associates existing virtual drives to CacheCade virtual drives. You can use the following options to create the CacheCade virtual drive.

Table 31. Add RAID Configuration Input Options

Option	Value Range	Description
cachecade	—	Creates a CacheCade virtual drive.
type	0, 1, 10	Sets the RAID type of the CacheCade virtual drive.
drives	Valid enclosure number and valid slot number	See the <code>drives</code> row in the previous table for format.
wt *wb	wt: Enables write through. wb: Enables write back.	Enables or disables write cache.
assignvds	Valid virtual drive number (0 to 63)	Specifies the list of virtual drives associated with the new CacheCade virtual drives.

Input example:

```
storcli /c0 add vd type=raid10 size=2gb,3gb,4gb
names=tmp1,tmp2,tmp3 drives=252:2-3, 7
```

10.5.4.2 Delete Virtual Drives Commands

The Storage Command Line Tool supports the following virtual drive delete commands:

```
storcli /cx/vx|vall del
storcli /cx/vx|vall del cachecade
storcli /cx/vx|vall del force
```

Note: *If the virtual drive has user data, you must use the force option to delete the virtual drive. A virtual drive with a valid master boot record (MBR) and a partition table is considered to contain user data.*

If you delete a virtual drive with a valid MBR without erasing the data and then create a new virtual drive using the same set of physical drives and the same RAID level as the deleted virtual drive, the old uneraser MBR still exists at block0 of the new virtual drive, which makes it a virtual drive with valid user data. Therefore, you must provide the `force` option to delete this newly created virtual drive.

The detailed description for each command follows.

```
storcli /cx/vx|vall del
```

This command deletes a particular virtual drive or, when the `vall` option is used, all the virtual drives on the controller are deleted.

Input example:

```
storcli /c0/v2 del
```

Note: *This command deletes virtual drives. Data located on these drives will no longer be accessible.*

```
storcli /cx/vx|vall del cachecade
```

This command deletes a specific CacheCade virtual drive on a controller, or all the CacheCade configuration for a controller.

Input example:

```
storcli /c0/vall del cachecade
```

Note: *This command deletes virtual drives. Data located on these drives will no longer be accessible.*

```
storcli /cx/vx|vall del force
```

This command deletes a virtual drive only after the cache flush is completed. With the `force` option, the command deletes a virtual drive without waiting for the cache flush to complete.

Input example:

```
storcli /c0/v2 del force
```

Note: *This command deletes the virtual drive where the operating system is present. Data located on these drives and the operating system of the drive will no longer be accessible.*

10.5.4.3 Virtual Drive Show Commands

The Storage Command Line Tool supports the following virtual drive show commands:

```
storcli /cx/vx show  
storcli /cx/vx show all
```

The detailed description for each command follows.

storcli /cx/vx show

This command shows the summary of the virtual drive information.

Input example:

```
storcli /c0/v0 show
```

storcli /cx/vx show all

This command shows all virtual drive information, which includes virtual drive information, physical drives used for the virtual drives, and virtual drive properties.

Input example:

```
storcli /c0/v0 show all
```

10.5.4.4 Preserved Cache Commands

If a virtual drive becomes offline or is deleted because of missing physical disks, the controller preserves the dirty cache from the virtual disk. The Storage Command Line Tool supports the following commands for preserved cache:

```
storcli /cx/vx delete preservedCache [force]  
storcli /cx show preservedCache
```

The detailed description for each command follows.

storcli /cx/vx delete preservedcache

This command deletes the preserved cache for a particular virtual drive on the controller in missing state. Use the `force` option to delete the preserved cache of a virtual drive in offline state.

Input example:

```
storcli /c0/v1 delete preservedcache
```

storcli /cx show preservedCache

This command shows the virtual drive that has preserved cache and whether the virtual drive is offline or missing.

Input example:

```
storcli /c0 show preservedCache
```

10.5.4.5 Change Virtual Properties Commands

The Storage Command Line Tool supports the following commands to change virtual drive properties:

```
storcli /cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd>
storcli /cx/vx set iopolicy=<cached|direct>
storcli /cx/vx set name=<namestring>
storcli /cx/vx set pdcache=<on|off|default>
storcli /cx/vx set rdcache=<ra|nora>
storcli /cx/vx|vall set ssdcaching=<on|off>
storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess
storcli /cx/vx set wrcache=<wt|wb|awb>
storcli /cx/vx set emulationType=0|1
storcli /cx/vx set ds=Default|Auto|None|Max|MaxNoCache
storcli /cx/vx set autobgi=On|Off
storcli /cx/vx set pi=Off
storcli /cx/vx set bootdrive=<On|Off>
```

The detailed description for each command follows.

storcli /cx/vx set accesspolicy=<rw|ro|blocked|rmvblkd>

This command sets the access policy on a virtual drive to read write, read only, or blocked or rmvblkd (remove blocked).

Input example:

```
storcli /c0/v0 set accesspolicy=rw
```

storcli /cx/vx set iopolicy=<cached|direct>

This command sets the I/O policy on a virtual drive to cached I/O or direct I/O.

Input example:

```
storcli /c0/v0 set iopolicy=cached
```

storcli /cx/vx set name=<namestring>

This command names a virtual drive. The name is restricted to 15 characters.

Input example:

```
storcli /c1/v0 set name=testdrive123
```

storcli /cx/vx set pdcache=<on|off|default>

This command sets the current disk cache policy on a virtual drive to on, off, or default setting.

Input example:

```
storcli /c0/v0 set pdcache=on
```

```
storcli /cx/vx set rdcache=<ra|nora>
```

This command sets the read cache policy on a virtual drive to read ahead, no read ahead, or adaptive read ahead.

Input example:

```
storcli /c0/v0 set rdcache=nora
```

```
storcli /cx/vx|vall set ssdcaching=<on|off>
```

This command assigns CacheCade virtual drives. If ssdcaching=off, the CacheCade virtual drive is removed.

Input example:

```
storcli /c0/v0 set ssdcaching=on
```

```
storcli /cx/vx|vall set HostAccess=ExclusiveAccess|SharedAccess
```

This command sets the host access policy for the virtual drive. When the host access policy is exclusive access, a server has exclusive access to the virtual drive. The virtual drive cannot be shared between servers. If the host policy is shared access, the virtual drive can be shared between servers.

Input example:

```
storcli /c0/v0 set HostAccess=ExclusiveAccess
```

```
storcli/cx/vx set wrcache=<wt|wb|awb>
```

This command sets the write cache policy on a virtual drive to write back, write through, or always write back.

Input example:

```
storcli /c0/v0 set wrcache=wt
```

10.5.4.6 Virtual Drive Initialization Commands

The Storage Command Line Tool supports the following commands to initialize virtual drives:

```
storcli /cx/vx show init
```

```
storcli /cx/vx start init [full][Force]
```

```
storcli /cx/vx stop init
```

Note: *If the virtual drive has user data, you must use the `force` option to initialize the virtual drive.*

A virtual drive with a valid MBR and partition table is considered to contain user data.

The detailed description for each command follows.

```
storcli /cx/vx show init
```

This command shows the initialization progress of a virtual drive in percentage.

Input example:

```
storcli /c0/v2 show init
```

```
storcli /cx/vx start init [full]
```

This command starts the initialization of a virtual drive. The default initialization type is fast initialization. If the `fulloption` is specified, full initialization of the virtual drive starts.

Input example:

```
storcli /cx/vx start init [full]
```

```
storcli /cx/vx stop init
```

This command stops the initialization of a virtual drive. A stopped initialization cannot be resumed.

Input example:

```
storcli /c0/v0 stop init
```

10.5.4.7 Virtual Drive Erase Commands

The Storage Command Line Tool supports the following commands to erase virtual drives:

```
storcli /cx/vx erase
```

```
storcli /cx/vx show erase
```

The detailed description for each command follows.

```
storcli /cx/vx erase
```

This command erases the data on the virtual drive.

Input example:

```
storcli /c0/v0 erase
```

```
storcli /cx/vx show erase
```

This command shows the status of the erase operation on the virtual drive.

Input example:

```
storcli /c0/v0 show erase
```

10.5.4.8 Virtual Drive Migration Commands

Note: *The virtual drive migration commands are not supported in Embedded Intel® RAID.*

The Storage Command Line Tool supports the following commands for virtual drive migration (reconstruction):

```
storcli /cx/vx show migrate
```

```
storcli /cx/vx start migrate <type=raidlevel> [option=<add|remove>  
disk=<e1/s1,e2/s2 ...> ]
```

The detailed description for each command follows.

```
storcli /cx/vx show migrate
```

This command shows the progress of the virtual drive migrate operation in percentage.

Input example:

```
storcli /c0/v0 show migrate
```

```
storcli /cx/vx start migrate <type=raidlevel> [option=<add |
remove> disk=<e1:s1,e2:s2 ...> ]
```

This command starts the reconstruction on a virtual drive to the specified RAID level by adding or removing disks from the existing virtual drive. You can use the following options with the start migrate command.

Table 32. Virtual Drive Migration Command Options

Option	Value Range	Description
<code>type=RAID level</code>	RAID [0 1 5 6]	The RAID level to which the virtual drive must be migrated.
<code>[option=<add remove> disk=<e1:s1,e2:s2, ...>]</code>	<p><code>add</code>: Adds disks to the virtual drive and starts reconstruction.</p> <p><code>remove</code>: Removes disks from the virtual drive and starts reconstruction.</p> <p><code>disk</code>: The enclosure number and the slot number of the disks to be added to the virtual drive.</p>	Adds or removes disks from the virtual drive.

Virtual drive migration can be done between the following RAID levels.

Table 33. Virtual Drive Migration Table

Initial RAID Level	Migrated RAID Level
RAID 0	RAID 1
RAID 0	RAID 5
RAID 0	RAID 6
RAID 1	RAID 0
RAID 1	RAID 5
RAID 1	RAID 6
RAID 5	RAID 0

Initial RAID Level	Migrated RAID Level
RAID 5	RAID 6
RAID 6	RAID 0
RAID 6	RAID 5

Input example:

```
storcli /c0/v3 start migrate type=r5 option=add disk=e5:s2,e5:s3
```

10.5.4.9 Virtual Drive Consistency Check Commands

The Storage Command Line Tool supports the following commands for virtual drive consistency checks:

```
storcli /cx/vx pause cc
storcli /cx/vx resume cc
storcli /cx/vx show cc
storcli /cx/vx start cc [force]
storcli /cx/vx stop cc
```

Note: *If enclosures are used to connect the physical drives to the controller, specify the IDs in the command.*

The detailed description for each command follows.

storcli /cx/vx pause cc

This command pauses an ongoing consistency check process. You can resume the consistency check at a later time. You can run this command only on a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 pause cc
```

storcli /cx/vx resume cc

This command resumes a suspended consistency check operation. You can run this command on a virtual drive that has a paused consistency check operation.

Input example:

```
storcli /c0/v4 resume cc
```

storcli /cx/vx show cc

This command shows the progress of the consistency check operation in percentage.

Input example:

```
storcli /c0/v5 show cc
```

```
storcli /cx/vx start cc force
```

This command starts a consistency check operation for a virtual drive. Typically, a consistency check operation is run on an initialized virtual drive. Use the `force` option to run a consistency check on an uninitialized drive.

Input example:

```
storcli /c0/v4 start cc
```

```
storcli /cx/vx stop cc
```

This command stops a consistency check operation. You can run this command only for a virtual drive that has a consistency check operation running.

Input example:

```
storcli /c0/v4 stop cc
```

Note: *You cannot resume a stopped consistency check process.*

10.5.4.10 Background Initialization Commands

The Storage Command Line Tool supports the following commands for background initialization:

```
storcli /cx/vx resume bgi  
storcli /cx/vx set autobgi=<on|off>  
storcli /cx/vx show autobgi  
storcli /cx/vx show bgi  
storcli /cx/vx stop bgi  
storcli /cx/vx suspend bgi
```

The detailed description for each command follows.

```
storcli /cx/vx resume bgi
```

This command resumes a suspended background initialization operation.

Input example:

```
storcli /c0/v0 resume bgi
```

```
storcli /cx/vx set autobgi=<on|off>
```

This command sets the auto background initialization setting for a virtual drive to on or off.

Input example:

```
storcli /c0/v0 set autobgi=on
```

```
storcli /cx/vx show autobgi
```

This command shows the background initialization setting for a virtual drive.

Input example:

```
storcli /c0/v0 show autobgi
```

```
storcli /cx/vx show bgi
```

This command shows the background initialization progress on the specified virtual drive in percentage.

Input example:

```
storcli /c0/v0 show bgi
```

```
storcli /cx/vx stop bgi
```

This command stops a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 stop bgi
```

```
storcli /cx/vx pause bgi
```

This command suspends a background initialization operation. You can run this command only for a virtual drive that is currently initialized.

Input example:

```
storcli /c0/v4 pause bgi
```

10.5.4.11 Virtual Drive Expansion Commands

The Storage Command Line Tool supports the following commands for virtual drive expansion:

```
storcli /cx/vx expand size=<value> [expandarray]  
storcli /cx/vx|vall show expansion
```

The detailed description for each command follows.

```
storcli /cx/vx expand size=<value> [expandarray]
```

This command expands the virtual drive within the existing array or if you replace the drives with drives larger than the size of the existing array. The value of the expand size is in GB. If the expandarray option is specified, the existing array is expanded. If this option is not specified, the virtual drive is expanded.

```
storcli /cx/vx show expansion
```

This command shows the expansion information on the virtual drive with and without array expansion.

Input example:

```
storcli /c0/v0 show expansion
```

10.5.5 Foreign Configurations Commands

The Storage Command Line Tool supports the following commands to view, import, and delete foreign configurations:

```
storcli /cx/fall|fall del|delete [ securitykey=ssssssssss ]
storcli /cx/fall|fall import [preview][ securitykey=ssssssssss ]
storcli /cx/fall|fall show [all] [ securitykey=ssssssssss ]
```

Note: Provide the security key when importing a locked foreign configuration created in a different machine that is encrypted with a security key.

The detailed description for each command follows.

```
storcli /cx/fall|fall del| delete [ securitykey=ssssssssss ]
```

This command deletes the foreign configuration of a controller. Input the security key if the controller is secured.

Input example:

```
storcli /c0/fall delete
```

```
storcli /cx/fall|fall import [preview] [ securitykey=ssssssssss ]
```

This command imports the foreign configurations of a controller. The `preview` option shows a summary of the foreign configuration before importing it.

Input example:

```
storcli /c0/fall import
```

```
storcli /cx/fall|fall show [all][ securitykey=ssssssssss ]
```

This command shows the summary of the entire foreign configuration for a particular controller. The `all` option shows all the information of the entire foreign configuration.

Note: The `EID:Slot` column is populated for the foreign PDs that are locked.

Input example:

```
storcli /c0/fall show preview
storcli /c0/fall import preview
storcli /c0/fall show all
```

10.5.6 BIOS-Related Commands

The Storage Command Line Tool supports the following BIOS commands:

```
storcli /cx set autobootselect(abs)=<on|off>
storcli /cx set bios=<on|off>
storcli /cx set BIOSMode=<SOE|BE|HCOE|HSM>
```

The detailed description for each command follows.

```
storcli /cx set autobootselect|abs=<on|off>
```

This command enables the BIOS to select the best logical drive as the boot drive.

Input example:

```
storcli /cx set autobootselect=on
```

```
storcli /cx set bios=<on|off>
```

This command enables or disables the Intel® RAID Controller's BIOS.

Note: *The legacy BIOS can load a limited number of the PCI device's BIOS. Disable the Intel® RAID BIOS to avoid issues during POST.*

Input example:

```
storcli /c0 set bios=enable
```

```
storcli /cx set BIOSMode=<SOE|BE|HCOE|HSM
```

This command sets the BIOS boot mode.

Input example:

```
storcli /c0/ set BIOSMode=SOE
```

10.5.6.1 OPROM BIOS Commands

The Storage Command Line Tool supports the following OPROM BIOS commands:

```
storcli /cx/ex/sx set bootdrive=on|off
```

```
storcli /cx/vx set bootdrive=on|off
```

```
storcli /cx show bootdrive
```

The detailed description for each command follows.

```
storcli /cx/ex/sx set bootdrive=on|off
```

This command sets the specified physical drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified physical drive.

Input example:

```
storcli /c0/e32/s4 set bootdrive=on
```

```
storcli /cx/vx set bootdrive=on|off
```

This command sets the specified virtual drive as the boot drive. During the next reboot, the BIOS looks for a boot sector in the specified virtual drive.

Input example:

```
storcli /c0/v0 set bootdrive=on
```

```
storcli/cx/vx show bootdrive
```

This command shows the boot drive for the controller. The boot drive can be a physical drive or a virtual drive.

Input example:

```
storcli /c0/v0 show bootdrive
```

10.5.7 Drive Group Commands

This section describes the drive group commands.

10.5.7.1 Drive Group Show Commands

The Storage Command Line Tool supports the following drive group commands:

```
storcli /cx/dall show
storcli /cx/dall show all
storcli /cx/dall show cachecade
storcli /cx/dx show
storcli /cx/dx show all
storcli /cx/dx set security=on
```

The detailed description for each command follows.

```
storcli /cx/dall show
```

This command shows the topology information of all the drive group.

Input example:

```
storcli /c0/dall show
```

```
storcli /cx/dall show all
```

This command shows all available configurations in the controller which includes topology information, virtual drive information, physical drive information, free space, and free slot information.

Input example:

```
storcli /c0/dall show all
```

```
storcli /cx/dall show cachecade
```

This command shows all CacheCade virtual drive information.

Input example:

```
storcli /c0/dall show cachecade
```

```
storcli /cx/dx show
```

This command shows the topology information of the drive group.

Input example:

```
storcli /c0/dx show
```

```
storcli /cx/dx show all
```

This command shows the physical drive and the virtual drive information for the drive group.

Input example:

```
storcli /c0/dx show all
```

```
storcli /cx/dx set security=on
```

This command enables security on the specified drive group.

Input example:

```
storcli /c0/dx set security=on all
```

10.5.8 Dimmer Switch Commands

10.5.8.1 Change Virtual Drive Power Settings Commands

The Storage Command Line Tool supports the following command to change the Dimmer Switch* setting. The Dimmer Switch is the power-saving policy for the virtual drive.

```
storcli /cx/vx set ds=<default | auto | none | max | maxnocache>
```

This command changes the power-saving properties on a virtual drive. See `dimmerswitch` in the following table for values.

Input example:

```
storcli /cx/vx set ds=default
```

Note: Only the `ds3` dimmer switch option cannot be selected in the Storage Command Line Tool.

You can use the following combinations for the dimmer switch commands:

```
storcli /cx set ds=off type=1|2|3|4
```

```
storcli /cx set ds=on type=1|2 [properties]
```

```
storcli /cx set ds=on type=3|4 defaultldtype=<value> [properties]
```

```
storcli /cx set ds=on [properties]
```

The following table describes the power-saving options.

Table 34. Dimmer Switch Input Options

Option	Value Range	Description
<code>dimmerswitch</code> or <code>ds</code>	<code>on off</code>	Turns the dimmer switch option on.

Option	Value Range	Description
type	1: Unconfigured 2: Hot spare 3: Virtual drive 4: All	Specifies the type of drives that the dimmer switch feature is applicable. By default, it is activated for unconfigured drives, hot spare drives, and virtual drives.
defaultldtype	auto: Logical device power savings are managed by the firmware. none: No power saving policy. max: Logical device uses maximum power savings. maxnocache: Logical device does not cache write to maximise power savings.	Specifies the default logical drive type that is created by the dimmer switch option; set to none automatically.
properties	disableldps: Interval in hours or time in <i>hh:mm</i> format spinupdrivecount: Valid enclosure number (0 to 255) SpinUpEncDelay: Valid time in seconds	Sets the interval or time in which the power-saving policy for the logical drive is turned off. Specifies the number of drives in the enclosure that are spun up. Specifies the delay of spin-up groups within an enclosure in seconds.

storcli/cx show DimmerSwitch(ds)

This command shows the current dimmer switch setting for the controller.

Input example:

```
storcli/c0 show ds
```

10.5.9 BBU Commands

The Storage Command Line Tool supports the following battery backup unit (BBU) commands:

Note: *To increase the life of a battery, the battery is not fully charged. Band Gap charging keeps the maximum battery charge within a band comfortably above the data retention time requirement instead of keeping the battery charged to the maximum level. However, when a learn cycle is required, the battery is fully charged because a learn cycle starts only after the battery is fully charged.*

```
storcli /cx/bbu show
storcli /cx/bbu show all
storcli /cx/bbu set autolearnmode=<value>
storcli /cx/bbu set bbuMode=<value>
storcli /cx/bbu set learndelayinterval=<value>
storcli /cx/bbu set powermode=sleep
```

```
storcli /cx/bbu set writeaccess=sealed
storcli /cx/bbu show modes
storcli /cx/bbu show properties
storcli /cx/bbu show status
storcli /cx/bbu start learn
```

The detailed description for each command follows.

storcli /cx/bbu show

This command shows the summary information for the BBU of a controller.

Input example:

```
storcli /c0/bbu show
```

storcli /cx/bbu show all

This command shows all the information of the BBU.

Input example:

```
storcli /c0/bbu show all
```

storcli /cx/bbu set autolearnmode=<value>

This command starts the automatic learn cycle on the battery. The possible values are 0 - Enabled, 1- Disabled, and 2 - WarnViaEvent.

Input example:

```
storcli /c0/bbu set autolearnmode=0
```

storcli /cx/bbu set bbuMode=<value>

This command sets the BBU mode for the BBU. The following table shows the various BBU modes.

Table 35. BBU Mode

Mode	Description
0	48 hours of retention ^A at 60 °C, 1-year Service Life.
1	12 hours of retention at 45 °C, 5-year Service Life, transparent learn ^B .
2	12 hours of retention at 55 °C, 3-year Service Life, transparent learn.
3	24 hours of retention at 45 °C, 3-year Service Life, transparent learn.
4	48 hours of retention at 45 °C, 3-year Service Life.
5	48 hours of retention at 55 °C, 1-year Service Life.
6	The same as the description for BBU mode 5. The BBU mode 6 enables you to receive events when the battery capacity reaches suboptimal and critical thresholds.

- A. Indicates how long the battery can hold data in the controller's memory in case of accidental system shutdown.
- B. The controller's performance is not affected during the battery's learn cycle.

Input example:

```
storcli /c0/bbu set bbuMode=2
```

Note: *BBU modes are supported on any BBU8/9 bbu/controller combo and later-generation controllers.*

```
storcli /cx/bbu set learndelayinterval=<value>
```

This command sets the learn delay interval for the BBU in hours. The value must be between 0 to 168 hours (7 days).

Input example:

```
storcli /c0/bbu set learnDelayInterval=30
```

```
storcli /cx/bbu set powermode=sleep
```

This command places the battery in low-power storage mode. The battery automatically exits this state after five seconds.

Input example:

```
storcli /c0/bbu set powermode=sleep
```

```
storcli /cx/bbu set writeaccess=sealed
```

This command seals the gas gauge EEPROM write access.

Note: *Use the set writeaccess=sealed command at manufacturing time.*

Input example:

```
storcli /c0/bbu set writeaccess=sealed
```

```
storcli /cx/bbu show modes
```

This command shows the bbu mode information that includes the bbu mode number, retention time, service life, maximum temperature, and battery learn information.

Input example:

```
storcli /c0/bbu show modes
```

```
storcli /cx/bbu show properties
```

This command shows the BBU Learn properties for a controller.

Input example:

```
storcli /c0/bbu show properties
```

```
storcli /cx/bbu show status
```

This command shows the battery information, firmware status, and the gas gauge status.

Input example:

```
storcli /c0/bbu show status
```

```
storcli /cx/bbu start learn
```

This command starts the BBU learning cycle. The battery learn cycle is immediately started and no other parameters are required for this command.

Input example:

```
storcli /c0/bbu start learn
```

10.5.10 Enclosure Commands

The Storage Command Line Tool supports the following enclosure commands:

```
storcli /cx/ex download src=filepath [forceActivate]
```

```
storcli /cx/ex show all
```

```
storcli /cx/ex show status
```

The detailed description for each command follows.

```
storcli /cx/ex download src=filepath [forceactivate]
```

This command flashes the firmware with the file specified at the command line. The enclosure performs an error check after the operation. The following option can be used with the enclosure firmware download command.

Table 36. Enclosure Firmware Download Command Options

Option	Value Range	Description
forceactivate	—	Issues a command descriptor block (CDB) with write command with no data with command mode 0x0F (flash download already in progress). Note: This option is used primarily to activate Scotch Valley Enclosures.

Note: The firmware file that is used to flash the enclosure can be of any format. The StorCLI utility assumes that you provide a valid firmware image.

Input example:

```
storcli /c0/e0 download src=c:\file2.bin
```

```
storcli /cx/ex show all
```

This command shows all enclosure information, which includes general enclosure information, enclosure inquiry data, a count of enclosure elements, and information about the enclosure elements.

Input example:

```
storcli /c0/e0 show all
```

```
storcli /cx/ex show status
```

This command shows the enclosure status and the status of all the enclosure elements.

Input example:

```
storcli /c0/e0 show status
```

10.5.11 PHY Commands

The Storage Command Line Tool supports the following PHY commands:

```
storcli /cx/px|pall set linkspeed=0(auto)|1.5|3|6|12
```

```
storcli /cx/px|pall show
```

```
storcli /cx/px|pall show all
```

The detailed description for each command follows.

```
storcli /cx/px|pall set linkspeed=0(auto)|1.5|3|6|12
```

This command sets the PHY link speed. You can set the speed to 1.5 Gb/s, 3 Gb/s, 6 Gb/s, or 12 Gb/s. The linkspeed is set to auto when you specify linkspeed = 0.

Input example:

```
storcli /c0/p0 set linkspeed=1.5
```

```
storcli /cx/px|pall show
```

This command shows the basic PHY layer information.

Input example:

```
storcli /c1/p0 show
```

```
storcli /cx/px|pall show all
```

This command shows all the PHY layer information.

Input example:

```
storcli /c1/p0 show all
```

10.5.12 Logging Commands

The Storage Command Line Tool supports the following commands to generate and maintain log files:

```
storcli /cx clear events
```

```
storcli /cx delete termlog
```

```
storcli /cx show events file=<absolute path>
```

```
storcli /cx show eventloginfo
```

```
storcli /cx show termlog type=config|contents
```

The detailed description for each command follows.

```
storcli /cx delete events
```

This command deletes all records in the event log.

Input example:

```
storcli /c0 delete events
```

```
storcli /cx delete termlog
```

This command clears the TTY (firmware log for issue troubleshooting) logs.

Input example:

```
storcli /c0 delete termlog
```

```
storcli /cx show events file=<absolute path>
```

This command prints the system log to a text file and saves the file in the specified location.

Input example:

```
storcli /c0 show events file=C:\Users\brohan\test\eventreports
```

```
storcli /cx show eventloginfo
```

This command shows the history of log files generated.

Input example:

```
storcli /c0 show eventloginfo type=config
```

```
storcli /cx show termlog type=config|contents
```

This command shows the firmware logs. The `config` option shows the term log configuration (settings of TTY BBU buffering), the `contents` option shows the term log. The `contents` option is the default.

Input example:

```
storcli /c0 show termlog type=contents
```

10.6 Frequently Used Tasks

10.6.1 Showing the Version of the Storage Command Line Tool

The following command shows the version of the command line tool:

```
storcli -v
```

10.6.2 Showing StorCLI Help

The following command shows the command line tool help:

```
storcli -h
```

Help appears for all the StorCLI commands.

10.6.3 Showing System Summary Information

The following command shows the summary of all the controller information:

```
storcli -show [all]
```

10.6.4 Showing Free Space in a Controller

The following command shows the free space available in the controller:

```
storcli /cx show freespace
```

10.6.5 Adding Virtual Drives

The following command creates a virtual drive:

```
storcli /cx add vd  
type=raid[0|1|5|6|10|50|60][Size=<VD1_Sz>,<VD2_Sz>,...|*all]  
[name=<VDNAME1>,...] drives=e:s|e:s-x|e:s-x,y [PDperArray=x|auto*]  
[SED] [pdcache=on|off|*default][pi]  
[DimmerSwitch(ds)=default|automatic(auto)|  
*none|maximum(max)|MaximumWithoutCaching(maxnocache)] [wt|*wb]  
[nora|*ra] [*direct|cached] [CachedBadBBU|*NoCachedBadBBU]  
[strip=<8|16|32|64|128|256|512|1024] [AfterVd=x]  
[Spares=[e:]s|[e:]s-x|[e:]s-x,y] [force]
```

The following inputs can be used when adding virtual drives:

- The controller in which the virtual drives are created.
- The RAID type of the virtual drives. The supported RAID types are 0, 1, 5, 6, 10, 50, 60.
- The size of each virtual drive.

- The drives that are used to create the virtual drives.
drives = e:s|e:s-x|e:s-x,y
Where:
 - e specifies the enclosure id.
 - s represents the slot in the enclosure.
 - e:s-x is the range conventions used to represents slots s to x in the enclosure e.
- The physical drives per array. The physical drives per array can be set to a particular value.
- The SED option creates security-enabled drives.
- The PDcache option can be set to on or off.
- The pi option enables protection information.
- The dimmer switch is the power save policy. It can be set to default or automatic *, none, maximum(max), or MaximumWithoutCaching(maxnocache).
- The wt option disables write back.
- The nora option disables read ahead.
- The cached option enables the cached memory.
- The CachedBadBBU option enables caching when BBU is not functional.
- The strip option sets the strip size. It can take the values 8, 16, 32, 64, 128, 256, 512, 1024.
- The AfterVdX option creates the virtual drives in the adjacent free slot next to the specified virtual drives.

Note: The * indicates default values used in the creation of the virtual drives. If values are not specified, the default values are taken.

Example: `/cxadd vd type=r1 drives=0:10-15 WB Direct strip=64`

This command creates a RAID volume of RAID 1 type from drives in slots 10 to slot 15 in enclosure 0. The strip size is 64kb.

10.6.6 Setting the Cache Policy in a Virtual Drive

The following command sets the write cache policy of the virtual drive:

```
storcli /cx/v(x|all) set wrcache=wt|wb|awb
```

The command sets the write cache to write back, write through, or always write back.

10.6.7 Showing Virtual Drive Information

The following command shows the virtual drive information for all the virtual drives in the controller:

```
storcli /cx show [all]
```

10.6.8 Deleting Virtual Drives

The following command deletes virtual drives:

```
storcli /cx/v(x|all) del [cc|cachecade]
```

The following inputs are required when deleting a virtual drive:

- The controller on which the virtual drive or virtual drives is present.
- The virtual drives that must be deleted; or you can delete all the virtual drives on the controller using the `vall` option.
- The `cc` or `cachecade` option to confirm that the deleted drive is a CacheCade drive.

10.6.9 Flashing Controller Firmware

The following command is used to flash the controller firmware.

```
storcli /cx download file=filepath [fwtype=<value>] [nosigchk]  
[noverchk][resetnow]
```

For more information, see “[Flashing Controller Firmware Command](#)” on page 267.

10.7 Unsupported StorCLI Commands

The commands in the following table are not supported for the Intel[®] RAID Controllers such as Intel[®] RAID Controller RS3WC080.

Table 37. Unsupported Commands for Intel[®] RAID Controller

Command Group	Command
Jbod	<code>storcli /c0 set jbod=<on off></code>
	<code>storcli /c0/s2 set jbod</code>
	<code>storcli /c0/s2 set bootdrive=<on off></code>
DS	<code>storcli /cx(x all) set ds=OFF type=1 2 3 4</code>
	<code>storcli /cx(x all) set ds=ON type=1 2 [properties]</code>
	<code>storcli /cx(x all) set ds=ON type=3 4 DefaultLdType=<val> [properties]</code>
	<code>storcli /cx(x all) set ds [properties]</code>
	<code>storcli /cx/v(x all) set ds=Default Auto None Max MaxNoCache</code>
Security	<code>storcli /cx delete security key</code>
	<code>storcli /cx set securitykey=xxxxxxxx {passphrase=xxxx} {keyid=xxx}</code>
	<code>storcli /cx set securitykey keyid=xxx</code>
	<code>storcli /cx compare securitykey=xxxxxxxx</code>
	<code>storcli /cx set securitykey=xxxxxxxx oldsecuritykey=xxxxxxxx</code>

Command Group	Command
ASO	storcli /cx(x all) set aso key=<keyvalue> preview
	storcli /cx(x all) set aso key=<key value>
	storcli /cx(x all) set aso transfertovault
	storcli /cx(x all) set aso rehostcomplete
	storcli /cx(x all) set aso deactivatetrialsec
	storcli /cx(x all) show safeid
	storcli /cx(x all) show rehostinfo
	storcli /c0 set time =<yyyymmdd hh:mm:ss system>
	storcli /c0 show cc consistencycheck
	storcli /c0/vall show expansion
	storcli /c0 set jbod
	storcli /cx download src=<filepath> [forceActivate]
Copy back	storcli /cx[/ex]/sx show copyback
	storcli /cx[/ex]/sx start copyback target=eID:sID
	storcli /cx[/ex]/sx stop copyback
	storcli /cx[/ex]/sx pause copyback
	storcli /cx[/ex]/sx resume copyback
Migrate	storcli /cx/v(x all) show migrate
	storcli /cx/vx start migrate type=raidx [option=add remove drives=[e:]s [e:]s- x [e:]s-x,y] [Force]
Cache	storcli /cx/v(x all) set ssdcaching=on off
	storcli /cx(x all) show preservedcache
	storcli /cx/v(x all) delete preservedcache[force]
BBU	storcli /cx/bbu show
	storcli /cx/bbu show all
	storcli /cx/bbu set [learnDelayInterval=<val> bbuMode=<val>
	storcli /cx/bbu start learn

Command Group	Command
Secure erase	storcli /cx/sx secureerase [force]
	storcli /cx/sx start erase [simple normal thorough][erasepatternA=<val>]
	storcli /cx/sx stop erase
	storcli /cx/sx show erase
Consistency check	storcli /cx show cc/ConsistencyCheck
Controller	storcli /cx show cc

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 the screenshots may change, and the following 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 175](#).

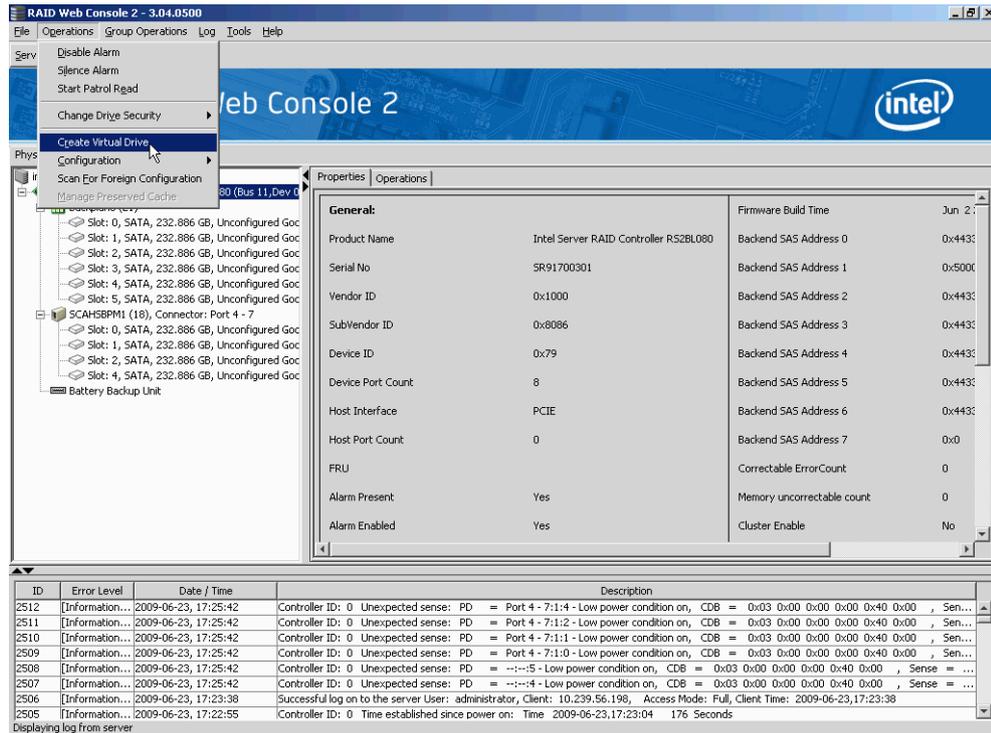


Figure 175. Virtual Drive Creation Menu

The dialog box shown in Figure 176 appears.

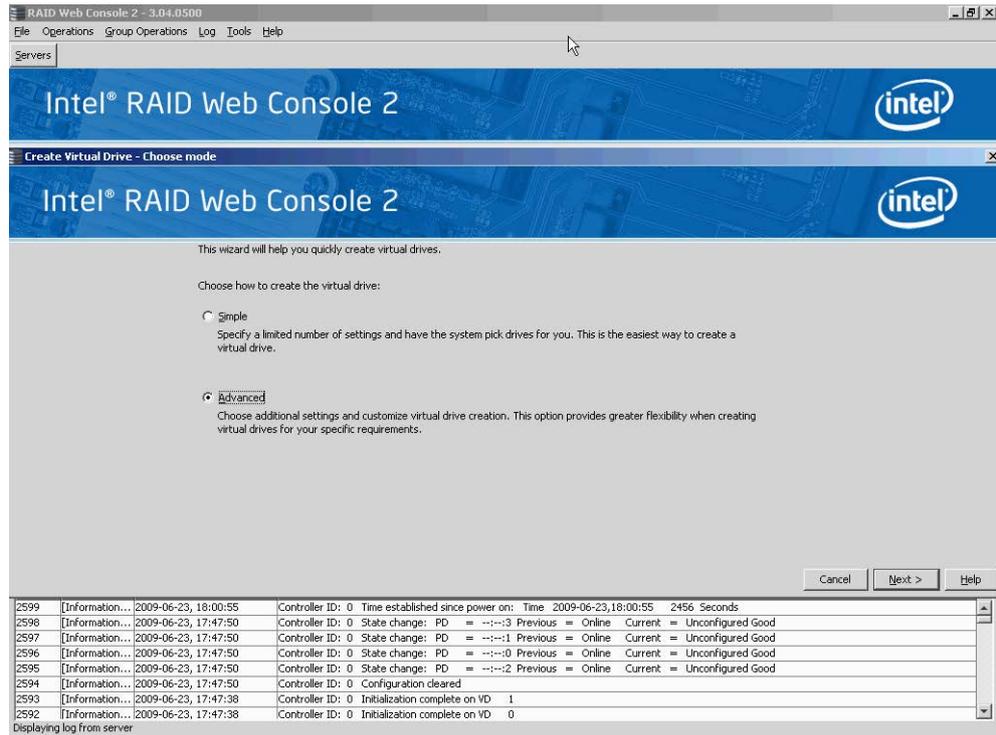


Figure 176. Virtual Drive Creation Mode

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

The Create Drive Group Settings screen appears, as shown in [Figure 177](#).

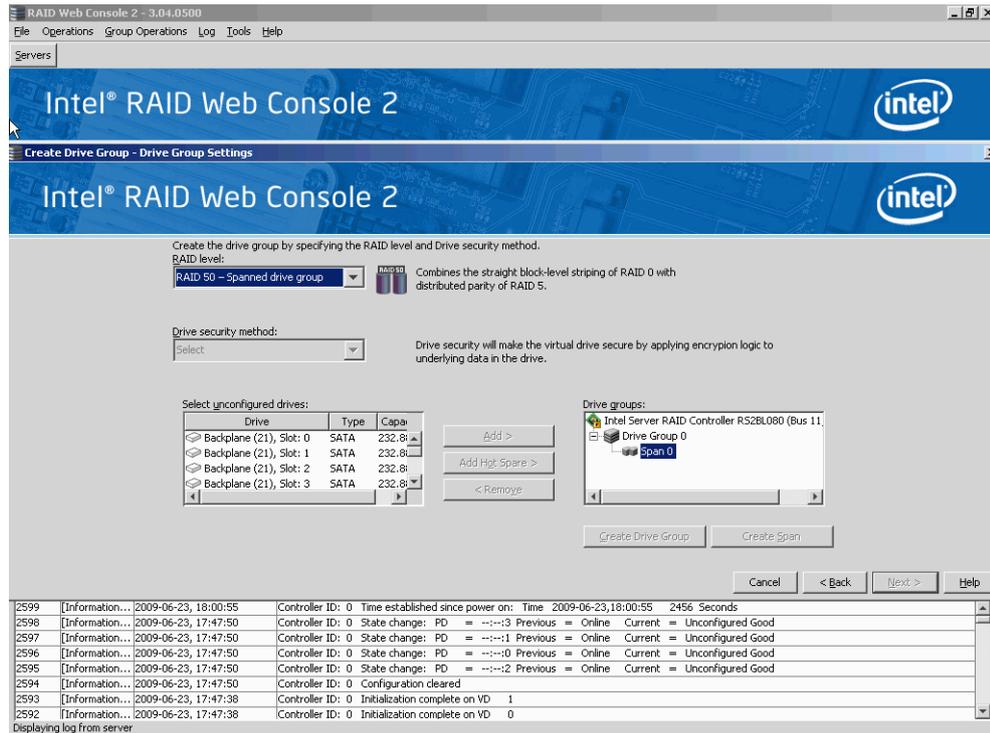


Figure 177. Create Drive Group Settings Screen

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 178](#)

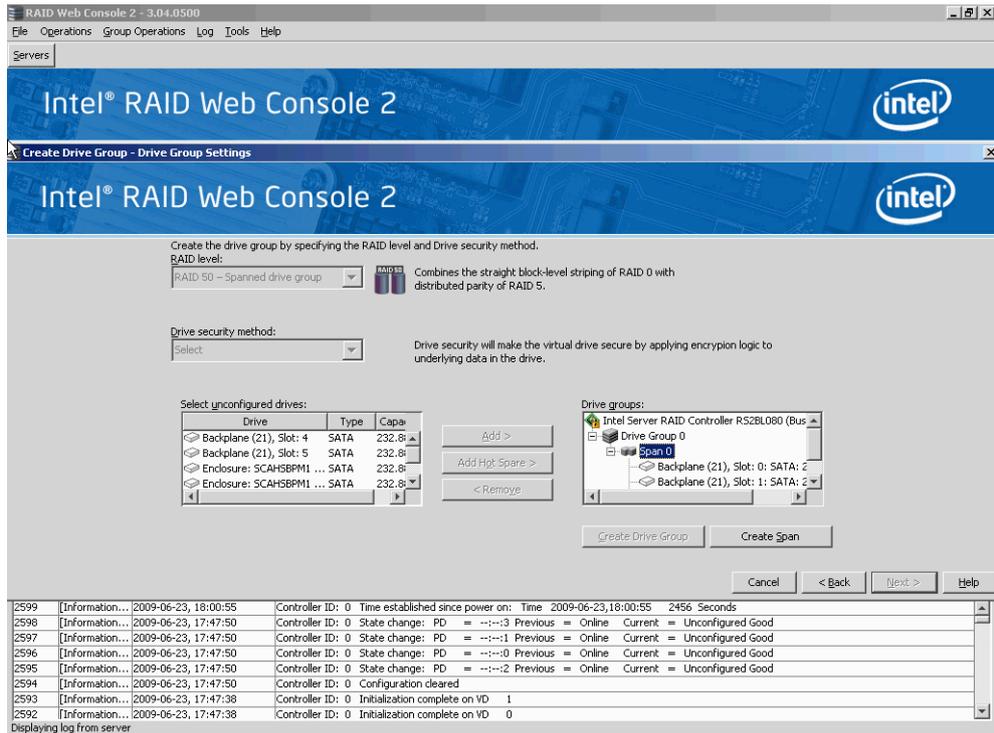


Figure 178. Span 0 of Drive Group 0

- 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 179](#)

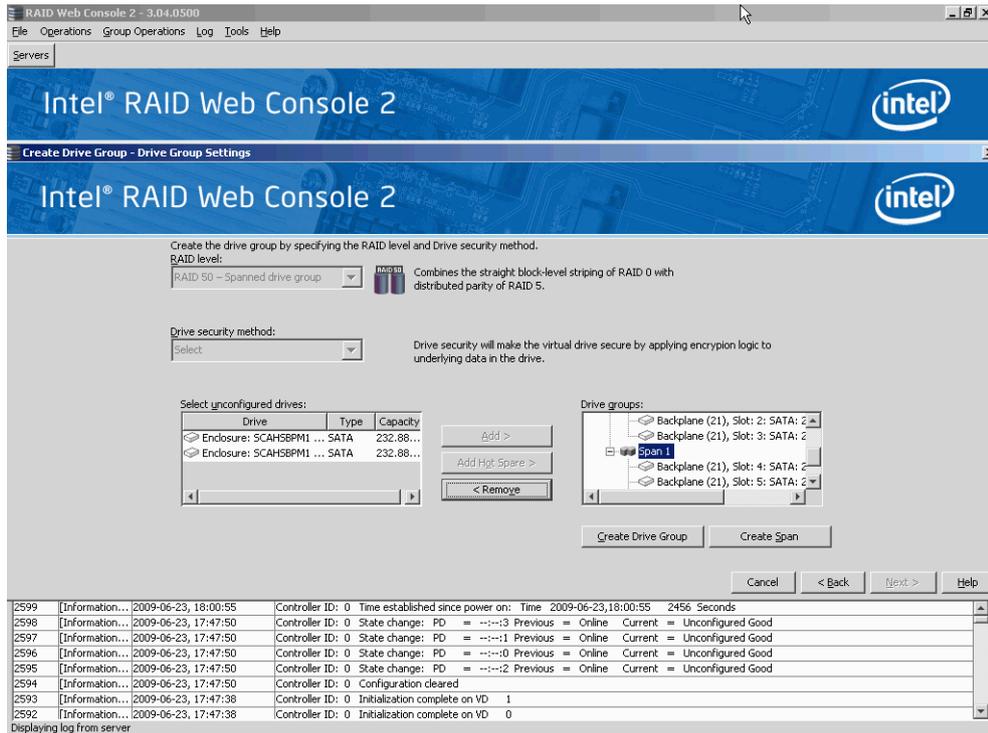


Figure 179. Span 0 and Span 1 of Drive Group 0

- 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 180](#). 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.

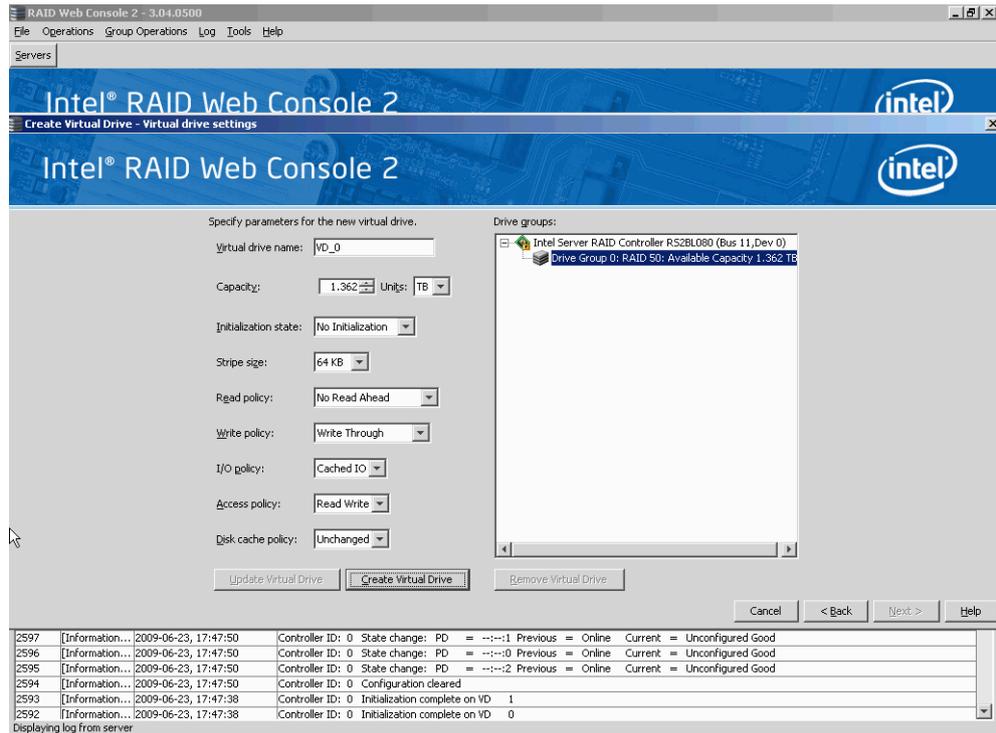


Figure 180. Virtual Drive Settings Window

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 181](#). 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.

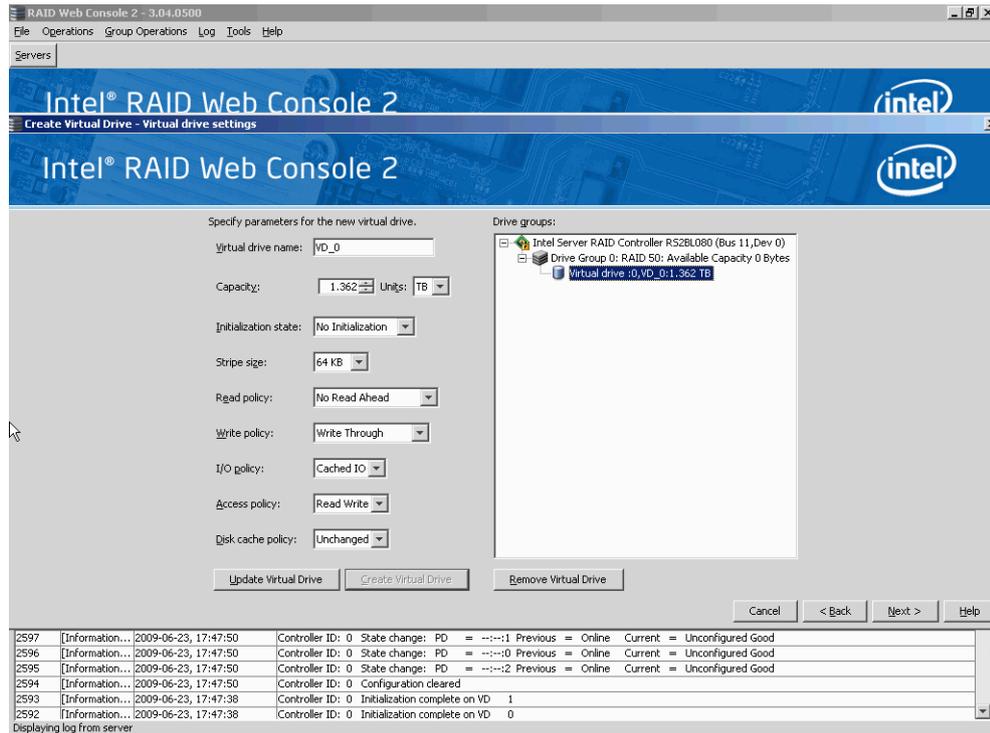


Figure 181. New Virtual Drive 0

6. Click Next.

The **Create Virtual Drive - Summary** window appears, as shown in Figure 182. This window shows the selections you made for advanced configuration.

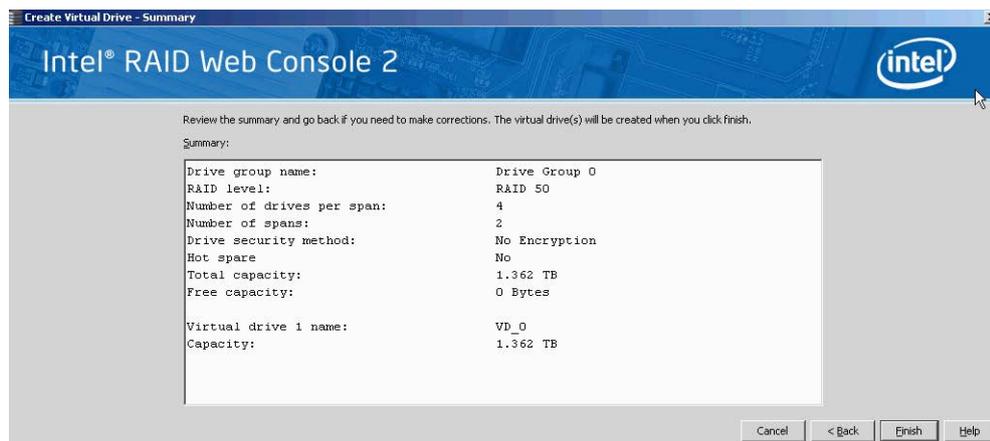


Figure 182. Create Virtual Drive Summary Window

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 notifies you that the virtual drives were created successfully, as shown in [Figure 183](#). 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 183. 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 96](#). If you select **No**, the utility asks whether you want to close the wizard, as [Figure 184](#).
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 184. 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 38. 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

Number	Type	Description
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%%
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

Number	Type	Description
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
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

Number	Type	Description
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

Number	Type	Description
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
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

Number	Type	Description
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
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

Number	Type	Description
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
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

Number	Type	Description
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
221	Warning	NVRAM is corrupt; reinitializing
222	Warning	NVRAM mismatch occurred
223	Warning	SAS wide port %d lost link on PHY %d
224	Info	SAS wide port %d restored link on PHY %d
225	Warning	SAS port %d, PHY %d has exceeded the allowed error rate
226	Warning	Bad block reassigned on %s at %lx to %lx
227	Info	Controller Hot Plug detected
228	Warning	Enclosure %s temperature sensor %d differential detected
229	Info	Drive test cannot start. No qualifying drives found
230	Info	Time duration provided by host is not sufficient for self check
231	Info	Marked Missing for %s on drive group %d row %d
232	Info	Replaced Missing as %s on drive group %d row %d
233	Info	Enclosure %s Temperature %d returned to normal
234	Info	Enclosure %s Firmware download in progress
235	Warning	Enclosure %s Firmware download failed
236	Warning	%s is not a certified drive
237	Info	Dirty cache data discarded by user
238	Info	Drives missing from configuration at boot
239	Info	Virtual drives (VDs) missing drives and will go offline at boot: %s

Number	Type	Description
240	Info	VDs missing at boot: %s
241	Info	Previous configuration completely missing at boot
242	Info	Battery charge complete
243	Info	Enclosure %s fan %d speed changed
244	Info	Dedicated spare %s imported as global due to missing arrays
245	Info	%s rebuild not possible as SAS/SATA is not supported in an array
246	Info	SEP %s has been rebooted as a part of enclosure firmware download. SEP will be unavailable until this process completes.
247	Info	Inserted PD: %s Info: %s
248	Info	Removed PD: %s Info: %s
249	Info	VD %s is now OPTIMAL
250	Warning	VD %s is now PARTIALLY DEGRADED
251	Critical	VD %s is now DEGRADED
252	Fatal	VD %s is now OFFLINE
253	Warning	Battery requires reconditioning; please initiate a LEARN cycle
254	Warning	VD %s disabled because RAID-5 is not supported by this RAID key
255	Warning	VD %s disabled because RAID-6 is not supported by this controller
256	Warning	VD %s disabled because SAS drives are not supported by this RAID key
257	Warning	PD missing: %s
258	Warning	Puncturing of LBAs enabled
259	Warning	Puncturing of LBAs disabled
260	Critical	Enclosure %s EMM %d not installed
261	Info	Package version %s
262	Warning	Global affinity Hot Spare %s commissioned in a different enclosure
263	Warning	Foreign configuration table overflow
264	Warning	Partial foreign configuration imported, PDs not imported:%s
265	Info	Connector %s is active
266	Info	Board Revision %s
267	Warning	Command timeout on PD %s, CDB:%s

Number	Type	Description
268	Warning	PD %s reset (Type %02x)
269	Warning	VD bad block table on %s is 80% full
270	Fatal	VD bad block table on %s is full; unable to log block %lx (on %s at %lx)
271	Fatal	Uncorrectable medium error logged for %s at %lx (on %s at %lx)
272	Info	VD medium error corrected on %s at %lx
273	Warning	Bad block table on PD %s is 100% full
274	Warning	VD bad block table on PD %s is 100% full
275	Fatal	Controller needs replacement, IOP is faulty
276	Info	CopyBack started on PD %s from PD %s
277	Info	CopyBack aborted on PD %s and src is PD %s
278	Info	CopyBack complete on PD %s from PD %s
279	Progress	CopyBack progress on PD %s is %s
280	Info	CopyBack resumed on PD %s from %s
281	Info	CopyBack automatically started on PD %s from %s
282	Critical	CopyBack failed on PD %s due to source %s error
283	Warning	Early Power off warning was unsuccessful
284	Info	BBU FRU is %s
285	Info	%s FRU is %s
286	Info	Controller hardware revision ID %s
287	Warning	Foreign import shall result in a backward incompatible upgrade of configuration metadata
288	Info	Redundant path restored for PD %s
289	Warning	Redundant path broken for PD %s
290	Info	Redundant enclosure EMM %s inserted for EMM %s
291	Info	Redundant enclosure EMM %s removed for EMM %s
292	Warning	Patrol Read can't be started, as PDs are either not ONLINE, or are in a VD with an active process, or are in an excluded VD
293	Info	Copyback aborted by user on PD %s and src is PD %s
294	Critical	Copyback aborted on hot spare %s from %s, as hot spare needed for rebuild

Number	Type	Description
295	Warning	Copyback aborted on PD %s from PD %s, as rebuild required in the array
296	Fatal	Controller cache discarded for missing or offline VD %s When a VD with cached data goes offline or missing during runtime, the cache for the VD is discarded. Because the VD is offline, the cache cannot be saved.
297	Info	Copyback cannot be started as PD %s is too small for src PD %s
298	Info	Copyback cannot be started on PD %s from PD %s, as SAS/SATA is not supported in an array
299	Info	Microcode update started on PD %s
300	Info	Microcode update completed on PD %s
301	Warning	Microcode update timeout on PD %s
302	Warning	Microcode update failed on PD %s
303	Info	Controller properties changed
304	Info	Patrol Read properties changed
305	Info	CC Schedule properties changed
306	Info	Battery properties changed
307	Warning	Periodic Battery Relearn is pending. Please initiate manual learn cycle as Automatic learn is not enabled.
308	Info	Drive security key created
309	Info	Drive security key backed up
310	Info	Drive security key from escrow, verified
311	Info	Drive security key changed
312	Warning	Drive security key, re-key operation failed
313	Warning	Drive security key is invalid
314	Info	Drive security key destroyed
315	Warning	Drive security key from escrow is invalid
316	Info	VD %s is now secured
317	Warning	VD %s is partially secured
318	Info	PD %s security activated
319	Info	PD %s security disabled
320	Info	PD %s is reprovisioned

Number	Type	Description
321	Info	PD %s security key changed
322	Fatal	Security subsystem problems detected for PD %s
323	Fatal	Controller cache pinned for missing or offline VD %s
324	Fatal	Controller cache pinned for missing or offline VDs: %s
325	Info	Controller cache discarded by user for VDs: %s
326	Info	Controller cache destaged for VD %s
327	Warning	Consistency Check started on an inconsistent VD %s
328	Warning	Drive security key failure, cannot access secured configuration
329	Warning	Drive security password from user is invalid
330	Warning	Detected error with the remote battery connector cable
331	Info	Power state change on PD %s from %s to %s
332	Info	Enclosure %s element (SES code 0x%x) status changed
333	Info	PD %s rebuild not possible as HDD/CacheCade software mix is not supported in a drive group
334	Info	Copyback cannot be started on PD %s from %s, as HDD/CacheCade software mix is not supported in a drive group
335	Info	VD bad block table on %s is cleared
336	Caution	SAS topology error: 0x%lx
337	Info	VD cluster of medium errors corrected for %s at %lx (on %s at %lx)
338	Info	Controller requests a host bus rescan
339	Info	Controller repurposed and factory defaults restored
340	Info	Drive security key binding updated
341	Critical	Controller encountered a fatal error and was reset
342	Info	Snapshots enabled on %s (Repository %s)
343	Info	Snapshots disabled on %s (Repository %s) by the user
344	Critical	Snapshots disabled on %s (Repository %s), due to a fatal error
345	Info	Snapshot created on %s at %s
346	Info	Snapshot deleted on %s at %s
347	Info	View created at %s to a snapshot at %s for %s
348	Info	View at %s is deleted, to snapshot at %s for %s

Number	Type	Description
349	Info	Snapshot rollback started on %s from snapshot at %s
350	Fatal	Snapshot rollback on %s internally aborted for snapshot at %s
351	Info	Snapshot rollback on %s completed for snapshot at %s
352	Info	Snapshot rollback progress for snapshot at %s, on %s is %s
353	Warning	Snapshot space for %s in snapshot repository %s, is 80%% full
354	Critical	Snapshot space for %s in snapshot repository %s, is full
355	Warning	View at %s to snapshot at %s, is 80%% full on snapshot repository %s
356	Critical	View at %s to snapshot at %s, is full on snapshot repository %s
357	Critical	Snapshot repository lost for %s
358	Warning	Snapshot repository restored for %s
359	Critical	Snapshot encountered an unexpected internal error: 0x%lx
360	Info	Auto Snapshot enabled on %s (snapshot repository %s)
361	Info	Auto Snapshot disabled on %s (snapshot repository %s)
362	Critical	Configuration command could not be committed to disk, please retry
363	Info	COD on %s updated as it was stale
364	Warning	Power state change failed on %s (from %s to %s)
365	Warning	%s is not available
366	Info	%s is available
367	Info	%s is used for CacheCade with capacity 0x%lx logical blocks
368	Info	%s is using CacheCade %s
369	Info	%s is no longer using CacheCade %s
370	Critical	Snapshot deleted due to resource constraints for %s in snapshot repository %s
371	Warning	Auto Snapshot failed for %s in snapshot repository %s
372	Warning	Controller reset on-board expander
373	Warning	CacheCade (%s) capacity changed and is now 0x%lx logical blocks
374	Warning	Battery cannot initiate transparent learn cycles
375	Info	Premium feature %s key was applied for - %s
376	Info	Snapshot schedule properties changed on %s

Number	Type	Description
377	Info	Snapshot scheduled action is due on %s
378	Info	Performance Metrics: collection command 0x%lx
379	Info	Premium feature %s key was transferred - %s
380	Info	Premium feature serial number %s
381	Warning	Premium feature serial number mismatched. Key-vault serial num - %s
382	Warning	Battery cannot support data retention for more than %d hours. Please replace the battery
383	Info	%s power policy changed to %s (from %s)
384	Warning	%s cannot transition to max power savings
385	Info	Host driver is loaded and operational
386	Info	%s mirror broken
387	Info	%s mirror joined
388	Warning	%s link %d failure in wide port
389	Info	%s link %d restored in wide port
390	Info	Memory module FRU is %s
391	Warning	Cache-vault power pack is sub-optimal. Please replace the pack
392	Warning	Foreign configuration auto-import did not import any drives
393	Warning	Cache-vault microcode update required
394	Warning	CacheCade (%s) capacity exceeds maximum allowed size, extra capacity is not used
395	Warning	LD (%s) protection Info lost
396	Info	Diagnostics passed for %s
397	Critical	Diagnostics failed for %s
398	Info	Server Power capability Diagnostic Test Started
399	Info	Drive Cache settings enabled during rebuild for %s
400	Info	Drive Cache settings restored after rebuild for %s
401	Info	Drive %s commissioned as Emergency spare
402	Warning	Reminder: Potential non-optimal configuration due to drive %s commissioned as emergency spare
403	Info	Consistency Check suspended on %s

Number	Type	Description
404	Info	Consistency Check resumed on %s
405	Info	Background Initialization suspended on %s
406	Info	Background Initialization resumed on %
407	Info	Reconstruction suspended on %s
408	Info	Rebuild suspended on %
409	Info	Copyback suspended on %s
410	Info	Reminder: Consistency Check suspended on %
411	Info	Reminder: Background Initialization suspended on %s
412	Info	Reminder: Reconstruction suspended on %s
413	Info	Reminder: Rebuild suspended on %s
414	Info	Reminder: Copyback suspended on %s
415	Info	Reminder: Patrol Read suspended
416	Info	Erase aborted on %s
417	Critical	Erase failed on %s (Error %02x)
418	Progress	Erase progress on %s is %s
419	Info	Erase started on %s
420	Info	Erase completed on %s
421	Info	Erase aborted on %s
422	Critical	Erase failed on %s
423	Progress	Erase progress on %s is %s
424	Info	Erase started on %s
425	Info	Erase complete on %s
426	Warning	Potential leakage during erase on %s
427	Warning	Battery charging was suspended due to high battery temperature
428	Info	NVCache firmware update was successful
429	Warning	NVCache firmware update failed
430	Fatal	%s access blocked as cached data in CacheCade is unavailable
431	Info	CacheCade disassociate started on %s
432	Info	CacheCade disassociate completed on %s

Number	Type	Description
433	Critical	CacheCade disassociate failed on %s
434	Progress	CacheCade disassociate progress on %s is %s
435	Info	CacheCade disassociate aborted by user on %s
436	Info	Link speed changed on SAS port %d and PHY %d
437	Warning	Advanced Software Options was deactivated for - %s
438	Info	%s is now accessible
439	Info	%s is using CacheCade
440	Info	%s is no longer using CacheCade
441	Info	Patrol Read aborted on %s
61440	Info	Monitor has shutdown successfully
61441	Fatal	No controllers detected
61442	Info	Testing the Email Configuration
61443	Fatal	Incompatible Controller
61444	Info	Successful log on to the server
10000	Info	Successful log out from the server
10001	Warning	Full access denied on the server
10002	Info	Server log cleared
65529	Critical	Temperature has exceeded Critical level on Sensor
65530	Warning	Temperature has exceeded Warning level on Sensor
65531	Critical	The Backup Rail Monitor has failed
65532	Critical	SSD Life is exhausted on Slot
65533	Warning	SSD Life is at warning level on Slot
65534	Info	SSD throttling is now removed on Slot
65535	Info	SSD is being throttled on Slot

Glossary

This chapter provides a glossary for terms used in this document.

A

Absolute state of charge	Predicted remaining battery capacity expressed as a percentage of Design Capacity. Note that the Absolute State of Charge operation can return values greater than 100 percent.
Access policy	A virtual drive property indicating what kind of access is allowed for a particular virtual drive. The possible values are Read/Write, Read Only, or Blocked.
Alarm enabled	A controller property that indicates whether the controller's on-board alarm is enabled.
Alarm present	A controller property that indicates whether the controller has an on-board alarm. If present and enabled, the alarm is sounded for certain error conditions.
Array	See Drive group .
Auto learn mode	The controller performs the learn cycle automatically in this mode. This mode offers the following options: <ul style="list-style-type: none">• BBU Auto Learn: Firmware tracks the time since the last learn cycle and performs a learn cycle when due.• BBU Auto Learn Disabled: Firmware does not monitor or initiate a learn cycle. You can schedule learn cycles manually.• BBU Auto Learn Warn: Firmware warns about a pending learn cycle. You can initiate a learn cycle manually. After the learn cycle is complete, the firmware resets the counter and warns you when the next learn cycle time is reached.
Auto learn period	Time between learn cycles. A learn cycle is a battery calibration operation performed periodically by the controller to determine the condition of the battery.
Average time to empty	One-minute rolling average of the predicted remaining battery life.
Average time to full	Predicted time to charge the battery to a fully charged state based on the one-minute rolling average of the charge current.

B

Battery module version	Current revision of the battery pack module.
Battery replacement	Warning issued by firmware that the battery can no longer support the required data retention time.
Battery retention time	Time, in hours, that the battery can maintain the contents of the cache memory.

Battery status	Operating status of the battery. Possible values are Missing, Optimal, Failed, Degraded (need attention), and Unknown.
Battery type	Possible values are intelligent Battery Backup Unit (BBU), intelligent Battery Backup Unit (iBBU), intelligent Transportable Battery Backup Unit (iTBBU), and ZCR Legacy.
BBU present	A controller property that indicates whether the controller has an on-board battery backup unit to provide power in case of a power failure.
BGI rate	A controller property indicating the rate at which the background initialization of virtual drives will be carried out.
BIOS	Basic Input/Output System. The computer BIOS is stored on a flash memory chip. The BIOS controls communications between the microprocessor and peripheral devices, such as the keyboard and the video controller, and miscellaneous functions, such as system messages.

C

Cache	Fast memory that holds recently accessed data. Use of cache memory speeds subsequent access to the same data. When data is read from or written to main memory, a copy is also saved in cache memory with the associated main memory address. The cache memory software monitors the addresses of subsequent reads to see whether the required data is already stored in cache memory. If it is already in cache memory (a cache hit), it is read from cache memory immediately and the main memory read is aborted (or not started). If the data is not cached (a cache miss), it is fetched from main memory and saved in cache memory.
Cache flush interval	A controller property that indicates how often the data cache is flushed.
Caching	The process of using a high-speed memory buffer to speed up a computer system's overall read/write performance. The cache can be accessed at a higher speed than a drive subsystem. To improve read performance, the cache usually contains the most recently accessed data, as well as data from adjacent drive sectors. To improve write performance, the cache can temporarily store data in accordance with its write back policies.
Capacity	A property that indicates the amount of storage space on a drive or virtual drive.
Coerced capacity	A drive property indicating the capacity to which a drive has been coerced (forced) to make it compatible with other drives that are nominally the same capacity. For example, a 4-GB drive from one manufacturer might be 4,196 MB, and a 4-GB from another manufacturer might be 4,128 MB. These drives could be coerced to a usable capacity of 4,088 MB each for use in a drive group in a storage configuration.
Coercion mode	A controller property indicating the capacity to which drives of nominally identical capacity are coerced (forced) to make them usable in a storage configuration.

Consistency check	An operation that verifies that all stripes in a virtual drive with a redundant RAID level are consistent and that automatically fixes any errors. For RAID 1 drive groups, this operation verifies correct mirrored data for each stripe.
Consistency check rate	The rate at which consistency check operations are run on a computer system.
Controller	A chip that controls the transfer of data between the microprocessor and memory or between the microprocessor and a peripheral device such as a drive. RAID controllers perform RAID functions such as striping and mirroring to provide data protection.
Copyback	<p>The procedure used to copy data from a source drive of a virtual drive to a destination drive that is not a part of the virtual drive. The copyback operation 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). The copyback operation can be run automatically or manually.</p> <p>Typically, a drive fails or is expected to fail, and the data is rebuilt on a hot spare. The failed drive is replaced with a new drive. 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.</p>
Current	Measure of the current flowing to (+) or from (-) the battery, reported in milliamperes.
Current write policy	<p>A virtual drive property that indicates whether the virtual drive currently supports Write Back mode or Write Through mode.</p> <ul style="list-style-type: none"> • In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. • In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Cycle count	The count is based on the number of times the near fully charged battery has been discharged to a level below the cycle count threshold.
D	
Default write policy	<p>A virtual drive property indicating whether the default write policy is Write Through or Write Back.</p> <ul style="list-style-type: none"> • In Write Back mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a transaction. • In Write Through mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data in a transaction.
Design capacity	Designed charge capacity of the battery, measured in milliampere-hour units (mAh).
Design charge capacity remaining	Amount of the charge capacity remaining, relative to the battery pack design capacity.

Design voltage	Designed voltage capacity of the battery, measured in millivolts (mV).
Device chemistry	Possible values are NiMH (nickel metal hydride) and LiON (lithium ion).
Device ID	A controller or drive property indicating the manufacturer-assigned device ID.
Device port count	A controller property indicating the number of ports on the controller.
Drive cache policy	A virtual drive property indicating whether the virtual drive cache is enabled, disabled, or unchanged from its previous setting.
Drive group	A group of drives attached to a RAID controller on which one or more virtual drives can be created. All virtual drives in the drive group use all of the drives in the drive group.
Drive state	A drive property indicating the status of the drive. A drive can be in one of the following states: <ul style="list-style-type: none"> • Unconfigured Good – A drive accessible to the RAID controller but not configured as a part of a virtual drive or as a hot spare. • Hot Spare – A drive that is configured as a hot spare. • Online – A drive that can be accessed by the RAID controller and will be part of the virtual drive. • Rebuild – A drive to which data is being written to restore full redundancy for a virtual drive. • Failed – A drive that was originally configured as Online or Hot Spare, but on which the firmware detects an unrecoverable error. • Unconfigured Bad – A drive on which the firmware detects an unrecoverable error; the drive was Unconfigured Good or the drive could not be initialized. • Missing – A drive that was Online, but which has been removed from its location. • Offline – A drive that is part of a virtual drive but which has invalid data as far as the RAID configuration is concerned. • None – A drive with an unsupported flag set. An Unconfigured Good or Offline drive that has completed the prepare for removal operation.
Drive state drive subsystem	A collection of drives and the hardware that controls them and connects them to one or more controllers. The hardware can include an intelligent controller, or the drives can attach directly to a system I/O bus controller.
Drive type	A drive property indicating the characteristics of the drive.

E

EKM	External Key Management
Estimated time to recharge	Estimated time necessary to complete recharge of the battery at the current charge rate.
Expected margin of error	Indicates how accurate the reported battery capacity is in terms of percentage.

F

Fast initialization	A mode of initialization that quickly writes zeroes to the first and last sectors of the virtual drive. This allows you to immediately start writing data to the virtual drive while the initialization is running in the background.
Fault tolerance	The capability of the drive subsystem to undergo a single drive failure per drive group without compromising data integrity and processing capability. The Intel® RAID Controllers provide fault tolerance through redundant drive groups in RAID levels 1, 5, 6, 10, 50, and 60. They also support hot spare drives and the auto-rebuild feature.
Firmware	Software stored in read-only memory (ROM) or programmable ROM (PROM). Firmware is often responsible for the behavior of a system when it is first turned on. A typical example is a monitor program in a system that loads the full operating system from a drive or from a network and then passes control to the operating system.
Foreign configuration	A RAID configuration that already exists on a replacement set of drives that you install in a computer system. The GUI management utility software allows you to import the existing configuration to the RAID controller, or you can clear the configuration so you can create a new one.
Formatting	The process of writing a specific value to all data fields on a drive, to map out unreadable or bad sectors. Because most drives are formatted when manufactured, formatting is usually done only if a drive generates many media errors.
Full charge capacity	Amount of charge that can be placed in the battery. This value represents the last measured full discharge of the battery. This value is updated on each learn cycle when the battery undergoes a qualified discharge from nearly full to a low battery level.

G

Gas gauge status	Hexadecimal value that represents the status flag bits in the gas gauge status register.
------------------	--

H

Hole	In the GUI management utility, a <i>hole</i> is a block of empty space in a drive group that can be used to define a virtual drive.
Host interface	A controller property indicating the type of interface used by the computer host system; for example, <i>PCIX</i> .
Host port count	A controller property indicating the number of host data ports currently in use.
Host system	Any computer system on which the controller is installed. Mainframes, workstations, and standalone desktop systems can all be considered host systems.

Hot spare	<p>A standby drive that can automatically replace a failed drive in a virtual drive and prevent data from being lost. A hot spare can be dedicated to a single redundant drive group or it can be part of the global hot spare pool for all drive groups controlled by the controller.</p> <p>When a drive fails, the GUI management utility software automatically uses a hot spare to replace it and then rebuilds the data from the failed drive to the hot spare. Hot spares can be used in RAID 1, 5, 6, 10, 50, and 60 storage configurations.</p>
I	
Initialization	The process of writing zeros to the data fields of a virtual drive and, in fault-tolerant RAID levels, generating the corresponding parity to put the virtual drive in a Ready state. Initialization erases all previous data on the drives. Drive groups work without initialization, but they can fail a consistency check because the parity fields have not been generated.
IO policy	A virtual drive property indicating whether Cached I/O or Direct I/O is being used. In Cached I/O mode, all reads are buffered in cache memory. In Direct I/O mode, reads are not buffered in cache memory. Data is transferred to cache and the host concurrently. If the same data block is read again, it comes from cache memory. (The IO Policy applies to reads on a specific virtual drive. It does not affect the read ahead cache.)
L	
Learning cycle	A battery calibration operation performed by a RAID controller periodically to determine the condition of the battery. You can start battery learn cycles manually or automatically.
Learn delay interval	Length of time between automatic learn cycles. You can delay the start of the learn cycles for up to 168 hours (seven days).
Learn mode	Mode for the battery auto learn cycle. Possible values are Auto, Disabled, and Warning.
Learn state	Indicates that a learn cycle is in progress.
Load-balancing	A method of spreading work between two or more computers, network links, CPUs, drives, or other resources. Load balancing is used to maximize resource use, throughput, or response time.
Low-power storage mode	Storage mode that causes the battery pack to use less power, which saves battery power consumption.
LKM	Local Key Management
M	
Manufacturing date	Date on which the battery pack assembly was manufactured.
Manufacturing name	Device code that indicates the manufacturer of the components used to make the battery assembly.

Max error	<p>Expected margin of error (percentage) in the state of charge calculation.</p> <p>For example, when Max Error returns 10 percent and Relative State of Charge returns 50 percent, the Relative State of Charge is more likely between 50 percent and 60 percent. The gas gauge sets Max Error to 100 percent on a full reset. The gas gauge sets Max Error to 2 percent on completion of a learn cycle, unless the gas gauge limits the learn cycle to the +512/-256-mAh maximum adjustment values. If the learn cycle is limited, the gas gauge sets Max Error to 8 percent unless Max Error was already below 8 percent. In this case Max Error does not change. The gas gauge increments Max Error by 1 percent after four increments of Cycle Count without a learn cycle.</p>
Maximum learn delay from current start time	The maximum length of time between automatic learn cycles. You can delay the start of a learn cycle for a maximum of 168 hours (7 days).
Media error count	A drive property indicating the number of errors that have been detected on the drive media.
Migration	The process of moving virtual drives and hot spare drives from one controller to another by disconnecting the drives from one controller and attaching them to another one. The firmware on the new controller will detect and retain the virtual drive information on the drives.
Mirroring	The process of providing complete data redundancy with two drives by maintaining an exact copy of one drive's data on the second drive. If one drive fails, the contents of the other drive can be used to maintain the integrity of the system and to rebuild the failed drive.
Multipathing	The firmware provides support for detecting and using multiple paths from the RAID controllers to the SAS devices that are in enclosures. Devices connected to enclosures have multiple paths to them. With redundant paths to the same port of a device, if one path fails, another path can be used to communicate between the controller and the device. Using multiple paths with load balancing, instead of a single path, can increase reliability through redundancy.

N

Name	A virtual drive property indicating the user-assigned name of the virtual drive.
Next learn time	Time at which the next learn cycle starts.
Non-redundant configuration	A RAID 0 virtual drive with data striped across two or more drives but without drive mirroring or parity. This provides for high data throughput but offers no protection in case of a drive failure.
NVRAM	Acronym for nonvolatile random access memory. A storage system that does not lose the data stored on it when power is removed. NVRAM is used to store firmware and configuration data on the RAID controller.
NVRAM present	A controller property indicating whether an NVRAM is present on the controller.
NVRAM size	A controller property indicating the capacity of the controller's NVRAM.

O

Offline A drive is offline when it is part of a virtual drive but its data is not accessible to the virtual drive.

P

Patrol read A process that checks the drives in a storage configuration for drive errors that could lead to drive failure and lost data. The patrol read operation can find and sometimes fix any potential problem with drives before host access. This enhances overall system performance because error recovery during a normal I/O operation might not be necessary.

Patrol read rate The user-defined rate at which patrol read operations are run on a computer system.

Predicted battery capacity status (hold 24hr charge) Indicates whether the battery capacity supports a 24-hour data retention time.

Product info A drive property indicating the vendor-assigned model number of the drive.

Product name A controller property indicating the manufacturing name of the controller.

R

RAID A group of multiple, independent drives that provide high performance by increasing the number of drives used for saving and accessing data.

A RAID drive group improves input/output (I/O) performance and data availability. The group of drives appears to the host system as a single storage unit or as multiple virtual drives. Data throughput improves because several drives can be accessed simultaneously. RAID configurations also improve data storage availability and fault tolerance. Redundant RAID levels (RAID levels 1, 5, 6, 10, 50, and 60) provide data protection.

RAID 0 Uses data striping on two or more drives to provide high data throughput, especially for large files in an environment that requires no data redundancy.

RAID 00 Uses data striping on two or more drives in a spanned drive group to provide high data throughput, especially for large files in an environment that requires no data redundancy.

RAID 1 Uses data mirroring on pairs of drives so that data written to one drive is simultaneously written to the other drive. RAID 1 works well for small databases or other small applications that require complete data redundancy.

RAID 5 Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access.

RAID 6	Uses data striping and parity data across three or more drives (distributed parity) to provide high data throughput and data redundancy, especially for applications that require random access. RAID 6 can survive the failure of two drives.
RAID 10	A combination of RAID 0 and RAID 1 that uses data striping across two mirrored drive groups. It provides high data throughput and complete data redundancy.
RAID 50	A combination of RAID 0 and RAID 5 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy.
RAID 60	A combination of RAID 0 and RAID 6 that uses data striping across two drive groups with parity data. It provides high data throughput and complete data redundancy. RAID 60 can survive the failure of two drives in each RAID set in the spanned drive group.
RAID level	A virtual drive property indicating the RAID level of the virtual drive. Intel® RAID Controllers support RAID levels 0, 1, 5, 6, 10, 50, and 60.
RAID migration	A feature in RAID subsystems that allows changing a RAID level to another level without powering down the system.
Raw capacity	A drive property indicating the actual full capacity of the drive before any coercion mode is applied to reduce the capacity.
Read policy	A controller attribute indicating the current Read Policy mode. In Always Read Ahead mode, the controller reads sequentially ahead of requested data and stores the additional data in cache memory, anticipating that the data will be needed soon. This speeds up reads for sequential data, but there is little improvement when accessing random data. In No Read Ahead mode (known as Normal mode in WebBIOS), read ahead capability is disabled.
Rebuild	The regeneration of all data to a replacement drive in a redundant virtual drive after a drive failure. A drive rebuild normally occurs without interrupting normal operations on the affected virtual drive, though some degradation of performance of the drive subsystem can occur.
Rebuild rate	The percentage of central processing unit (CPU) resources devoted to rebuilding data onto a new drive after a drive in a storage configuration has failed.
Reclaim virtual drive	A method of undoing the configuration of a new virtual drive. If you highlight the virtual drive in the Configuration Wizard and click Reclaim, the individual drives are removed from the virtual drive configuration.
Reconstruction rate	The user-defined rate at which a drive group modification operation is carried out.
Redundancy	A property of a storage configuration that prevents data from being lost when one drive fails in the configuration.

Redundant configuration	<p>A virtual drive that has redundant data on drives in the drive group can be used to rebuild a failed drive. The redundant data can be parity data striped across multiple drives in a drive group, or it can be a complete mirrored copy of the data stored on a second drive.</p> <p>A redundant configuration protects the data in case a drive fails in the configuration.</p>
Relative state of charge	Predicted remaining battery capacity expressed as a percentage of Full Charge Capacity.
Remaining capacity	The amount of remaining charge capacity of the battery as stated in milliamp hours. This value represents the available capacity or energy in the battery at any given time. The gas gauge adjusts this value for charge, self-discharge, and leakage compensation factors.
Revertible hot spare	When you use the Replace Member procedure, after data is copied from a hot spare to a new drive, the hot spare reverts from a rebuild drive to its original hot spare status.
Revision level	A drive property that indicates the revision level of the drive's firmware.
Run time to empty	Predicted remaining battery life at the present rate of discharge in minutes.

S

SAS	Acronym for Serial-Attached SCSI. SAS is a serial, point-to-point, enterprise-level device interface that leverages the Small Computer System Interface (SCSI) protocol set. The SAS interface provides improved performance, simplified cabling, smaller connectors, lower pin count, and lower power requirements when compared to parallel SCSI.
SATA	Acronym for Serial Advanced Technology Attachment. A physical storage interface standard. SATA is a serial link that provides point-to-point connections between devices. The thinner serial cables allow for better airflow within the system and permit smaller chassis designs.
SCSI device type	A drive property indicating the type of the device, such as drive.
Serial no.	A controller property indicating the manufacturer-assigned serial number.
Strip size	The portion of a stripe that resides on a single drive in the drive group.
Stripe size	A virtual drive property indicating the length of the interleaved data segments that the RAID controller writes across multiple drives, not including parity drives. For example, consider a stripe that contains 64 KB of drive space and has 16 KB of data residing on each drive in the stripe. In this case, the stripe size is 64 KB and the strip size is 16 KB. The user can select the stripe size.

Striping	<p>A technique used to write data across all drives in a virtual drive.</p> <p>Each stripe consists of consecutive virtual drive data addresses that are mapped in fixed-size units to each drive in the virtual drive using a sequential pattern. For example, if the virtual drive includes five drives, the stripe writes data to drives one through five without repeating any of the drives. The amount of space consumed by a stripe is the same on each drive. Striping by itself does not provide data redundancy. Striping in combination with parity does provide data redundancy.</p>
Subvendor ID	<p>A controller property that lists additional vendor ID information about the controller.</p>
T	
Temperature	<p>Temperature of the battery pack, measured in Celsius.</p>
U	
Uncorrectable error count	<p>A controller property that lists the number of uncorrectable errors detected on drives connected to the controller. If the error count reaches a certain level, a drive will be marked as failed.</p>
V	
Vendor ID	<p>A controller property indicating the vendor-assigned ID number of the controller.</p>
Vendor info	<p>A drive property listing the name of the vendor of the drive.</p>
Virtual drive	<p>A storage unit created by a RAID controller from one or more drives. Although a virtual drive can be created from several drives, it is seen by the operating system as a single drive. Depending on the RAID level used, the virtual drive can retain redundant data in case of a drive failure.</p>
Virtual drive state	<p>A virtual drive property indicating the condition of the virtual drive. Examples include Optimal and Degraded.</p>
W	
Write-back	<p>In Write-Back Caching mode, the controller sends a data transfer completion signal to the host when the controller cache has received all of the data in a drive write transaction. Data is written to the drive subsystem in accordance with policies set up by the controller.</p> <p>These policies include the amount of dirty/clean cache lines, the number of cache lines available, and elapsed time from the last cache flush.</p>
Write policy	<p>See Default write policy.</p>
Write-through	<p>In Write-Through Caching mode, the controller sends a data transfer completion signal to the host when the drive subsystem has received all of the data and has completed the write transaction to the drive.</p>