

Intel® Rapid Storage Technology Release 16.x with Intel® Optane™ Memory

OEM Technical Guide

For the Intel® Rapid Storage Technology Release Version 16.0.1, 16.0.2, 16.3 with Intel® Optane™ Memory Support (For OEMs, ODMs, IHVs, ISVs, IBVs, SIs)

Revision 0.93

Intel Confidential



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at Intel.com, or from the OEM or retailer.

No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses.

The products described 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.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Copies of documents which have an order number and are referenced in this document may be obtained by calling 1-800-548-4725 or visit www.intel.com/design/literature.htm.

By using this document, in addition to any agreements you have with Intel, you accept the terms set forth below.

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

Intel, the Intel logo, are trademarks of Intel Corporation in the U.S. and/or other countries.

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

Copyright © 2018, Intel Corporation. All rights reserved.

Contents

1 About This Document 11
1.1 Purpose and Scope of this Document 11
1.1.1 CPU Attached Storage 11
1.1.2 Intel® Optane™ Memory Support 11
1.1.3 Intel® Smart Response Technology (Intel® SRT) (EOL) 11
1.1.4 Intel® RST Premium Feature Support 11
1.2 Intended Audience 11

2 Intel® Rapid Storage Technology 13
2.1 Overview 13
2.1.1 Product Release Numbering Scheme 13
2.1.2 RAID Levels 14
2.1.3 Supported Platforms for This Release 15
2.1.4 Intel® RST Key Feature/Functionality Matrix 16
2.1.5 Supported Operating Systems for This Release 18
2.1.6 Supported MSFT* Performance Debug Tools 18

3 Intel® Rapid Storage Technology Suite 19
3.1 Intel® Rapid Storage Technology Software 19
3.2 Intel® Rapid Storage Technology Option ROM 20
3.3 Intel Pre-OS RAID Configuration Utilities 20
3.3.1 RCfgSata Utility for MS-DOS* and UEFI 20
3.4 RSTCLI (32/64 bit) Windows Utilities 22
3.5 UEFI System BIOS and Intel® RST UEFI/RAID Package 30
3.5.1 Specification References 30
3.5.2 What Intel® RST Provides to OEMs/BIOS Vendors 31
3.5.3 UEFI System BIOS Requirements for Platform Compatibility with Intel® RST UEFI 33
3.5.4 How-to-Enable the Platform with Intel® RST UEFI Driver/HII_GUI 35
3.5.5 Known Compatibility Issues with UEFI Self Certification Test (UEFI SCT) Tool 36
3.6 Pre-OS 38
3.6.1 BIOS PCH SATA Controller Modes 38

4 New in Release Version 16.x 39
4.1 Feature Specification Update in Release 16.3 39
4.1.1 Intel® Optane™ Memory User Pinning 39
4.2 Feature Specification Update in Release 16.0.2 39
4.2.1 Acceleration of Data Drives with Intel® Optane™ Memory 39
4.3 New Features/Specifications in Release 16.0.1 39
4.3.1 CPU Attached Storage- PCIe NVMe 39
4.4 What's New 39
4.4.1 Intel® Optane™ Memory UI/Installer 39

5 CPU Attached Storage: PCIe NVMe 42
5.1 Overview 42
5.2 Limitations and Minimum Requirements: 42
5.2.1 HW 42
5.2.2 SW/FW 42
5.3 Usage models 42

6 System Acceleration with Intel® Optane™ Memory 44



6.1 Limitations: 44
6.1.1 HW limits 44
6.1.2 Functional limits 44
6.2 Minimum Requirements: 45
6.2.1 HW 45
6.2.2 SW/FW 45
6.2.3 Operating System 45
6.2.4 "Optane™ Memory Ready" Logo: Intel® RST 5MB Metadata Requirement 45
6.3 User Pinning Files/Folders/Applications 46
6.3.1 Pinning using the Intel® Optane™ Memory UI 46
6.3.2 User Pinning using Windows File Explorer application 48
6.4 Configuring Intel® Optane™ Memory (Post-Factory Build Environment[1]) 54
6.4.1 Enable Optane™ Volume on a Brand New System 55
6.4.2 System Upgrade (Windows*10 OS Already Installed) 60
6.4.3 System Upgrade (Windows* 10 OS Already Installed with BIOS in AHCI mode) 64
6.5 Intel® Optane™ Memory Member Disk Upgrades 66
6.5.1 Slow Disk (SATA HDD, SSD, SSHD) Upgrade/Replacement 67
6.5.2 Intel® Optane™ Memory Module Upgrade/Replacement 68
6.6 Managing Intel® Optane™ Memory 69
6.6.1 During the Pre-OS (Pre-Boot) 69
6.6.2 During the OS Runtime 70
6.6.3 Roaming Optane™ Volumes 72
6.7 I/O Error Handling 73
6.7.1 Host I/O Failures 73
6.7.2 RST I/O Failures: 73
6.7.3 Missing Media Failures: 74
6.7.4 S.M.A.R.T. Events 77
6.7.5 Fast Media: Block Cache Errors 78
6.8 Intel® Optane™ Memory Member Disk Failures Requiring Replacement of Disk 79
6.8.1 Slow Disk (SATA HDD, SSD, SSHD) Replacement 79
6.8.2 Intel® Optane™ Memory Module Replacement 80
6.9 Windows Recovery Environment Support 82
6.9.1 Setting Up the Windows Recovery Environment (WinRE) 82
6.9.2 WinRE Menu Items 83

7 OEM HVM Factory Support for Intel® Optane™ System Acceleration 85
7.1 Intel® Optane™ Pre-Factory Recommendations for the 'ReferencePC' 85
7.1.1 Optane™ Recommendations for the 'ReferencePC' 85
7.2 Intel® Optane™ Factory-Enabled System Build 88
7.2.1 Requirements/Considerations 89
7.2.2 PHASE 1: OEM Pre-Optane™-Config Factory Processes 89
7.2.3 PHASE 2: Optane™ Storage Sub-System Configuration 89
7.2.4 PHASE 3: OS and SW Download and Installation and OOBE Preparation 90
7.3 Intel® Optane™ Post-Factory-Enabled Builds 91
7.3.1 Requirements/Considerations 91
7.3.2 Example Build Flow 91
7.4 OEM Application Pinning in the Factory Flow 91
7.4.1 Manufacturing Considerations for Pinning 92
7.4.2 Using the JSON files in the Factory Flow 93

7.4.3 OEM Pinning Methods 93
8 Intel Rapid Storage Technology for PCIe NVMe Storage Devices 97
8.1 OEM System BIOS Requirements 97
8.2 General Requirements 97
8.3 Feature Limitations 97
8.4 PCIe NVMe Device Usage Model 98
8.5 Intel® RST for PCIe NVMe Storage Use cases 98
8.6 Intel® Rapid Storage Technology UEFI Compliance Utility for PCIe Storage 100
9 Intel Rapid Storage Technology for PCIe AHCI Storage Devices 102
9.1 OEM System BIOS Requirements 102
9.2 General Requirements 102
9.3 Warnings 102
9.3.1 Features Limitations 102
9.3.2 PCIe Device Usage Model 103
9.4 Intel Rapid Storage Technology for PCIe Storage Use cases 103
9.5 Intel® Customer Reference Board BIOS Settings 103
9.6 Intel® Rapid Storage Technology UEFI Compliance Utility for PCIe Storage 104
10 How to Enable the Platform for Intel® RST Support of BIOS Fast Boot 106
10.1.1 OEM System BIOS Vendors' Requirements 106
10.1.2 Supported System Configurations 106
11 Creating a RAID Volume 108
11.1 Minimum Requirements 108
11.2 Feature Limitations 108
11.3 Using Intel® Rapid Storage Technology UI 108
11.4 Using Intel® Rapid Storage Technology Legacy Option ROM User Interface . 109
11.5 Using Intel® Rapid Storage Technology UEFI User Interface 110
11.6 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows*) 115
12 Deleting a RAID Volume 116
12.1 Using the Windows* User Interface Utility 116
12.2 Using the Option ROM User Interface 116
12.3 Using the Intel® Rapid Storage Technology UEFI User Interface 117
12.4 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows*) 117
13 Common RAID Setup Procedures 118
13.1 Build a SATA RAID 0, 1, 5 or 10 System 118
13.1.1 Using the Legacy OROM User Interface 118
13.1.2 Using the UEFI HII User Interface 119
13.2 Build a "RAID Ready" System 120
13.3 Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System 120
13.4 Migrate an Existing Data Hard Drive to a RAID 0 or RAID 1 Volume 121
13.5 Migrating From one RAID Level to Another 121
13.6 Create a RAID Volume on Intel® SATA Controller While Booting to Different Controller 122
13.7 Build a RAID 0 or RAID 1 System in an Automated Factory Environment 123
13.7.1 Part 1: Create Master Image 123
13.7.2 Part 2: Apply Master Image 123
14 RAID Volume Data Verification and Repair Feature 125
14.1 Verify and Repair Volume Feature 125
14.2 Verify and Repair Scheduler 125



15 Intel® Rapid Recover Technology 127
15.1 Creating a Recovery Volume Through the RAID Option ROM 127
15.2 Creating a Recovery Volume Using Intel® Rapid Storage Technology UEFI User Interface 128
15.3 Creating a Recovery Volume Through Intel® RST UI 128
15.4 Changing Recovery Volume Modes 129
15.5 Update Recovery Volume in On Request Update Policy 129
15.6 Access Recovery Drive Files 129
15.7 Hide Recovery Drive Files 130
15.8 Scenarios of Recovering Data 130
15.9 System Running from Recovery Drive 132
15.10 Drive Offline or Missing 132

16 Pre-OS Installation of Intel® Rapid Storage Technology Driver 134
16.1 Pre-OS Driver Installation Using the "Load Driver" Method 134

17 Determining the Version of the RAID Driver 135
17.1 Using Intel® Rapid Storage Technology User Interface (UI) 135
17.2 Using Intel® RST File Properties (Alternate) 135
17.3 Determining the Version of the Option ROM 135
17.3.1 Using Intel® Rapid Storage Technology UI 135
17.3.2 Using Intel® RST Option ROM User Interface 136
17.3.3 Using the EFI Shell 136

18 Un-Installation 137
18.1 Un-Installing Intel® RST Software (except the RAID Driver) 137
18.2 Disabling the RAID Driver by Disabling the RAID Controller 137

19 Registry Customizations 139
19.1 Zero Power ODD Settings 139
19.2 E-mail Notification UI Visible Enable/Disable 140
19.3 Disabling Maximized Mode Option for Intel® SRT 141
19.4 Rebuild On Hot Insert 141
19.5 SATA Asynchronous Notification 142
19.6 Runtime D3 (RTD3) 142
19.7 Hybrid Hinting 143
19.7.1 Instructions to Disable Hybrid Hinting 143
19.7.2 Hybrid Hint Reset 143
19.7.3 Disable Hybrid Hinting During Hibernation 143

20 Power Savings with Intel® Rapid Storage Technology 145
20.1 Link Power Management (LPM) 145
20.1.1 Instructions to Disable/Enable LPM 145
20.1.2 LPM Updates in 14.0 (APS, SIPM) 146
20.2 Runtime D3 (RTD3) 147
20.2.1 Adapter RTD3 Support 147
20.3 New for RST 14.0 Release 148
20.3.1 RTD3 Support - RAID HDD/SSD/SSHDs Unit Support 148
20.3.2 RTD3 Support - RAID w Mixed RTD3 Capable/Non-Capable Ports 148
20.3.3 RTD3 Support - RRT 148

20.3.4 RTD3 Support - SRT 148
20.3.5 RTD3 Support - Hot Spares 148
20.3.6 RTD3 Support – Migrations and Rebuilds 148
20.4 DEVSLP 148
20.4.1 DEVSLP Registry Key Setting 150
20.5 DevSleep Tool 150
20.5.1 CsDeviceSleepIdleTimeoutInMS 151
20.5.2 DeviceSleepIdleTimeoutInMS 151
20.5.3 DeviceSleepExitTimeoutInMS 151
20.5.4 MinimumDeviceSleepAssertionTimeInMS 152
20.5.5 DevSleep Tool Usage 152
20.6 L1.2 Support 153
20.7 InstantGo* Device Notification Support 153
20.7.1 Requirements 154
20.7.2 Detail Description 154
20.7.3 Registry Settings 154
20.8 New in 14.5 Release 155
20.8.1 Connected Standby Power State Support for SSHD 155
20.8.2 CONNECTED STANDBY Power Model 156
20.8.3 Adaptive D3 Idle Timeout 156
20.8.4 Connected Standby Power Model Support for SRT 157
20.8.5 SATA Link Power Management Support 159

21 Legacy RAID Option ROM and Utilities 160

22 HDD Password Support with RAID Volumes 161
22.1 HDD Password Use Cases 161
22.2 Unlocking Password Protected Disks 162

23 Intel® Rapid Storage Technology UI 163
23.1 Introduction 163
23.1.1 Getting Started 163
23.1.2 Understanding the Application 165
23.1.3 Notification Area 167
23.2 Storage System Status 168
23.2.1 Understanding the Status 168
23.2.2 Storage System View 169
23.3 Creating a Volume 171
23.3.1 Volume Requirements 171
23.3.2 Creation Process 172
23.3.3 Creating Additional Volumes 176
23.4 Managing the Storage System 177
23.4.1 Managing Arrays 178
23.4.2 Managing Volumes 181
23.4.3 Managing Disks 195
23.4.4 Managing Ports 200
23.4.5 Managing ATAPI Devices 200
23.4.6 Managing Solid-State Hybrid Drives (SSHD) 201
23.5 Accelerating the Storage System 201
23.5.1 Cache Device Properties 201
23.5.2 Enabling Acceleration 203
23.5.3 Disabling Acceleration 205
23.5.4 Changing Acceleration Mode 205
23.5.5 Accelerating a Disk or Volume 206



23.5.6 Resetting a Cache Device to Available 207
23.5.7 Disassociating the Cache Memory 207
23.6 Preferences 208

24 Opal Drive Support 210
24.1 Opal ATA Commands Supported 210
24.2 Intel® Rapid Storage Technology UI Support 210
24.3 Opal SCSI Commands Supported 210
24.4 Opal UEFI Protocol Supported 210
24.5 Intel® RST PCIe NVMe Support for Opal 210
24.6 NVMe Device Requirements for Intel® RST Opal Support 211

25 eDrive Support 212
25.1 What is an eDrive? 212
25.2 SCSI Standard Inquiry Command Response 212
25.3 Configurations Supported 212

26 Using the BCFS to Differentiate Platform SKUs 213
26.1 Configuring the Platform's RAID Related Features 213
26.1.1 Configuring the Standard Supported RAID Levels 213
26.1.2 Configuring Intel® RRT Related RAID Features 214
26.1.3 Configuring the Behavior of the OROM UI and Banner 214
26.1.4 Configuring Intel® RST UI Capabilities 215
26.1.5 Configuring the Platform to Support Intel® SRT (EOL'd) 216
26.1.6 Configuring the Platform to Support Intel® Optane™ Memory System Acceleration 216
26.1.7 Configuring the Platform to Support CPU Attached Storage 216
26.2 BIOS Control Feature Set 217

27 Testing, Certification Notes 219

28 Glossary 220

29 Troubleshooting 227
29.1 Failed Volumes 227
29.2 Degraded Volumes 229
29.3 Other Volume States 232
29.4 Disk Events 236
29.5 Caching Issues 238
29.6 Software Errors 241

Appendix A RST SATA Port Bitmap Implementation 243

A.1 Legacy OROM 243

A.2 UEFI Driver 243

A.2.1 EFI_DEVICE_PATH_PROTOCOL 244

Appendix B Common Storage Management Interface Support (CSMI) 245

Appendix C Drive and Volume Encryption Support 246

C.1 ATA Security Commands and HDD Password Support 246

C.2 Self-Encrypting Drives (SED) 246
C.3 Solid State Hybrid Drives (SSHD's) with Encryption 246
C.4 RAID Volume and Drive Partition Encryption 246
Appendix D Remapping Guidelines for RST PCIe Storage Devices 248
D.1 Remapping Reference Documentation 248
D.1.1 Remapping HW and BIOS Requirements 249
D.1.2 Remapping Configuration Rules 249
D.1.3 PCH-H Remapping Configurations 251
D.1.4 Configurations With 1 x2 PCIe Port Remapped 251
D.1.5 Configurations With 2 x2 PCIe Ports Remapped 252
D.1.6 Configurations With 3 x2 PCIe Ports Remapped 253
D.1.7 Configurations with 1, 2, and 3 x4 PCIe Ports Remapped 254
D.1.8 Configurations with (1 x2 + 1 x4) PCIe Ports Remapped 255
D.1.9 Configurations with (2 x2 + 1 x4) PCIe Ports Remapped 256
D.1.10 Configurations with (1 x2 + 2 x4) PCIE Ports Remapped 257
D.1.11 PCH-LP Premium-U Remapping Configurations 257
D.1.12 Configurations with 1 x2 PCIe Port Remapped 258
D.1.13 Configurations With 2 x2 PCIe Ports Remapped 259
D.1.14 Configurations with 1 and 2 x4 PCIe Ports Remapped 260
D.1.15 Configurations with (1 x2 + 1 x4) PCIe Ports Remapped 261
D.1.16 PCH-LP Premium-Y Remapping Configurations 261
D.1.17 Configurations With 1 x2 PCIe Port Remapped 262
D.1.18 Configurations With 2 x2 PCIe Ports Remapped 263
D.1.19 Configurations With x4 PCIe Ports Remapped 264
D.2 Examples of Configurations to Meet Design Specifications 265
D.2.1 Example #1: SPT-H HM170 SKU With 1x2 + 1x4 + 1 SATA 265

§ §



Revision History

Revision Number	Description	RST Release Version	Revision Date
0.93	▪Added 'User Pinning' sections	16.3	Feb 2018
0.92	▪Added release version 16.0.2 ▪Includes Data Drive Acceleration ▪Updated chapters 4 and 6 to include data drive acceleration procedures	16.0.2	Jan 2018
0.9	▪Adjusted release version to 16.0.1 ▪Removed Pinning feature	16.0.1	Jan 2018
0.88	▪Updated Chapter 7 "OEM Pinning Methods" section with corrected registry path as follows in .reg file example: [HKEY_LOCAL_MACHINE\MyKey2\Software\Microsoft\Windows\CurrentVersion\RunOnce] ▪	16.0	Nov 2017
0.87	▪Updated section 6.3 ▪Updated section 7.4	16.0	Nov 2017
0.85	▪Added table 4-3 ▪Updated section 7.4	16.0	Oct 2017
0.8	▪Initial Release	16.0	Oct 2017

§ §



1 About This Document

1.1 Purpose and Scope of this Document

This document will assist customers in evaluating, testing, configuring, and enabling CPU Attached Storage, Intel® Optane™ Memory, RAID 0/1/5/10, and AHCI functionality on Intel-based platforms using the *Intel® Rapid Storage Technology* (Intel® RST) software for the chipset components as listed in the product's Readme.txt file.

This document also describes installation procedures, caching acceleration techniques, other RST features, RAID volume management such as creating, deleting, and modifying volumes, common usage models, and any special notes necessary to enable customers to develop their RST-compatible computer systems.

1.1.1 CPU Attached Storage

- Minimum requirements
- How to enable
- How to manage
- High Volume Manufacturing requirements and guidelines

1.1.2 Intel® Optane™ Memory Support

- Minimum requirements
- How to enable
- How to manage
- High Volume Manufacturing requirements and guidelines

1.1.3 Intel® Smart Response Technology (Intel® SRT) (EOL)

- Minimum requirements
- How to enable
- How to manage
- High Volume Manufacturing requirements and guidelines

1.1.4 Intel® RST Premium Feature Support

- RAID 0/1/5/10
- Rapid Recovery Technology
- Manufacturing tools
- Power savings support

1.2 Intended Audience

This document is targeted for the following audience under proper Intel NDA guidelines and thus can be shared amongst the following provided they have proper NDA in place (e.g. an OEM can share with an ODM that has Intel NDA in place) (e.g. an OEM can share with an ODM that has Intel NDA in place):

- OEMs (Original Equipment Manufacturers)

- ODMs (Original Design Manufacturers)
- IHVs (Independent Hardware Vendors)
- ISVs (Independent Software Vendors)
- IBVs (Independent BIOS Vendors
- SIs (System Integrators)

§ §

2 Intel® Rapid Storage Technology

Intel® Rapid Storage Technology (Intel® RST) provides added performance and reliability for supported systems equipped with serial ATA (SATA) hard drives (HDD's) and/or solid state drives (SSD's) and PCIe AHCI/NVMe SSD's to enable an optimal PC storage solution. It offers value-add features such as RAID and advanced Serial ATA* capabilities for the Microsoft* Windows* operating systems (for detailed OS support, review the Release Notes for each software release). Also overall system responsiveness is boosted by the Intel® Optane™ Memory system acceleration caching feature or the legacy Intel® Smart Response Technology (Intel® SRT) system acceleration caching solution.

The RAID feature supports RAID level 0 (striping), RAID level 1 (mirroring), RAID level 5 (striping with parity) and RAID level 10 (striping and mirroring). A configuration supporting two RAID levels can also be achieved by having two volumes in a single RAID array that use Intel® RST. This configuration is known as a Matrix array. The RAID capability addresses the demand for high-performance or data-redundant desktop and mobile platforms.

The product also includes premium-RST features as well as support for other premium-platform technologies.

2.1 Overview

2.1.1 Product Release Numbering Scheme

The product release **version** is divided into 4 sections or numbers (**AA**.**B**.**CC**.**DDDD**, e.g. **12.0.0.1001**).

Number / Section	Description
AA: **Major Release Number**	This section represents the major release version of the product. It usually is usually associated with a major change in features or new platform/chipset launch.
B: **Minor (Maintenance) Release Number**	This section represents the minor release version of the product. If this number is **non-zero**, then the release is a minor release of the **AA** major release version. This can represent a maintenance release with several bug fixes or it can align with a platform refresh as example.
CC: **Hot Fix Release Number**	This section represents customer specific hot fixes. If this number is **non-zero**, then the release is a customer specific hot fix release to resolve a customer specific issue.
DDDD: **Release Build Number**	This section represents the build number of release **AA.B.CC**, **Note:** for production releases, the build number always begins with the number '1' (e.g. AA.B.CC.1001)

2.1.2 RAID Levels

RAID 0 (striping)

RAID level 0 combines two to six drives so that all data is divided into manageable blocks called strips. The strips are distributed across the array members on which the RAID 0 volume resides. This improves read/write performance, especially for sequential access, by allowing adjacent data to be accessed from more than one hard drive simultaneously. However, data stored in a RAID 0 volume is not redundant. Therefore, if one hard drive fails, all data on the volume is lost.

The RAID 0 volume appears as a single physical hard drive with a capacity equal to twice the size of the smaller hard drive.

The Intel® SATA AHCI/RAID controllers with Intel Rapid Storage Technology allows up to six** drives to be combined into a single RAID 0 array, providing additional scaling of storage performance.

**Note: the number of drives supported in a RAID 0 array is dependent upon the chipset model. Consult the specification for your chipset to determine the maximum number of drives supported in a RAID array.

RAID 1 (mirroring)

RAID level 1 combines two hard drives so that all data is copied concurrently across the array members that the RAID 1 volume resides on. In other words, the data is mirrored across the hard drives of the RAID 1 volume. This creates real-time redundancy of all data on the first drive, also called a mirror. RAID 1 is usually used in workstations and servers where data protection is important.

The RAID 1 volume appears as a single physical hard drive with a capacity equal to that of the smaller hard drive.

RAID 5 (striping with parity)

RAID level 5 combines three to six drives so that all data is divided into manageable blocks called strips. RAID 5 also stores parity, a mathematical method for recreating lost data on a single drive, which increases fault tolerance. The data and parity are striped across the array members. The parity is striped in a rotating sequence across the members.

Because of the parity striping, it is possible to rebuild the data after replacing a failed hard drive with a new drive. However, the extra work of calculating the missing data will degrade the write performance to the volumes. RAID 5 performs better for smaller I/O functions than larger sequential files.

RAID 5, when enabled with volume write-back cache with Coalescer, will enhance write performance. This combines multiple write requests from the host into larger more efficient requests, resulting in full stripe writes from the cache to the RAID5 volume.

RAID 5 volume provides the capacity of (N-1) * smallest size of the hard drives, where N >= 3 and <=4.

For example, a 3-drive RAID 5 will provide capacity twice the size of the smallest drive. The remaining space will be used for parity information.

RAID 10 (striping and mirroring)

RAID level 10 uses four hard drives to create a combination of RAID levels 0 and 1. The data is striped across a two-disk array forming a RAID 0 component. Each of the drives in the RAID 0 array is mirrored to form a RAID 1 component. This provides the performance benefits of RAID 0 and the redundancy of RAID 1. The RAID 10 volume appears as a single physical hard drive with a capacity equal to two drives of the four drive configuration (the minimum RAID 10 configuration). The space on the remaining two drives will be used for mirroring.

*RAID 10 is currently NOT supported on RST PCIe Storage member disks.

2.1.2.1 Typical Usage Model for RAID Levels

RAID 0

This provides end-users the performance necessary for any disk-intensive applications; these include video production and editing, image editing, and gaming applications.

RAID 1

This provides end-users with data redundancy by mirroring data between the hard drives.

RAID 5

This provides end-users with good performance and data redundancy by striping data and parity across all the hard drives. The write performance is enhanced with volume write-back cache.

RAID 10

This provides end-users with the benefits of RAID 0 (performance) and RAID 1 (data mirroring).

2.1.3 Supported Platforms for This Release

Note: All platform/chipset SKUs listed in this section support AHCI mode. Premium chipsets also support RST Premium/Optane™ mode. Certain non-Premium SKUs (see following table) support AHCI along with RST/Optane™ mode.

Table 2-1: Platform/Chipset SKUs Supported in this Release

Initial RST Release Version	Chipset Name	Platform / PCH / (Segment)	PCH SKU Details
16.x	Intel® 300/240 Series Chipset Family	Cannon Lake (CNL)/Coffee Lake (CFL) PCH: Cannon Point-H (CNP-H) (DT, HEDT)	- H310[A] - H370 - Z390 - Q370 - Q360[O] - B360[O] - X399
		CNL/CFL PCH: CNP-H (WS)	- C246
		CNL/CFL PCH: CNP-H (Mobile Halo)	- QM370 - HM370 - CM246
	Intel® 8th Generation Core Processor Family Platform I/O SATA AHCI/RAID Controller	CNL PCH: CNP-LP (Mobile LP)	- Premium-U - Base-U[A]
15.8 / 15.9	Intel® 200 Series Chipset Family	Coffee Lake (CFL-S) PCH: Kaby Point (KBP-H) (Desktop)	- Z370***

	N - 1					
		15.7			Intel® 200 Series Chipset Family	Basin Falls (w/ KBL-X, SKL-X CPU) PCH: KBP-H (HEDT)	- X299
					Intel® 8th Generation Core Processor Family Platform I/O SATA AHCI/RAID Controller	Kaby Lake Refresh (KBL-R) PCH: SPT-LP (Mobile-LP)	- Base-U - Premium-U - Premium-Y
			15.5		Intel® 100/C230 Series Chipset Family	Greenlow-Refresh (w/ KBL CPU) PCH: SPT PCH-H (WS)	- C236
				15.2	Intel® 200 Series Chipset Family	KBL PCH: KBP-H (Desktop)	- Q250(O) - B250(O) - Z270 - H270 - Q270
					Intel® 100/C230 Series Chipset Family	(w/ KBL CPU) PCH: SPT-H (Mobile Halo)	- HM175 - QM175 - CM238
				15.0	Intel® 7th Generation Core Processor Family Platform I/O SATA AHCI/RAID Controller	(w/ KBL CPU) PCH: SPT-LP (Mobile-LP)	- Base-U(A) - Premium-U - Premium-Y

(A) This SKU of the PCH supports AHCI mode Only
(O) This SKU of the PCH supports both AHCI mode and Optane™ non-Premium mode (non-RAID)
***This SKU is considered platform generation N (not N – 1) of CNL/CFL family

2.1.4 Intel® RST Key Feature/Functionality Matrix

Table 2-2. Key Feature/Functionality Key

Short Name	Feature Name
RAID[1]	SATA controller supports RST RAID feature set (RAID 0, 1, 5[2], 10[3])
ZPODD	Zero Power Optical Disk Drive feature
DDA	Data Drive Acceleration
UEFI	UEFI pre-OS driver
BFB	BIOS Fast Boot
DEVSLP	SATA Device Sleep
RTD3[4]	Runtime D3 (Includes SATA, PCIe NVMe pass-through disk)
PCIe-PCH	Intel® RST PCIe (AHCI/NVMe) PCH attached Storage device support (Remapping Required)
OPTANE[5]	System Acceleration with Intel® Optane™ Memory feature
CPU-AS	CPU Attached Storage – SW remapped Intel-only PCIe NVMe devices
HMB	Host Memory Buffer

[1] *Supports SATA disks and remapped PCH and CPU attached PCIe NVMe/AHCI disks*

[2] *PCH SKU dependent requiring 3 PCH SATA controller ports (SATA and remapped PCIe)*
[3] *PCH SKU dependent requiring 4 PCH SATA controller ports (SATA only)*
[4] *Supports SATA disks and only PCIe NVMe disks (no PCIe AHCI, and no RAID/SRT/Optane configurations)*
[5] *Added support for data drive acceleration beginning release version 16.0.1*

Table 2-3: Key Feature/Functionality Matrix

Platform / Chipset	Segment	PCH SKU	RAID	ZPODD	SRT (EOL)	BFB	DEVSLP	RTD3#	PCIe-PCH	OPTANE	HMB	CPU-AS	DDA
Cannon Lake(CNL)/Coffee Lake(CFL) PCH: Cannon Point (CNP)	DT	H310		Y		Y	Y	Y					
		H370	Y	Y		Y	Y	Y	Y	Y	Y		Y
		Z390	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
		Q370	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
		Q360		Y		Y	Y	Y	Y	Y			Y
		B360		Y		Y	Y	Y	Y	Y			Y
CNL/CFL PCH: CNP-H	MB-Halo	QM370	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
		HM370	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
		CM246	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
CNL/CFL PCH: CNP-H	HEDT	X399											
CNL/CFL PCH: CNP-H	WS	C246	Y	Y		Y	Y	Y	Y	Y	Y	Y	Y
CNL PCH: CNP-LP Intel® 8th Generation Core Processor Family Platform I/O	MB-LP	Premium-Y	Y	Y		Y	Y	Y	Y		Y		
		Premium-U	Y	Y		Y	Y	Y	Y	Y	Y		
		Base-U		Y			Y	Y			Y		
Coffee Lake-S PCH: KBP-H	DT	Z370	Y	Y		Y	Y	Y	Y	Y	Y	Y	
N - 1													
Basin Falls PCH: KBP-X	HEDT	X299	Y	Y	Y	Y	Y	Y	Y	Y	Y		
KBL Refresh (KBL-R) PCH: SPT-LP Intel® 8th Generation Core Processor Family Platform I/O	MB-LP	Base-U		Y			Y	Y					
		Premium-U	Y	Y	Y	Y	Y	Y	Y	Y	Y		
		Premium-Y	Y	Y	Y	Y	Y	Y	Y		Y		
Greenlow-Refresh	WS	C236	Y	Y	Y	Y	Y	Y	Y	Y	Y		
KBL-S PCH: KBP-H	DT	Q250 [SK]		Y		Y	Y	Y	Y	Y	Y		

		B250 SK		Y		Y	Y	Y	Y	Y	Y	
		H270 SK	Y	Y		Y	Y	Y	Y	Y	Y	
		Z270 SK	Y	Y	Y	Y	Y	Y	Y	Y	Y	
		Q270 SK	Y	Y	Y	Y	Y	Y	Y	Y	Y	
KBL-S **PCH: SPT-H**	MB-Halo	HM175 SK	Y	Y	Y	Y	Y	Y	Y	Y	Y	
		QM175 SK	Y	Y	Y	Y	Y	Y	Y	Y	Y	
		CM238 SK	Y	Y	Y	Y	Y	Y	Y	Y	Y	
KBL-LP **PCH: SPT-LP** Intel® 7th Generation Core Processor Family Platform I/O	MB-LP	Premium-U	Y	Y	Y	Y	Y	Y	Y	Y	Y	
		Base-U		Y			Y	Y				
		Premium-Y	Y	Y	Y	Y	Y	Y	Y		Y	

SK – *Supports both Kabylake CPU and Skylake CPU*

\# - *Supports SATA disks and only PCIe NVMe disks (no PCIe AHCI, and no RAID/SRT/Optane configurations)*

2.1.5 Supported Operating Systems for This Release

Table 2-4: Supported OS Matrix

16.x Release Supported	Operating System	Version	Supported Editions
16.0.x, 16.3	Windows* Redstone3	32@/64 bit	All editions with latest available RTM service release#
16.0.x, 16.3	Windows* Server 2016	64 bit	All editions with latest available RTM service release
16.0.x, 16.3	Windows PE~	32/64 bit	All editions

\# **Note**: that the 32 bit client OS versions will not be WHQL certified.

@ **Note**: Not WHQL certified

~ **Note:** OEM manufacturing support only

2.1.6 Supported MSFT* Performance Debug Tools

Beginning with Intel® Rapid Storage Technology 13.0, Intel® RST Driver supports the MSFT* "Crimson" Event Tracing for Windows*(ETW) performance measurements for responsiveness. Refer MSFT* for information on this ETW.

§ §

3 Intel® Rapid Storage Technology Suite

The Intel® Rapid Storage Technology Suite contains these core components:

1. Intel® Rapid Storage Technology (Intel® RST) OS runtime software package:
    a. AHCI/RAID driver (and filter driver for backwards compatibility)
    b. Graphical User Interface (Intel® RST UI) , optional
    c. Event Monitor service (IAStorDataMgrSvc) optional; interfaces with:
        i. Intel® RST UI (graphical user interface)
        ii. Event Notification Tray Icon (IAStorIcon)
        iii. Windows system NT Event log
    d. Intel® Optane™ Memory Service[1] (iastorAfsService)
    e. Intel® Optane™ Memory Minifilter[1] driver (iastorAfs)
    f. Intel® Optane™ Memory Native App[1] (iastorAfsNative)
2. Intel® Rapid Storage Technology BIOS components:
    a. Intel® Rapid Storage Technology RAID Option ROM (legacy support)
    b. UEFI driver (with HII-compliant UI)

Note[1]: Although installed in all cases, only runs when 32GB or larger SKUs of the Intel® Optane™ module are being used to enable System Acceleration.

The following components are available for OEM manufacturing use only; NOT to be distributed to end-users!

3. Intel® Rapid Storage Technology RAID mode utilities
    a. Intel® RSTCLI 32/64-bit Windows/WinPE command line interface utilities
    b. DEVSLP Tool - command line utility for configuring DEVSLP register values
    c. RcfgSata
        i. DOS-based command line interface utility (legacy support)
        ii. UEFI Shell-based command line interface utility
    d. RcmpSata compliance utility
        i. DOS-based Intel® RST RAID mode compliance check utility (legacy support)
        ii. UEFI Shell-based Intel® RST RAID mode compliance check utility

3.1 Intel® Rapid Storage Technology Software

The Intel® RST software is the major component of the Intel® Rapid Storage Technology Suite. The software includes the Intel® RST AHCI and RAID 32 and 64 bit drivers for supported Windows* operating systems. The driver supports several Intel® Serial ATA AHCI/RAID controllers and will recognize each unique device ID and sub-class code. Because of this, the driver must be installed before the Windows operating system is installed onto a RAID volume or a single SATA hard drive connected to the RAID controller. The OS runtime driver, in conjunction with the Intel® Rapid Storage Technology option ROM or the RST pre-OS UEFI driver, will provide boot capability for all supported RAID levels. The driver, in conjunction with the Intel® RST UI, provides RAID volume management (create, delete, migrate, etc.) within the Windows operating system. It also displays SATA* device and RAID volume information. Included with the software package is the RAID monitor service that monitors and reports various events of the storage subsystem. Other advanced features supported include System Acceleration with Intel® Optane technology.

3.2 Intel® Rapid Storage Technology Option ROM

The Intel® Rapid Storage Technology Option ROM is a standard Plug and Play option ROM that adds the Int13h services and provides a pre-OS user interface for the Intel® Rapid Storage Technology solution. The Int13h services allow a RAID volume to be used as a boot hard drive. They also detect any faults in the RAID volume being managed by the RAID controller. The Int13h services are active until the RAID driver takes over after the operating system is loaded.

The Intel® Rapid Storage Technology option ROM expects a BIOS Boot Specification (BBS) compliant BIOS. It exports multiple Plug and Play headers for each non-RAID hard drive or RAID volume, which allows the boot order to be selected from the system BIOS's setup utility. When the system BIOS detects the RAID controller, the *RAID option ROM* code should be executed.

The Intel Rapid Storage Technology option ROM is delivered as a single uncompressed binary image compiled for the 16-bit real mode environment. To conserve system flash space, the integrator may compress the image for inclusion into the BIOS. System memory is taken from conventional DOS memory and is not returned.

3.3 Intel Pre-OS RAID Configuration Utilities

The Intel RAID Configuration utility is an executable with capabilities similar to the Intel Rapid Storage Technology option ROM. Both Legacy OROM and UEFI configuration utilities can operate in 16-bit MS-DOS* mode. It provides customers with the ability to create, delete, and manage RAID volumes on a system within a DOS environment. For ease of use, the utility has command line parameters that make it possible to perform these functions by using DOS scripts or shell commands.

The RAID Configuration utilities use command line parameters. Below is a snapshot of the help text displayed when using the -? flag. It shows the usage for all supported command line flags necessary for creating, deleting, and managing RAID volumes.

3.3.1 RCfgSata Utility for MS-DOS* and UEFI

- Rcfgsata.exe = DOS application
- Rcfgsata.efi = UEFI application (UEFI shell required)

The command syntax for the Intel RAID Configuration utility is shown below:

==========

rcfgsata.efi (or rcfgsata.exe) [/?] [/Y] [/Q] [/C:vol_name] [/SS:strip_size] [/L:raid_level] [/S:vol_size] [/DS:disk_id] [/D:vol_name] [/X] [/I] [/P] [/U] [/ST] [/SP] [/V] [/RRT] [/Sync: <Auto | Manual>] [/M:disk_port] [/EM] [/ER] [/ACCEL:vol_name1 cache_vol mode] [/RA] [/SD] [CONCAT **D**evice_ID_fast-drive **D**evice_ID_slow-drive **C**oncatenation_boundary_in_MB] [DECONCAT]

/? Displays Help Screen. Other options ignored.

/Y Suppress any user input. Used with options /C, /D, /SP and /X.

/Q Quiet mode / No output. Should not be used with status commands.

COMMANDS - Only one command at a time unless otherwise specified.

/C	Create a volume with the specified name. /S, /DS, /SS, and /L can be specified along with /C.
/SS	Specify strip size in KB. Only valid with /C.
/L	Specify RAID Level (0, 1, 10, or 5). Only valid with /C.
/S	Specify volume size in GB or percentage if a '%' is appended. Percentage must be between 1-100. Only valid with /C.
/DS	Selects the disks to be used in the creation of volume. List should be delimited by spaces.
/D	Delete Volume with specified name.
/X	Remove all metadata from all disks. Use with /DS to delete metadata from selected disks.
/I	Display All Drive/Volume/Array Information. /P can be specified.
/P	Pause display between sections. Only valid with /I or /ST.
/U	Do not delete the partition table. Only valid with /C on RAID 1 volumes.
/SP	Marks the selected drive(s) as spare(s). Use with /DS
/ST	Display Volume/RAID/Disk Status.
/V	Display version information
/RRT	Create a recovery volume. Only valid with /C. Requires /M.
/Sync	Set sync type for 'Recovery' volume. Only valid with /RRT.
/M	Specify the port number of the Master disk for 'Recovery' volume. Only valid with /RRT.
/EM	Enable only master disk for recovery volume
/ER	Enable only recovery disk for recovery volume; /EM and /ER actions will result in change from Continuous Update mode to On-Request.
/ACCEL	Specify the volume to accelerate and acceleration mode vol_name1 - volume to accelerate cache_vol - the volume to use as cache mode - "enh" for enhanced, "max" - maximized
/RA	Removes the Disk/Volume Acceleration.
/SD	Synchronizes the data from the cache device to the Accelerated Disk/Volume.
/CONCAT	Concatenates drives into Optane Volume in postfix mode with specified migrations size in megabytes.

device_id_of_fast_drive - DeviceId of fast drive

device_id_of_slow_drive - DeviceId of slow drive

concatenation_boundary_in_MB - concatenation boundary in MB

/DECONCAT Disables the Optane™ Volume Acceleration.
==========

3.4 RSTCLI (32/64 bit) Windows Utilities

Note: : RSTCLI Commands are CasE SenSitiVe

The Intel RSTCLI 32/64 utility is an executable. It provides OEMs with the ability to create, delete, and manage RAID volumes on a system within a windows environment using command line parameters that make it possible to perform these functions by using scripts or shell commands. For use in all supported Windows OS including WinPE 32/64.

The command syntax for the Intel RSTCLI utilities is shown below:

USAGE: rstcli.exe (or rstcli64.exe)

Create Commands:

Flag	Name
-**C**	--create
-**E**	--create-from-existing
-l	--level
-n	--name
-s	--stripe-size
-z	--size
	--rrt
	--rrtMaster
	--rrtUpdate

Create Usage:
Creates a new volume and array or creates a new volume on an existing array.
--create --level x [--size y] [--stripe-size z] --name string [--create-from-existing diskId] diskId {[diskId]}
Create Examples:
-C -l 1 -n Volume 0-0-**1**-0 0-0-**2**-0 *(format of the disk ID is "0-0-**SATA_Port**-0" where the second digit from the left represents the SATA port on the platform where the disk is located; thus 0-0-**1**-0 represents SATA port # 1)*
--create -l 0 -z 5 --name RAID0Volume 0-0-3-0 0-0-4-0 0-0-5-0
-C -l 1 -E 0-0-1-0 -n VolumeWithData 0-0-2-0
-C --rrt -n RRTVolume 0-0-1-0 0-0-2-0 --rrtMaster 0-0-1-0
-C --rrt -n RRTVolume 0-1-0-0 0-2-0-0 --rrtUpdate Continuous
--create --help

Information Commands:

Flag	Name
-**I**	--information
-a	--array
-c	--controller
-d	--disk
-v	--volume

Information Usage:
Displays disk, volume, array, and controller information.
--information --controller|--array|--disk|--volume {[device]}
Information Examples:
-I -v Volume
-I -d 0-0-5-0
--information --array Array_0000
--information --help

Manage Commands:

Flag	Name
-**M**	--manage
-x	--cancel-verify
-D	--delete
-p	--verify-repair
-f	--normal-volume
-**F**	--normal
-i	--initialize
-**L**	--locate
-**T**	--delete-metadata
-Z	--delete-all-metadata**
-**N**	--not-spare
-**P**	--volume-cache-policy
-**R**	--rebuild
-**S**	--spare
-t	--target
-**U**	--verify
-w	--write-cache

****WARNING: Using this command deletes the metadata on ALL disks in the system. There is no option to select individual disks with this command and there is no warning prior to the command initiating and completing. To delete metadata on individual disks use the –D (--delete) command with either "volume_name" or "diskID".**

Manage Usage:
Manages arrays, volumes and disks present in the storage system.
--manage --cancel-verify volumeName
--manage --delete volumeName
--manage --verify-repair volumeName
--manage --normal-volume volumeName
--manage --normal diskId
--manage --initialize volumeName
--manage --locate diskId
--manage --delete-metadata diskId (deletes the metadata only on disks that are in a non-Normal state e.g. offline or unknown)
--manage --delete-all-metadata
--manage --not-spare diskId

--manage --volume-cache-policy off|wb --volume volumeName
--manage --rebuild volumeName --target diskId
--manage --spare diskId
--manage --verify volumeName
--manage --write-cache true|false --array arrayName

Manage Examples:
--manage --spare 0-0-3-0
-M -D VolumeDelete
-M --normal 0-0-2-0
--manage -w true -array Array_0000
-M -U VolumeVerify
-M -Z
--manage --help

Modify Commands:

Flag	Name
-m	--modify
-**A**	--Add
-**X**	--expand
-l	--level
-n	--name
-s	--stripe-size
-v	--volume

Modify Usage:
Modifies an existing volume or array.
--modify --volume VolumeName --add diskId {[diskId]}
--modify --volume VolumeName --expand
--modify --volume VolumeName --level L [--add diskId {[diskId]} [--stripe-size s] [--name N]
--modify --volume VolumeName --name n

Modify Examples:
-m -v Volume_0000 -A 0-0-3-0 0-0-4-0
-m --volume ModifyVolume --level 5
--modify -v Volume -n RenameVolume
--modify --help

Accelerate Commands (SRT):

Flag	Name
	--accelerate
	--createCache
	--setAccelConfig
	--disassociate
	--reset-to-available
	--accel-info
	--loadCache
	--stats

Accelerate Usage:

Accelerates a given disk or volume with the specified SSD disk.
--accelerate --createCache|--setAccelConfig|--disassociate|--reset-to-available|--accel-info
--accelerate --createCache --SSD <diskId> --cache-size **X** [where 16 ≤ **X** ≤ 64]
--accelerate --setAccelConfig --disk-to-accel <diskId> | --volume-to-accel <volume name> --mode [enhanced | maximized | off]
--accelerate --disassociate --cache-volume <volume name>
--accelerate --reset-to-available --cache-volume <volume name>
--accelerate --accel-info
--accelerate --loadCache <files or directory> --recurse
--accelerate --stats

Accelerate Examples:
--accelerate --createCache --SSD 0-0-3-0 --cache-size **X** [where 16 ≤ **X** ≤ 64]
--accelerate --setAccelConfig --disk-to-accel 0-0-5-0 --mode enhanced
--accelerate --setAccelConfig --volume-to-accel MyVolume --mode maximized
--accelerate --disassociate --cache-volume Cache_Volume
--accelerate --reset-to-available --cache-volume Cache_Volume
--accelerate --accel-info
--accelerate --loadCache C:\Windows*.* --recurse
--accelerate --stats
--accelerate --help

Intel® Optane™ Memory Commands:

Flag	Name
	--enable
	--disable
	--progress
	--info
	--clear-cache
	--file-cache-offset

Optane Memory Usage:
Options to manage Optane Memory drive:
--OptaneMemory --enable
--OptaneMemory --enable --fast-drive <driveId> --drive-to-accel <driveId>
--OptaneMemory --disable
--OptaneMemory --progress
--OptaneMemory --info <--verbose>
--OptaneMemory --clear-cache

Optane Memory Examples:
--OptaneMemory --enable
--OptaneMemory --enable --fast-drive 0-3-0-0 --drive-to-accel 0-0-5-0 –file-cache-offset <LBA>
--OptaneMemory --disable
--OptaneMemory --info --verbose
--OptaneMemory --help

OPTIONS:
--clear-cache
Clear Optane™ block cache; only recommended to use when running benchmarks.
--disable
Separate Optane™ Memory Volume into Optane device and slow capacity drive.
--enable
Create Optane™ Memory volume. No other parameters are required if the only 2 attached drives controlled by RST are the fast-drive and the drive-to-accel. The RSTCLI/driver will automatically select those two drives to enable Optane™.

-h, --help
Displays help documentation for command line utility modes, options, usage, examples, and return codes. When used with a mode switch (create, information, mange, modify, or accelerate), instructions for that mode display. For example, --create --help displays Create option help.

--file-cache-offset
This option is important and should be used with the --enable command to specify the start point on the slow drive, from which data will be transferred to fast drive. In Practice, --file-cache-offset must be equal to the beginning of Windows (C:\) partition on the slow drive, specified in LBA. If --file-cache-offset parameter is set to 0, or not defined properly, during the next boot (e.g. OOBE) the Optane™ Memory volume rebuilds. The rebuilds process consist of disable Optane™ Memory volume, and enable Optane™ Memory with proper --file-cache-offset value (set automatically to the beginning of C: partition).

--info
Lists information about Optane Memory settings.

--progress
Displays the progress of file cache migration.

CLI OPTIONS:

-A <<host>-<bus>-<target>-<lun>>, --add <<host>-<bus>-<target>-<lun>>
Adds new disks to an existing volume.

-a, --array
Lists information about the arrays in the storage system.

--accel-info
Lists information about Accelerate settings.

--accelerate
Accelerates a given disk or volume with the specified SSD disk.

-C, --create
Creates a new volume and array or creates a new volume on an existing array.

-c, --controller
Lists information about the controllers in the storage system.

--cache-size <MIN or MAX>
Sets a size in gigabytes for the cache memory. This is an optional switch. If the size is not specified, the complete size of the SSD will be used for acceleration.

--cache-volume <Volume name>
Specifies a name for the volume used as cache.

--clear-cache
Clear Optane block cache.

--createCache
Creates the cache.

-D <Volume name>, --delete <Volume name>
Deletes the specified volume.

-d, --disk
Lists information about the disks in the storage system.

--disable
Separate Optane Memory Volume into Optane device and capacity drive.

--disassociate
Disassociates the Cache volume from acceleration

--disk-to-accel <<host>-<bus>-<target>-<lun>>

Specifies a disk if accelerating a pass-through disk.

--drive-to-accel <<host>-<bus>-<target>-<lun>>

Specifies a disk if enabling an Optane volume on a pass-through disk.

--dynamic-storage-accelerator <true or false>

Enables/disables dynamic storage accelerator; using 'true' enables, 'false' disables.

-E <<host>-<bus>-<target>-<lun>>, --create-from-existing <<host>-<bus>-<target>-<lun>>

Identifies the disk if data is to be migrated from one of the disks. Disk identifier is SCSI address.

--enable

Create Optane Memory Volume

-F <<host>-<bus>-<target>-<lun>>, --normal <<host>-<bus>-<target>-<lun>>

Resets failed or SMART event disk to normal.

-f <Volume name>, --normal-volume <Volume name>

Resets failed RAID 0 volume to normal and recovers data.

--fast-drive <<host>-<bus>-<target>-<lun>>

Specifies the location of the Optane device that will be used to enable Optane Memory volume

--file-cache-offset

This option is important and should be used with the --enable command to specify the start point on the slow drive, from which data will be transferred to fast drive. In Practice, --file-cache-offset must be equal to the beginning of Windows (C:\) partition on the slow drive, specified in LBA. If --file-cache-offset parameter is set to 0, or not defined properly, during the next boot (e.g. OOBE) the Optane™ Memory volume rebuilds. The rebuilds process consist of disable Optane™ Memory volume, and enable Optane™ Memory with proper --file-cache-offset value (set automatically to the beginning of C: partition).

-h, --help

Displays help documentation for command line utility modes, options, usage, examples, and return codes. When used with a mode switch (create, information, mange, modify, or accelerate), instructions for that mode display. For example, --create --help displays Create option help.

-I, --information

Displays disk, volume, array, and controller information.

-i <Volume name>, --initialize <Volume name>

Initializes the redundant data on a RAID 1, 5 or 10 volume.

--info

Lists information about Optane Memory settings.

-L <<host>-<bus>-<target>-<lun>>, --locate <<host>-<bus>-<target>-<lun>>

Locates device and blinks the LED.

-l <0, 1, 5, 10>, --level <0, 1, 5, 10>

Changes the RAID type of an existing volume. Options are migrations from RAID 1 to RAID 0 or 5, RAID 0 to RAID 5, and RAID 10 to RAID 5.

--loadCache C:\Windows*.* --recurse

Used to pre-load files into the cache

-M, --manage

Manages arrays, volumes and disks present in the storage system.

-m, --modify

Modifies an existing volume or array.

--mode <Enhanced or Maximized mode>

Specifies Accelerate mode as Enhanced or Maximized.

-N <<host>-<bus>-<target>-<lun>>, --not-spare <<host>-<bus>-<target>-<lun>>

Resets a spare disk to available.

-n <Volume name>, --name <Volume name>

Specifies a name for the volume created. Renames an existing volume in Modify mode.

-P <Volume name>, --volume-cache-policy <Volume name>

Sets volume cache policy to either off or wb.

-p <Volume name>, --verify-repair <Volume name>

Verifies and repairs the volume.

--progress

Displays the progress of file cache migration.

-q, --quiet
Suppresses output for create, modify, and manage modes. Not valid on info mode.

-R <Volume name>, --rebuild <Volume name>
Rebuilds the degraded volume.

-r, --rescan
Forces the system to rescan for hardware changes.

--reset-to-available
Resets the cache volume to available.

--rrt
Creates a recovery volume using Intel(R) Rapid Recovery Technology (RRT).

--rrtMaster <<host>-<bus>-<target>-<lun>>
Optionally creates a recovery volume that allows you to select a specific disk as the master disk. Default is the first disk in the disk list.

--rrtUpdate <Continuous or OnRequest Update>
Specifies a data update setting when creating a recovery volume as Continuous or OnRequest. Default is Continuous.

-S <<host>-<bus>-<target>-<lun>>, --spare <<host>-<bus>-<target>-<lun>>
Marks a disk as a spare.

--SSD <<host>-<bus>-<target>-<lun>>
Specifies SSD disk that will be used as cache. If another SSD is being used as cache, then that volume needs to be deleted to use a new SSD disk.

-s <size in KB>, --stripe-size <size in KB>
Sets a stripe size in kilobytes (2^10 bytes) for a volume. Valid when creating or changing the type of a volume and for RAID 0, RAID 5 and RAID 10. Options are 4, 8, 16, 32, 64 and 128 KB.

--setAccelConfig
Sets the config for accelerating a volume or disk.

--stats
Indicates percentage of cache usage.

-T <<host>-<bus>-<target>-<lun>>, --delete-metadata <<host>-<bus>- <target>-<lun>>
Deletes the metadata from the specified disk. (deletes the metadata only on disks that are in a non-Normal state e.g. offline or unknown)

-t <<host>-<bus>-<target>-<lun>>, --target <<host>-<bus>-<target>-<lun>>
Indicates the pass-through disk to be used for rebuilding a degraded volume.

-U <Volume name>, --verify <Volume name>
Verifies data on the volume.

-u <password>, --unlock <password>
Unlocks a disk.

-V, --version
Displays version information.

-v, --volume
Lists information about the volumes on the system. Stipulates the volume to act on when used in Modify or Manage mode.

--volume-to-accel <Volume name>
Specifies a name of the volume to be accelerated.

-w <true or false>, --write-cache <true or false>
Enables or disables write cache for all disks that are part of an array.

-X, --expand
Expands a volume to consume all available space in an array.

-x <Volume name>, --cancel-verify <Volume name>
Cancels a verify operation in progress.

-z <size in GB>, --size <size in GB>
Sets a size in gigabytes. This is an optional switch. If the size is not specified or specified to 0, then the maximum size available will be used.

-Z --delete-all-metadata
Deletes the metadata on **all disks** in the system without any warning prior to initiating and completing the action.

RETURN CODES:

0, Success
Request completed successfully.
1, Request Failed
Request is formatted correctly but failed to execute.
2, Invalid Request
Unrecognized command, request was formatted incorrectly.
3, Invalid Device
Request not formatted correctly, device passed in does not exist.
4, Request Unsupported
Request is not supported with the current configuration.
5, Device State Invalid
Request is not supported with the current device state.
20, Invalid Stripe Size
Stripe size is not supported.
21, Invalid Name
Volume name is too long, has invalid characters, or already exists.
Volume name cannot exceed 16 English characters.
22, Invalid Size
Size requested is invalid.
23, Invalid Number Disks
Number of disks requested is invalid.
24, Invalid RAID Level
RAID level requested is invalid.

3.5 UEFI System BIOS and Intel® RST UEFI/RAID Package

Beginning with the Intel® RST 11.5 Release version, the product provides a native UEFI driver for OEMs and their BIOS vendors to integrate into their RAID-enabled platforms (Not required for AHCI mode platforms).

3.5.1 Specification References

This document is not intended to be a go-to document for the UEFI specification. The specification is owned by the UEFI working group and detailed information regarding UEFI can be found in documents published by that organization. The Intel® RST UEFI driver implementation conforms to the UEFI specification and is in compliance with version 2.3.1.

Table 3-1: UEFI Specifications and Location

Specification	Location
UEFI Specification version 2.3.1	(http://www.uefi.org/specsandtesttools)
UEFI Platform Initialization Specification version 1.2	(http://www.uefi.org/specsandtesttools)

3.5.2 What Intel® RST Provides to OEMs/BIOS Vendors

3.5.2.1 Intel® Rapid Storage Technology UEFI Driver

This is the main component of the Intel® RST pre-OS EFI solution. It is provided in three different formats:

RaidDriver.efi (filename):

- UEFI driver that requires integration into the UEFI System BIOS by the OEM's BIOS vendor. This file can be placed into the OEMs' UEFI BIOS source build where their tools can integrate it.

RaidDriver.ffs (filename):

- The Intel® RST UEFI driver (RaidDriver.efi) is wrapped in the Firmware File System (.ffs)
- Useful for an external tool to integrate the binary into a compiled BIOS image. Firmware File System Details:
  - Firmware File Type - EFI_FV_FILETYPE_DRIVER (0x07)
  - File GUID - 90C8D394-4E04-439C-BA55-2D8CFCB414ED
  - 2 Firmware File Sections
    - EFI_SECTION_PE32 (0x10)
    - EFI_SECTION_USER_INTERFACE (0x15)
      Name "SataDriver"

RaidDriver.bin (filename):

- This is an optional format that is provided to OEMs that might want it delivered as a PCI 3.0 UEFI OROM
- Disadvantage of the UEFI OROM format is that it likely will require the BIOS to have a Compatibility Support Module (CSM) in order to function

3.5.2.2 Intel® RST UEFI User Interface

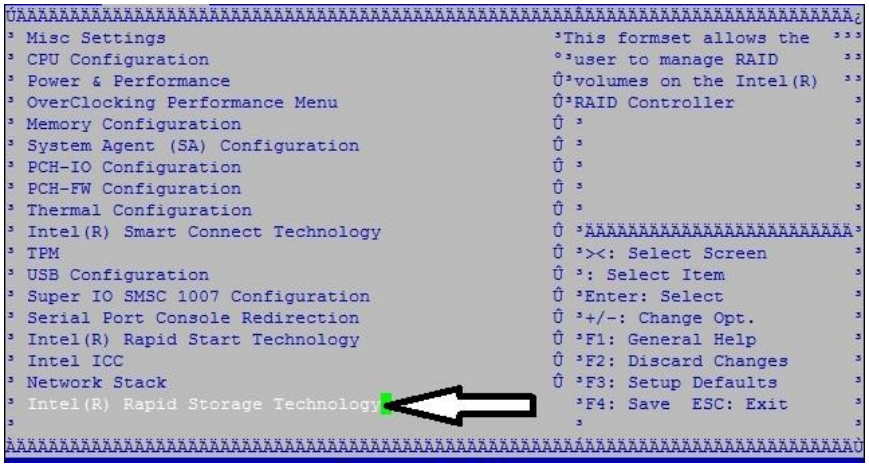
An **HII-compliant** user interface is provided for the pre-boot configuration of the RAID system, including management of the premium cacheing features Intel® Smart Response Technology and System Acceleration with Intel® Optane™ Technology. The same functionality provided in the legacy OROM UI is available in the HII UI.

- The UI is integrated within the UEFI driver binary (RaidDriver.efi, .ffs, and .bin files)
- Per the UEFI specification, we publish the UI as string and forms packages
- The UI is accessible from within the UEFI BIOS (How the user accesses it from within the BIOS is OEM-dependent upon implementation)
- The text string 'Intel(R) Rapid Storage Technology' will be displayed as the selection to enter the UI
- Some OEMs may want to hard assign where the Intel® RST UEFI GUI will be located within their BIOS.
- **The Intel® RST UEFI Driver FORMSET_GUID is:**

```
FORMSET_GUID { 0xd37bcd57, 0xaba1, 0x44e6, { 0xa9, 0x2c, 0x89, 0x8b, 0x15, 0x8f,
0x2f, 0x59 } }
```

{D37BCD57-ABA1-44e6-A92C-898B158F2F59}

Figure 3-1: Example location of RST UEFI UI in a System BIOS

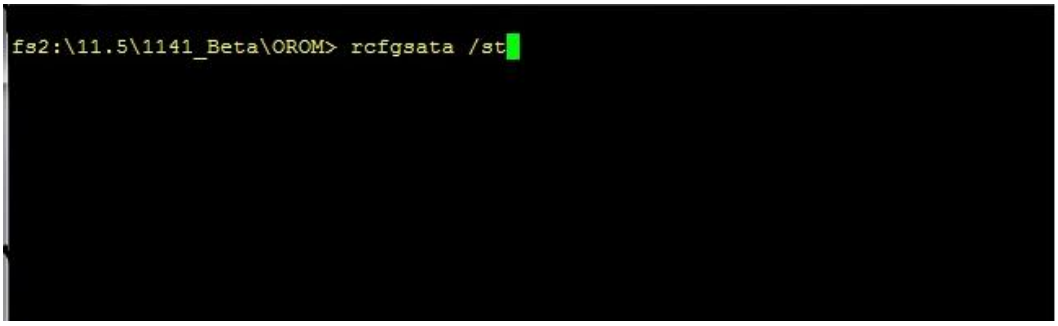


3.5.2.3 Command Line RAID Configuration Utility

RcfgSata.efi (filename):

- A UEFI application that requires booting to the UEFI Shell environment to run
- Same functionality and commands as have been provided by the legacy DOS version (***RcfgSata.exe***) in previous releases of the Intel® RST product.
- Requires the exact same version of the RST UEFI Driver to be loaded on system in order to function.

Figure 3-2: rcfgsata command line syntax



3.5.2.4 Command Line RAID Compliance Checking Utility

RcmpSata.efi (filename):

- A UEFI application that requires booting to the UEFI Shell environment to run.

- Investigates if the list of UEFI required protocols by the RST UEFI Driver are present. Also provides a list of the protocols published by the RST UEFI Driver and the capabilities/features of the RST UEFI Driver.

Figure 3-3: rcfgsata command line syntax



3.5.3 UEFI System BIOS Requirements for Platform Compatibility with Intel® RST UEFI

This section covers what the OEM/BIOS Vendor is required to accomplish in order to ensure that the platform is compatible the Intel® RST UEFI driver.

3.5.3.1 Required Protocols/Functions to be Provided by UEFI System BIOS

The Intel® RST UEFI driver requires the following protocols/functions to be provided by the BIOS:

EFI_BOOT_SERVICES:

- LocateHandleBuffer
- OpenProtocol
- CloseProtocol
- WaitForEvent
- HandleProtocol
- FreePool
- AllocatePages
- AllocatePool
- InstallMultipleProtocolInterfaces
- UninstallMultipleProtocolInterfaces
- Stall

EFI_RUNTIME_SERVICES:

- SetVariable
- GetVariable
- GetTime

Other Protocols:

- EFI_ACPI_TABLE_PROTOCOL (or EFI_ACPI_SUPPORT_PROTOCOL (EDK117))

3.5.3.2 Optional Protocols/Functions to be Provided by UEFI System BIOS

If the OEM plans to use the Intel® RST HII-compliant UI, then **the following protocols/functions are required to be provided by the BIOS**:

- Form Browser 2 Protocol
- Config Routing Protocol
- HII String Protocol
- HII Database Protocol

3.5.3.3 Protocols Provided by Intel® RST UEFI Driver

The Intel® RST UEFI driver provides the following protocols:

- *Driver Binding Protocol*
- *Component Name Protocol (English only)*
- *Component Name 2 Protocol (English only)*
- *Driver Supported EFI Version Protocol*
- *Device Path Protocol*
- *Config Access Protocol*
- EFI_BLOCK_IO_PROTOCOL
    - For Logical Devices
- EFI_STORAGE_SECURITY_PROTOCOL
    - For Non-RAID disks that support TCG Feature Set
- EFI_EXT_SCSI_PASS_THRU_PROTOCOL:
    - All SCSI commands are supported (for ATAPI devices)
- EFI_ATA_PASS_THRU_PROTOCOL:
    - Non-RAID disks:
        - All ATA commands are supported
    - RAID disks (only the following commands are supported):
        - EXECUTE DEVICE DIAGNOSTIC (0x90)
        - IDENTIFY DEVICE (0xEC)
        - IDLE (0xE3)
        - IDLE IMMEDIATE (0xE1)
        - SECURITY DISABLE PASSWORD (0xF6)
        - SECURITY ERASE PREPARE (0xF3)
        - SECURITY ERASE UNIT (0xF4)
        - SECURITY FREEZE (0xF5)
        - SECURITY SET PASSWORD (0xF1)
        - SECURITY UNLOCK (0xF2)
        - SET FEATURES (0xEF)

        - SMART READ DATA (0xB0 / 0xD0)
        - SMART READ LOG (0xB0 / 0xD5)
        - SMART RETURN STATUS (0xB0 / 0xDA)
        - STANDBY (0xE2)
        - STANDBY IMMEDIATE (0xE0)
    - All disk types:
        - EFI_ATA_PASS_THRU_PROTOCOL_ATA_NON_DATA
        - EFI_ATA_PASS_THRU_PROTOCOL_PIO_DATA_IN
        - EFI_ATA_PASS_THRU_PROTOCOL_PIO_DATA_OUT
        - EFI_ATA_PASS_THRU_PROTOCOL_DEVICE_DIAGNOSTIC
        - EFI_ATA_PASS_THRU_PROTOCOL_UDMA_DATA_IN
        - EFI_ATA_PASS_THRU_PROTOCOL_UDMA_DATA_OUT
        - EFI_ATA_PASS_THRU_PROTOCOL_RETURN_RESPONSE

3.5.4 How-to-Enable the Platform with Intel® RST UEFI Driver/HII_GUI

This section covers what the OEM/BIOS Vendor is required to accomplish in order to ensure that the platform is compatible with the Intel® RST UEFI driver.

3.5.4.1 Step1: Platform UEFI BIOS

1. Ensure that the UEFI System BIOS meets UEFI Specification 2.3.1 compliance
2. The BIOS must provide the following protocols:
    - EFI_Boot_Services Protocols (see section 3.5.3.1)
    - EFI_Runtime_Services Protocols (see section 3.5.3.1)
    - EFI_HII Protocols** (see section 3.5.3.2) *** Required for the Intel® RST UEFI UI

3.5.4.2 Step2: Download and Integrate Intel® RST UEFI Package

1. Download the latest kit from the Intel VIP (Validation Internet Portal) website. From the kit select the ***efi_sata.zip*** file which will contain the UEFI driver binary files (*RaidDriver.efi, RaidDriver.ffs, and RaidDriver.bin*)
2. Select and extract the binary file based on the planned integration method:
    - *RaidDriver**.efi***: Use this binary if planning to integrate at the time of the BIOS image build
    - *RaidDriver**.ffs***: Use this binary if planning to integrate into an already built BIOS image
    - *RaidDriver**.bin***: Use this binary if planning to integrate as legacy type OROM (CSM may also be required)
3. Use the proper integration tools based on the binary file selected above

3.5.4.3 Step3: Verify Compliance

1. From the ***efi_sata.zip*** downloaded in step 2, extract the ***RCmpSata.efi*** file.

2. Place the file on a USB thumb drive and insert the drive into the platform
3. Boot to the UEFI Shell environment.
4. Run the RCmpSata.efi application (it's a command line utility): at the prompt type the command:
   - To print to screen: ***rcmpsata.efi***
   - To print to a file: ***rcmpsata.efi > comply.txt***
5. Ensure all sections pass with no fails reported

3.5.5 Known Compatibility Issues with UEFI Self Certification Test (UEFI SCT) Tool

The following UEFI 2.3.1 SCT tests will appear as FAIL in reports generated using the "Report Generation" tool of the SCT framework. The "Report Generation" tool is the only method that should be used to determine if tests fail. Do not determine test failing test results by viewing the raw log files. The "Report Generation" tool will discard any test results that failed due to an invalid system configuration.

3.5.5.1 Bootable Image Support Test\Block IO Protocol Test

EFI_BLOCK_IO_PROTOCOL.Reset - Reset() returns EFI_SUCCESS with ExtendedVerification being TRUE

- *Test Index*: 5.7.5.1.1
- *Test GUID*: 61EE3A34-62A2-4214-B076-5073B177156C
- *Reason*: The Intel® RST UEFI driver does not support Reset – EFI_UNSUPPORTED is returned.

EFI_BLOCK_IO_PROTOCOL.Reset - Reset() returns EFI_SUCCESS with ExtendedVerification being FALSE

- *Test Index*: 5.7.5.1.2
- *Test GUID*: 98530F3D-8BD8-44A1-9D06-08039FDFEC63
- *Reason*: The Intel® RST UEFI driver does not support Reset – EFI_UNSUPPORTED is returned.

EFI_BLOCK_IO_PROTOCOL.ReadBlocks - ReadBlocks() returns EFI_SUCCESS with valid parameter

- *Test Index*: 5.7.5.2.1
- *Test GUID*: 9EFE26C2-C565-478A-A0B4-05A8FD2E7E3E
- *Reason*: Test called ReadBlocks() with a BufferSize of 0 so EFI_BAD_BUFFER_SIZE is returned. The UEFI 2.3.1 specification states for ReadBlocks, "The size of the Buffer in bytes. This must be a multiple of the intrinsic block size of the device."

3.5.5.2 ATA Bus Support Test\ATA Pass-Thru Protocol Test

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath - call BuildDevicePath with NULL DevicePath.

- *Test Index*: 5.7.8.2.1
- *Test GUID*: D72E6A78-5292-4493-9040-B0445A9C1714
- *Reason*: The Intel® RST UEFI driver does not support BuildDevicePath – EFI_UNSUPPORTED is returned

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath - call BuildDevicePath with invalid Port.
- *Test Index*: 5.7.8.2.2
- *Test GUID*: A42A0E01-7B80-46E4-A757-86C4EC53F4E4
- *Reason*: The Intel® RST UEFI driver does not support BuildDevicePath – EFI_UNSUPPORTED is returned

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath - call BuildDevicePath with invalid PortMultiplierPort
- *Test Index*: 5.7.8.2.3
- *Test GUID*: 322F00C1-F6BF-41ED-AEFD-AAC48F3FA9DB
- *Reason*: The Intel® RST UEFI driver does not support BuildDevicePath – EFI_UNSUPPORTED is returned

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath BuildDevicePath() with available device, device path
- *Test Index*: 5.7.8.2.4
- *Test GUID*: 230D44B6-CE53-42B6-9BA6-3D115D492B33 should be created.
- *Reason*: The Intel® RST UEFI driver does not support BuildDevicePath – EFI_UNSUPPORTED is returned

EFI_ATA_PASS_THRU_PROTOCOL.GetDevice GetDevice() with NULL device path.
- *Test Index*: 5.7.8.3.1
- *Test GUID*: 0F2F0849-690B-48EA-8E35-64363FAA8C5C
- *Reason*: The Intel® RST UEFI driver does not support GetDevice – EFI_UNSUPPORTED is returned

3.5.5.3 HII Test\HII Config Access Protocol Test

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig - RouteConfig() returns EFI_NOT_FOUND if no target was
- *Test Index*: 5.18.6.2.3
- *Test GUID*: 1F99EBC8-0253-455F-88AC-9E2BA6DCD729 found with the routing data.
- *Reason*: Intel® RST UEFI driver does not support RouteConfig – EFI_UNSUPPORTED is returned. RouteConfig is not supported so that Intel® RST HII form values are only modified by the Intel® RST driver itself.

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig - RouteConfig() returns EFI_INVALID_PARAMETER with Configuration been NULL
- *Test Index*: 5.18.6.2.1
- *Test GUID*: 495C99F3-0231-45A5-AFFA-D25C6F9A191C
- *Reason*: is caused by not using EFI HII Configuration Access Protocol (see section 31.4 in UEFI 2.4 specification) in RST UEFI Driver. The RST software does not use it as it is not necessary.

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig- RouteConfig() returns EFI_SUCCESS with valid parameters
- *Test Index*: 5.18.6.2.4
- *Test GUID*: 1A15DF85-6CC1-43F2-9B86-218BD5FDF4A0

- *Reason*: is caused by not using EFI HII Configuration Access Protocol (see section 31.4 in UEFI 2.4 specification) in RST UEFI Driver. The RST software does not use it as it is not necessary.

3.6 Pre-OS

3.6.1 BIOS PCH SATA Controller Modes

The following table explains the PCH SATA controller modes available on all Kaby Lake and later platforms:

Pre-Kaby Lake BIOS code SATA modes (Old)	Kaby Lake BIOS code SATA modes (New)	Supported Platforms	Major Feature Support
AHCI	AHCI (no change)	All	AHCI basic features
RAID	Intel® RST Premium and System Acceleration with Intel® Optane Technology	Premium Intel® PCH Chipsets (e.g. Q270, Z270, Z370)	• RAID 0/1/5/10 • Intel® Optane™ Memory • Intel® SRT** • PCIe NVMe by remapping
None/New Mode	Intel® RST and System Acceleration with Intel® Optane Technology	Non-Premium Intel® PCH Chipsets (e.g. B250, Q360)	• Intel® Optane™ Memory • PCIe NVMe by remapping

Discontinued beginning on Intel® 300/240 Series Chipset Family

§ §

4 New in Release Version 16.x

4.1 Feature Specification Update in Release 16.3

4.1.1 Intel® Optane™ Memory User Pinning

User pinning functionality has been added to the Intel® Optane™ Memory feature. Pinning utilizes the file cache portion of the Intel® Optane™ memory module. Pinning allows the user to select files, folders, and applications that they wish to pin to the file cache which is the highest performance level of the Intel® Optane™ Memory caching solution. All items pinned by the user are considered to be "hard" pinned which means they can only be removed/unpinned from the file cache by the user. The Intel® Optane™ Memory caching policies cannot evict 'hard pinned' items.

4.2 Feature Specification Update in Release 16.0.2

4.2.1 Acceleration of Data Drives with Intel® Optane™ Memory

The Intel® Optane™ Memory feature is enhanced to include support for the acceleration of non-boot system drives (data drives). Procedures for accelerating a data drive are the same as for OS drive acceleration. The major difference between data drive and OS drive acceleration is that data drive acceleration supports only block cache. There is no file cache support for data drive acceleration.

4.3 New Features/Specifications in Release 16.0.1

4.3.1 CPU Attached Storage- PCIe NVMe

CPU Attached Storage uses SW remapping by the Intel® RST SW to control Intel-branded PCIe NVMe drives that are attached directly to the PEG ports off the computer's CPU. These Intel NVMe devices/drives can be used to create RAID 0 or RAID 1 volumes. They also can be used as the fast drive (cache device) for System Acceleration with Intel® Optane™ Memory.

4.4 What's New

4.4.1 Intel® Optane™ Memory UI/Installer

4.4.1.1 New Functionality in the Intel® Optane™ Memory UI

The Intel® Optane™ Memory UI has added user pinning of applications, files and directories when the Optane™ memory module is 32GB or larger (user pinning is not supported on 16GB Optane™ memory modules).

NOTE! This feature is not supported in the Intel® RST UI (UI with RAID management support)

Figure 4-1: Intel® Optane™ Memory UI Pinning Main Page

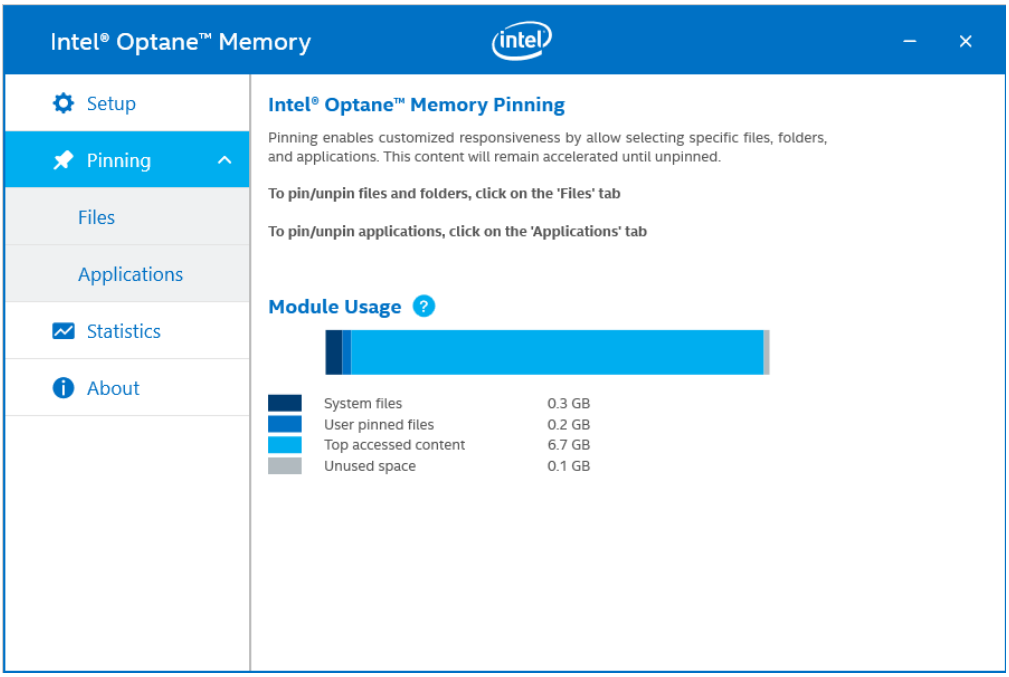


Figure 4-2: Intel® Optane™ Memory UI Currently Pinned Files and Folders Page

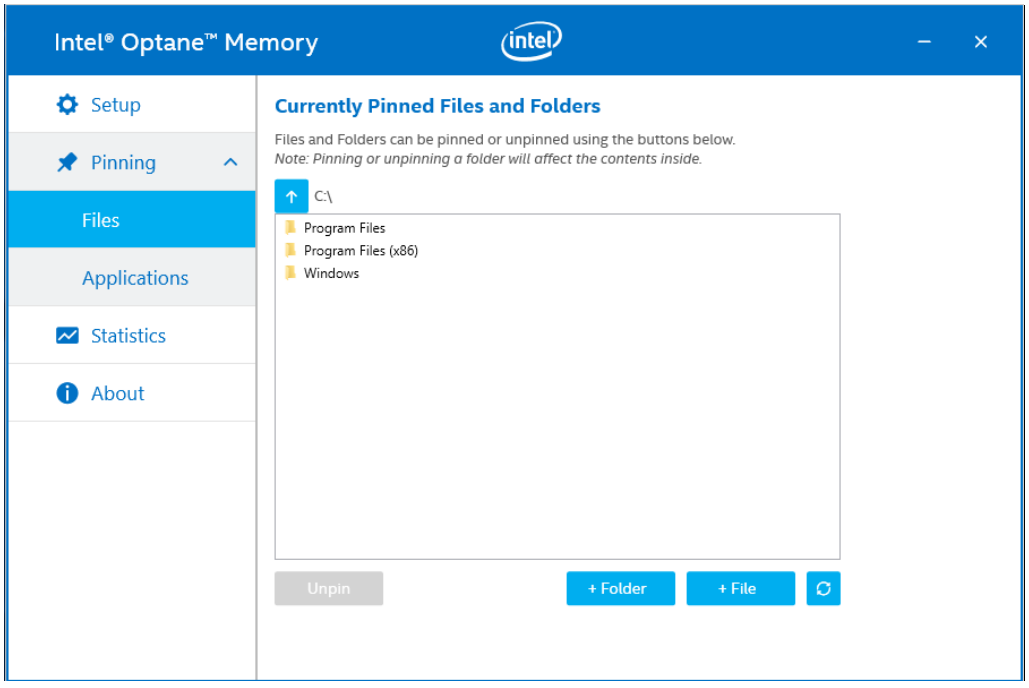
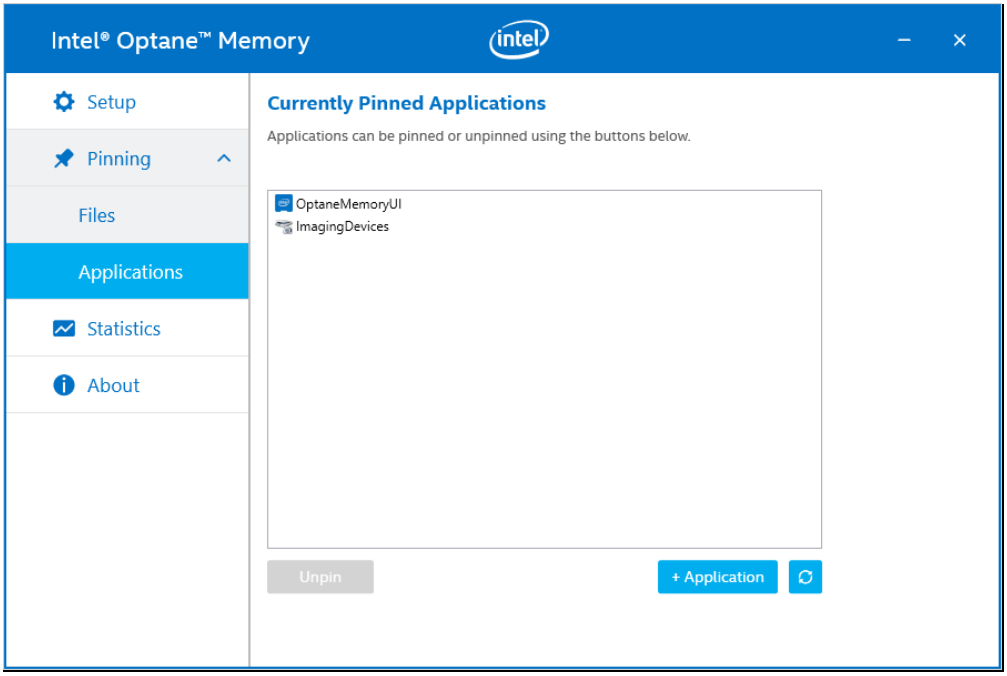


Figure 4-3: Intel® Optane™ Memory Currently Pinned Applications Page



§ §

5 CPU Attached Storage: PCIe NVMe

5.1 Overview

CPU Attached Storage feature eliminates existing platform HW restrictions: no Cycle Routers are connected to the CPU PCIe port. Using SW remapping technology, Intel-branded PCIe NVMe storage devices can be attached to the PCI Express Graphics port (PEG-port) with full support for Intel® RST premium features such as RAID and Intel® Optane™ Memory.

5.2 Limitations and Minimum Requirements:

5.2.1 HW

- Intel-branded PCIe NVMe devices only (3rd party devices not supported)
- A maximum of 4 Intel-branded PCIe NVMe devices
    - Requires 4 x2 or x4 PEG ports configured with either M.2 connectors or PCIe slots

5.2.2 SW/FW

- Intel® RST SW release version 16.0PV or later
- Intel® RST pre-OS version 16.0PV or later

5.3 Usage models

CPU attached Intel® Optane™ memory modules and SSDs can be used in the following configurations.

Configuration	Details
Pass-through data storage disk	Intel-branded PCIe NVME SSDs installed into a properly configured PEG port (CPU Attached) can be used as a pass-through data storage disk.

Bootable storage disk	Intel-branded PCIe NVME SSDs installed into a properly configured PEG port (CPU Attached) can be used as a bootable OS disk.
RAID 0/1 data volume	Two or more Intel-branded PCIe NVME SSDs installed into properly configured PEG ports (CPU Attached) can be configured into a RAID 0 or RAID volume and used as a data storage disk.
RAID 0/1 bootable volume	Two or more Intel-branded PCIe NVME SSDs installed into properly configured PEG ports (CPU Attached) can be configured into a RAID 0 or RAID 1 volume and used as a bootable OS disk.
RAID 0/1/5/10 data volume	At least one Intel-branded PCIe NVME SSD installed into a properly configured PEG port (CPU Attached) can be combined with 1 or more HW remapped (PCH attached) PCIe NVMe SSDs and configured into a RAID 0, 1, 5, or 10 volume. The RAID volume can then be used as a data storage disk. **Note:** CPU Attached (CPUa) PCIe NVMe SSDs can only be used in RAID 5 and RAID 10 volumes when combined with at least one PCH attached (PCHa) remapped PCIe NVMe SSD. Not supported: (CPUa + CPUa + CPUa) RAID 5 Supported: (CPUa + CPUa + PCHa) RAID 5
RAID 0/1/5/10 bootable volume	Two or more Intel-branded PCIe NVME SSDs installed into properly configured PEG ports (CPU Attached) can be combined with 1 or more HW remapped (PCH attached) PCIe NVMe SSDs and configured into a RAID 0, 1, 5, or 10 volume. The RAID volume can then be used as a bootable OS disk (same RAID 5 and RAID 10 restrictions as above).
Intel® Optane™ Memory fast media disk	An Intel-branded PCIe NVME SSD installed into a properly configured PEG port (CPU Attached) can be used as a fast media disk to enable System Acceleration with Intel® Optane™ memory. **Note:** The slow media (disk to be accelerated) must be a single pass through SATA disk (SSD, SSHD, or HDD) with the active OS GPT partition.

6 System Acceleration with Intel® Optane™ Memory

6.1 Limitations:

6.1.1 HW limits

1. **Fast disk** (cache device)
   1. Only a single HW or SW remapped Intel® Optane™ memory module
2. **Slow[1] disk** (disk to be accelerated must be larger capacity than the fast disk) only in the following configurations and drive interfaces:
   1. ***Boot Drive*** (Active boot volume):
      - SATA HDD – single pass-through[2]
      - SATA SSD - single pass-through
      - SSHD[3] - single pass-through
   2. ***Data Drive*** (No active boot volume):
      - SATA HDD – single pass-through[2]
      - SATA SSD - single pass-through
      - SSHD[3] - single pass-through
3. Optane™ Memory volume
   1. Only one Optane™ Memory volume allowed per boot instance. If more than one volume detected during boot, the second volume detected will be placed offline and not accessible to the system.

NOTES:

Note[1]: The slow disk (targeted disk to be accelerated) can be data drives as well as boot drives with the active boot volume beginning with release version 16.0.2.

Note[2]: Only single pass-through disks can be accelerated. The acceleration of RAID volumes or other multiple disk configurations is not supported.

Note[3]: Self-pinning SSHDs only, SSHDs that use the Hybrid Information Feature Set are not supported.

6.1.2 Functional limits

6.1.2.1 SATA controller mode switching to AHCI mode

Once the system is setup with Intel® Optane™ Memory enabled, the user cannot switch the PCH SATA controller mode from one of the Intel RST modes to the AHCI mode. If the mode is switched and the user attempts to boot the system, the Intel® RST configuration metadata can be corrupted making the system unbootable when PCH SATA mode is switched backed to the available Intel RST mode. To avoid this error condition the following precautions should be taken:

1. All users: If you require switching the PCH SATA controller mode to AHCI for some reason, before switching, disable Intel® Optane™ Memory prior to switching the mode. You can re-enable Intel® Optane™ Memory.
2. BIOS vendors: Upgrading BIOS should maintain pre-upgrade BIOS settings (specifically, the PCH SATA mode should remain in the 'Intel RST…' mode and not switch to AHCI mode)

6.1.2.2 Disk defragmentation of an Optane™ volume

The 'Disk Defragmentation' feature in Windows is not functional on Intel® Optane™ volumes.

6.2 Minimum Requirements:

6.2.1 HW

1. CPU: Intel® Kaby Lake or later CPU family of platforms
2. PCH: Intel® 100/C230 Series, 200 Series, C240 Series, or 300 Series family of chipsets
3. Motherboard supporting Intel® RST PCH HW remapping technology or a x2 or x4 PEG port
4. Intel® Optane™ technology NVMe SSD/memory module or SSD
    - An M.2 connector is required for the Optane™ memory module; it can be down on the motherboard or on a PCIe adapter card plugged into a PCH remapped PCIe x2 or x4 slot or CPU PEG port.

6.2.2 SW/FW

- UEFI-compliant system BIOS based on Intel® Kaby Lake BIOS reference code version 1.5.0 or later
- Intel® RST pre-OS UEFI driver version **15.5.0.2875 (16.3 PV for the latest features)** release or later integrated into system BIOS
- BIOS Intel PCH SATA controller mode set to "Intel® RST…" with remapping enabled on the HSIO lanes connecting the Intel® Optane™ technology NVMe SSD/memory module
- Intel® RST Windows runtime driver version **15.5.0.1051 (16.3 PV for the latest features)** release or later

6.2.3 Operating System

- Microsoft Windows 10 x64 bit Editions (latest RTM e.g. Redstone X)
- GPT disks only (MBR not supported)

6.2.4 "Optane™ Memory Ready" Logo: Intel® RST 5MB Metadata Requirement

This notice is for System Manufacturers and System Integrators who will ship "Optane™ Memory Ready" systems in the following configurations (or any system that they wish an end-user to be able to enable Optane™ Memory):W

- Complete systems with the OS already installed **AND**
- PCH SATA controller mode set to (AHCI) **OR** (Intel RST) **OR** (Intel RST Premium) **AND**

- The system does not have Optane™ Memory enabled

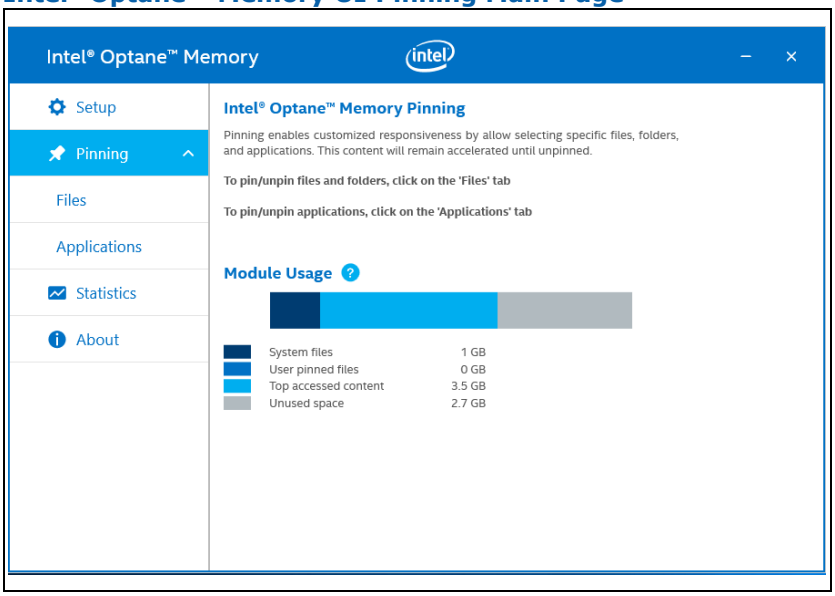
When all these conditions exist, especially if the manufacturer or integrator partitions the system drive with a hidden partition as the final partition (the end of the disk) or if the OS installation places the WinRE partition at the end of the disk, there must be unallocated space of 5MB reserved or left at the end of the disk (at the max LBA after the final partition). This space will be required for RST metadata when updating those systems at a later date to Optane™ Memory enabled accelerated systems.

6.3 User Pinning Files/Folders/Applications

6.3.1 Pinning using the Intel® Optane™ Memory UI

Limitations	The following objects cannot be pinned: 1. Quick Access 2. This PC 3. Any disk Root directory (e.g. Local Disk (C:)) 4. Control Panel 5. Libraries folder 6. Network objects 7. Recycle Bin
Pinning Tab:	1. Launch the Intel® Optane™ Memory UI 2. Click on the 'Pinning' tab in the left navigation pane. This will bring up the 'Intel Optane Memory Pinning' page of the UI The page displays information on pinning and unpinning of files/folders and applications. Also there is a graphical display of the current usage of the Optane™ memory module in terms of the capacity of files/content that is in the module and any un-used space remaining in the module.

Intel® Optane™ Memory UI Pinning Main Page



Files Sub-Tab: Click on the 'Files' tab to launch the 'Currently Pinned Files and Folders' page

How-to pin files:

1. Click [+Document] button; this launches file explorer
2. Browse to the file you wish to pin and do one of the following:
   a) double click on the file to pin it
   b) click on the file to highlight it, then click the [Open] button to pin it
3. To pin multiple files in the same folder, click on one file then use either the <Shift> key to highlight the multiple sequential files or use the <Ctrl> key to select multiple files

How-to pin folders:

1. Click [+Folder] button; this launches the 'Browse for Folder' window.
2. Browse to the folder you wish to pin and click to highlight it

Note: *Only the files directly in the folder will be pinned. Subfolders will not be pinned. If there are only subfolders and no individual files in the folder, then pinning the folder will only pin an empty folder.*

3. Click [OK] to pin the folder

How-to Unpin files/folders:

1. While in the 'Currently Pinned Files and Folders' window, use the left click button of the mouse to click through the file structure until you find the file/folder that you wish to unpin. Click on the file/folder to highlight it (for multiple files/folders use the <Shift> or <Ctrl> key to highlight the files/folders)
2. Click the [Unpin] button to unpin the file/folder

Applications Sub-Tab:	Click on the 'Applications' tab to launch the 'Currently Pinned Applications' page
How-to pin applications:	1. Click [+Application] button; this launches the Windows file Explorer window. 2. Browse to the application that you wish to pin and do one of the following: a) double click on the file to pin it b) click on the file to highlight it, then click the [Open] button to pin it
How-to Unpin applications:	1. While in the 'Currently Pinned Applications' window, use the left click button of the mouse to highlight the application you wish to unpin 2. Click the [Unpin] button to unpin the application

6.3.2 User Pinning using Windows File Explorer application

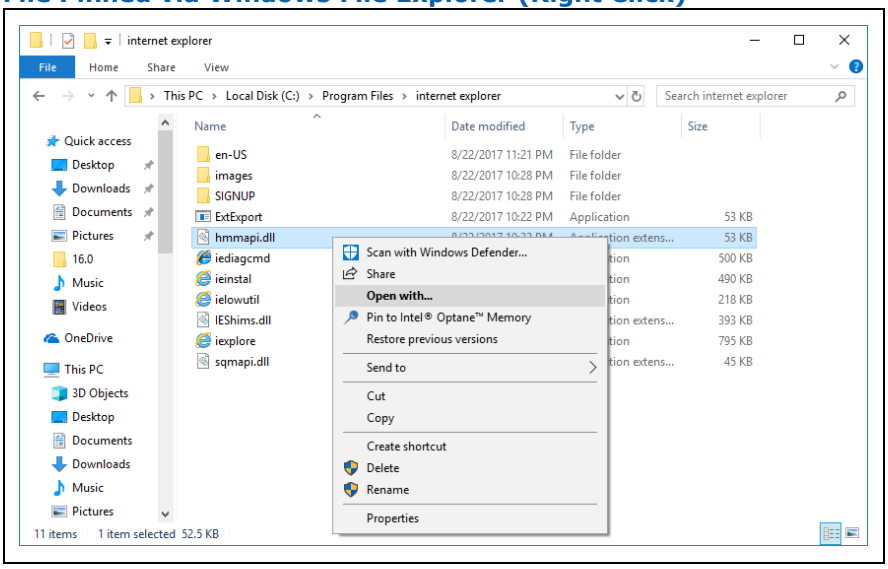
Pinning Files (non-application):

1. Launch the Windows File Explorer
2. Navigate to the file that you wish to pin to Intel® Optane™ memory
3. Use the mouse to right click on the file that you wish to pin to Intel® Optane™ memory
4. In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the file into the file cache. The file will be **hard pinned**:
   a) To pin multiple files in the same folder, left click on one file then use either the <Shift> key to highlight the multiple sequential files or use the <Ctrl> key to select multiple files that are not listed sequentially
   b) Now use the mouse to right click on one of the highlighted files
   c) In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the files into the file cache

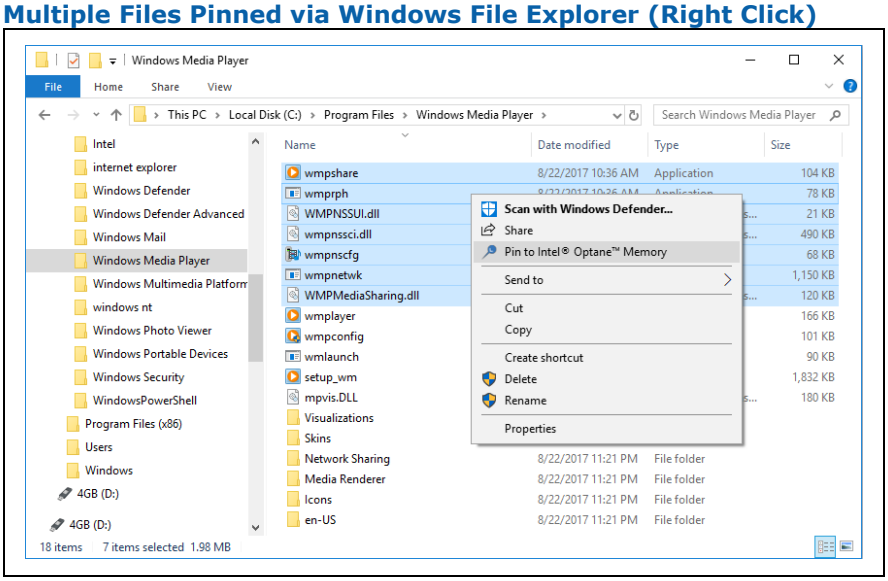
File Types:
All file types highlighted will be pinned, including applications.

Pinning Types:
All files pinned will be **hard pinned**, except application executables which will be **soft pinned**.

File Pinned via Windows File Explorer (Right Click)

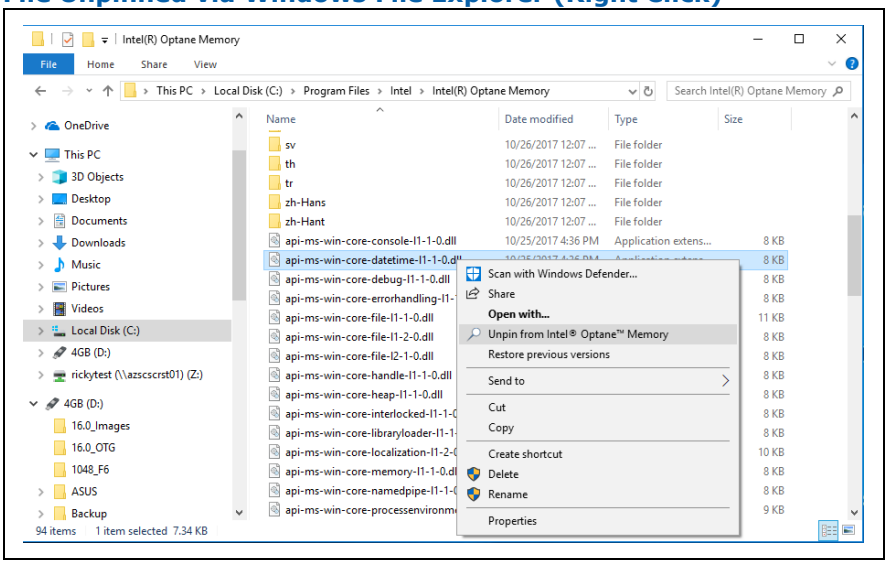


Multiple Files Pinned via Windows File Explorer (Right Click)



Unpinning Files:	1. Launch Windows File Explorer 2. Navigate to the file that you wish to unpin from Intel® Optane™ memory 3. Use the mouse to right click on the file that you wish to unpin 4. In the pop-up menu, left click the 'Unpin from Intel® Optane™ Memory' option to unpin the file

File Unpinned via Windows File Explorer (Right Click)



Pinning Folders:

1. Launch Windows File Explorer
2. Navigate to the folder that you wish to pin to Intel® Optane™ memory
3. Use the mouse to right click on the folder that you wish to pin to Intel® Optane™ memory
4. In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the folder

Note: *Only the files directly in the folder will be pinned. Subfolders will not be pinned. If there are only subfolders and no individual files in the folder, then pinning the folder will only pin an empty folder. In the example below, only the files in the red box will be pinned*

5. To pin multiple folders in the same folder, click on one folder then use either the <Shift> key to highlight the multiple sequential folders or use the <Ctrl> key to select multiple folders

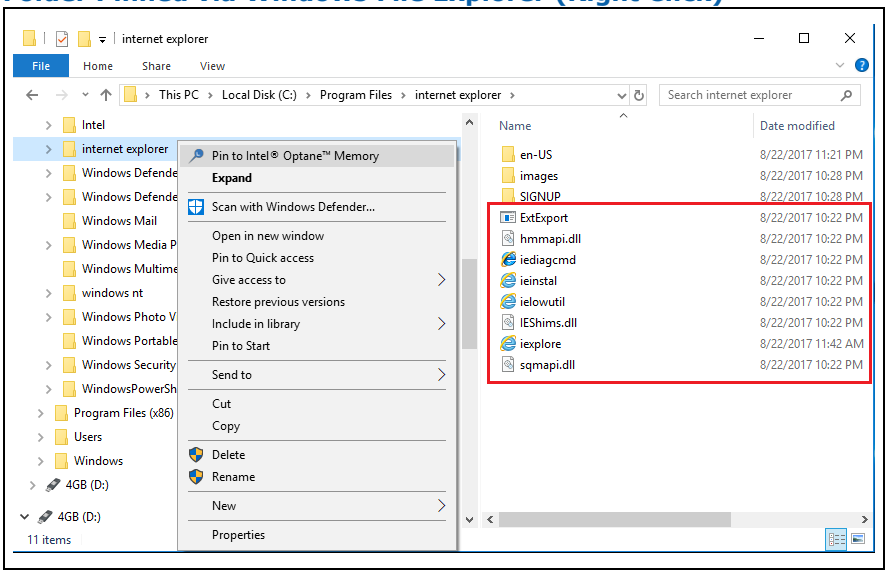
File Types:

All file types directly in the folder will be pinned, including applications.

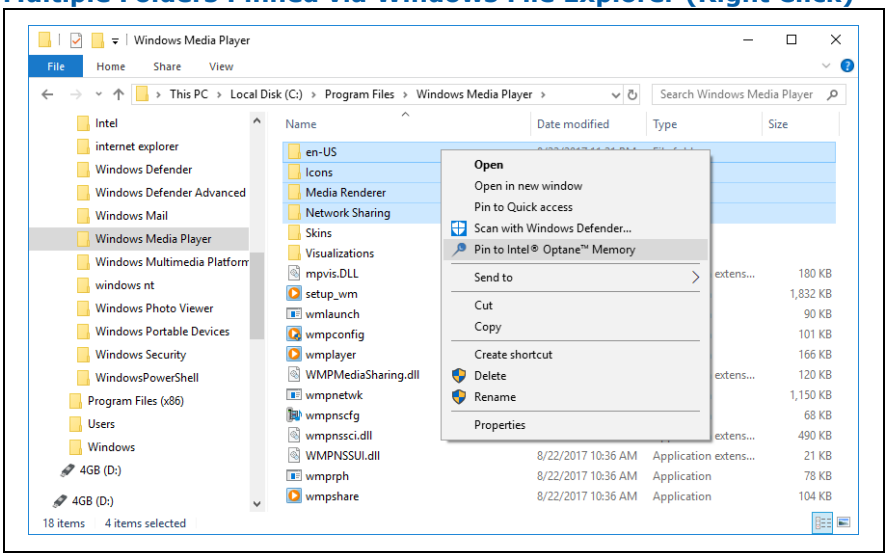
Pinning Types:

All files pinned will be **hard pinned**, including application.

Folder Pinned via Windows File Explorer (Right Click)



Multiple Folders Pinned via Windows File Explorer (Right Click)

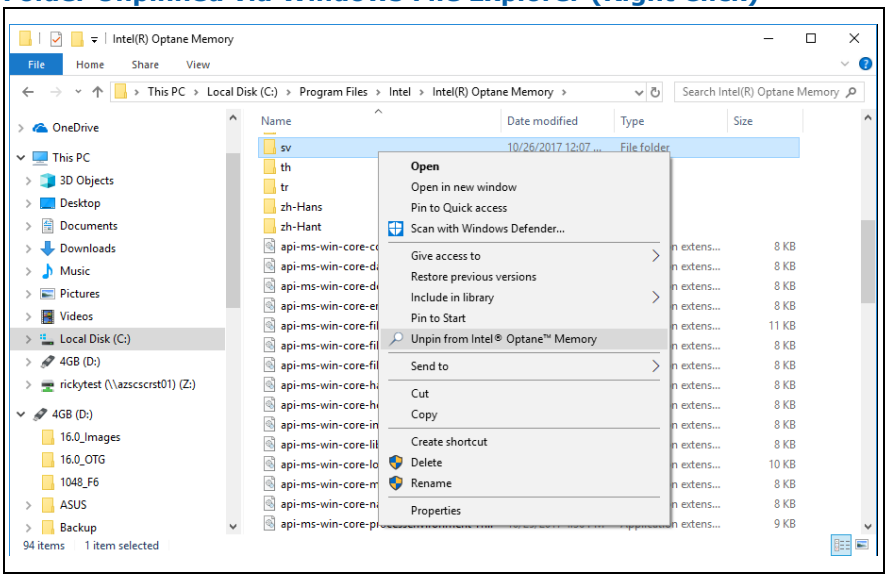


Unpinning folders:

1. Launch Windows File Explorer
2. Navigate to the folder that you wish to unpin from Intel® Optane™ memory
3. Use the mouse to right click on the folder that you wish to unpin to Intel® Optane™ memory
4. In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the folder

5. To unpin multiple folders in the same folder, click on one folder then use either the <Shift> key to highlight the multiple sequential folders or use the <Ctrl> key to select multiple folders

Folder Unpinned via Windows File Explorer (Right Click)



Pinning applications:

1. Launch the Windows File Explorer
2. Navigate to the application that you wish to pin to Intel® Optane™ memory
3. Use the mouse to right click on the application executable file that you wish to pin to Intel® Optane™ memory
4. In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the application into the file cache. The application will be **soft pinned**:
   d) To pin multiple applications in the same folder, left click on one application then use either the <Shift> key to highlight the multiple sequential apps or use the <Ctrl> key to select multiple apps that are not listed sequentially
   e) Now use the mouse to right click on one of the highlighted apps
   f) In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the apps into the file cache

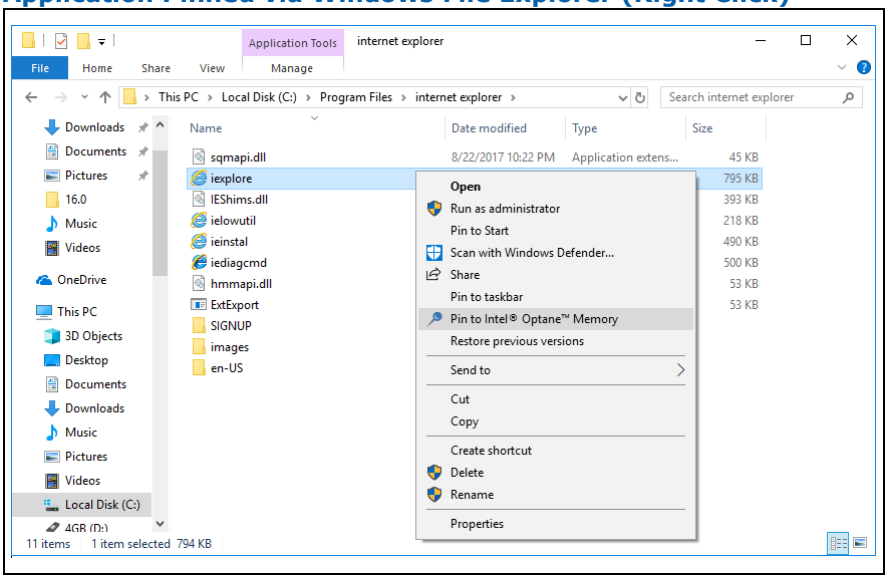
File Types:

All applications highlighted will be pinned.

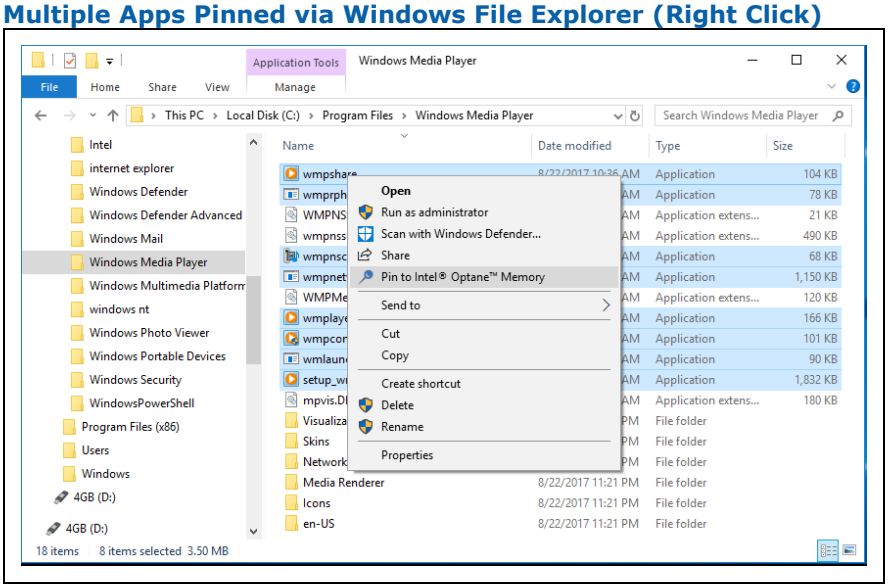
Pinning Types:

All applications pinned will be **soft pinned**.

Application Pinned via Windows File Explorer (Right Click)



Multiple Apps Pinned via Windows File Explorer (Right Click)



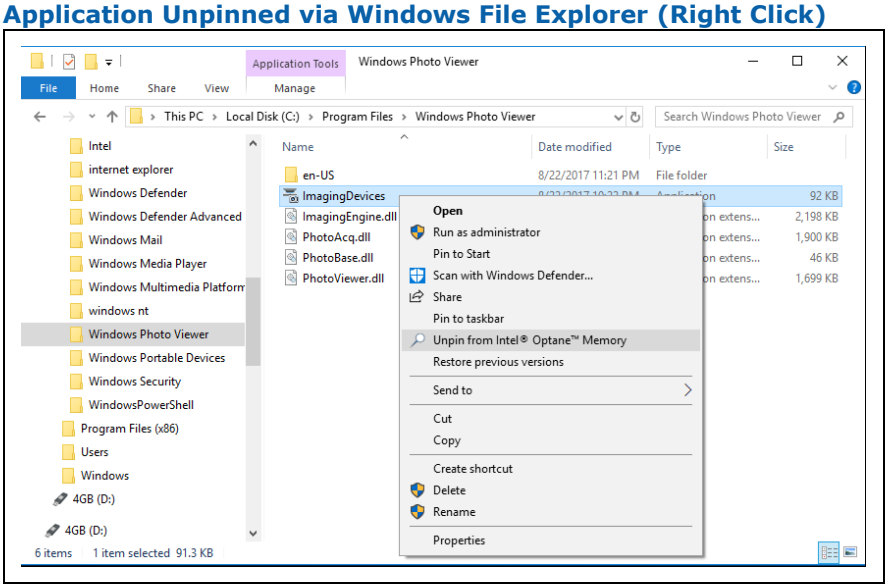
a)

How-to Unpin applications:

1. Launch Windows File Explorer
2. Navigate to the folder that you wish to unpin from Intel® Optane™ memory
3. Use the mouse to right click on the folder that you wish to unpin to Intel® Optane™ memory

4. In the pop-up menu, left click the 'Pin to Intel® Optane™ Memory' option to pin the folder
5. To unpin multiple folders in the same folder, click on one folder then use either the <Shift> key to highlight the multiple sequential folders or use the <Ctrl> key to select multiple folders

Application Unpinned via Windows File Explorer (Right Click)



6.4 Configuring Intel® Optane™ Memory (Post-Factory Build Environment[1])

This section documents how to enable System Acceleration with Intel® Optane™ Memory on systems that are not shipped with the feature enabled. It focuses on systems being installed/configured by end-users/system integrators/others outside of the factory environment.

Note: : For configuring Intel® Optane™ Memory in the factory, review the next *chapter*.

WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process

6.4.1 Enable Optane™ Volume on a Brand New System

6.4.1.1 BIOS

There are no pre-OS configuration tools for Optane™ available to the end-user.

6.4.1.2 OS runtime

This section instructs how to enable the System Acceleration with Optane™ memory after the installation of the OS. Users who want to have a completely automated Intel® Optane™ Memory configuration process after the OS installation are recommended to use this method.

6.4.1.2.1 Intel® Optane™ Memory UI/Installer (All User Levels)

This configuration method can be used by experienced and inexperienced users. Users who want to have a completely automated Intel® Optane™ Memory configuration process after the OS installation are recommended to use this method. It automates all of the BIOS configuration tasks involved in setting up the entire system to support System Acceleration with Intel® Optane™ Memory.

System prep:	Boot Drive Acceleration (requires 2 drives): • Ensure that you have installed: A. A SATA HDD/SSD/SSHD (the slow media to install OS) B. An Intel® Optane™ NVMe SSD or memory module (fast media to be the cache) in a PCH remapable port or a CPU PEG port
	Data Drive Acceleration (requires 3 drives): • Ensure that you have installed: A. A drive to install the OS to (no specific requirement) B. A SATA HDD/SSD/SSHD (the slow media data drive); must be larger than the fast media drive C. A PCIe Intel® Optane™ NVMe SSD or memory module (fast media to be the cache) in a PCH remapable port or a CPU PEG port.
BIOS Setup:	Enter the BIOS and configure the following: 1. Locate the menu to configure the SATA PCH controller 2. Set the SATA mode to one of the non-AHCI modes: a. Intel RST with System Acceleration with Optane Memory…, or b. Intel RST Premium with System Acceleration with Optane Memory…

	3. Locate the setting for enabling/disabling remapping for the PCIe port where the fast media drive is located 4. Set the port to remapped enabled (controlled by RST)
OS installation:	Obtain your copy of the supported Windows 10 installation media and install the OS using the default installation or any customizations desired.
Verify OS detects Optane™ module:	1. Complete the OS installation and boot to the Windows desktop 2. Launch Windows Device Manager 3. Expand 'Disk drives' and confirm that all drives are detected A. Intel® Optane™ model disk (fast media) B. HDD/SSD (slow media) C. If accelerating a data drive, then there should be a 3rd drive for the OS 4. Exit Device Manager
Optane™ volume creation:	*The following steps in this process will automatically accomplish all or some of the following tasks depending on the state of your system:* • *Install the required SW needed for Optane™ Memory* • *Select the proper mode for the PCH SATA controller* • *Enable remapping required for the Optane™ PCIe NVMe device to be detected by Intel® RST* • *Installs the Intel® Optane™ Memory UI* • *And finally, with user input, enables System Acceleration with Intel® Optane™ Memory* *If any of the above steps are already in the required state the installer will automatically skip those steps that are not required and complete the installation and configuration process* ***WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process***

	1. Obtain the Intel® Optane™ Memory SW/driver installation package and run the executable (**SetupOptaneMemory.exe**) 2. Install the defaults 3. When this part of the installation process completes the installer will prompt you to 'Restart', click [Finish] to restart the system: A. During this reboot the system will accomplish any required system configuration tasks (BIOS settings) and may reboot more than once B. When complete, the system will reboot back to the Windows desktop 4. There will be a few seconds delay; the installation process continues and completes driver installation 5. A pop-up ("You're Almost There …") will prompt you to continue to enable Optane™; click [Next] then [Yes] and the Intel® Optane™ UI will launch 6. To complete the Optane™ enabling process (*to combine the Optane™ memory device and the OS system drive into a single accelerated Optane™ volume*), go to the 'Setup' page, A. Select an Optane drive from the "Compatible drive" drop down box (if only a single Optane module is present in the system then skip this step, there will be no drop down box) B. If multiple compatible drives are in the system select the drive to be accelerated: 1. For OS drive acceleration select the disk labeled "SYSTEM" 2. For data drive acceleration select any disk not labeled "SYSTEM" C. Click the [Enable] button D. Click the "Erase all data on the Intel® Optane™ memory module" box to place a check in the box E. Click [Continue] 7. Once any necessary file migration completes, click the [Restart] button to reboot 8. Reboot to Windows to complete the configuration process
Verify the system is Optane™ Memory accelerated:	From Windows desktop, launch the Windows Device Manager 1. Go to 'Disk Drives' and click to expand • You should see a single drive labeled "Intel Optane+slow media" (e.g. "Intel Optane+1.0TBHDD") • If you are accelerating a data drive then you will see the OS drive listed as well 2. Go to Disk Manager

- The Intel® Optane™ memory module should not be listed. It will be combined with either the OS drive or a data drive.

3. Another method is to launch the Intel® Optane™ Memory UI and see the 'Status' page.
   These three items confirm that the Intel® Optane™ NVMe SSD/memory module and the large capacity HDD/SSD have been combined to form the accelerated Intel® Optane™ volume.

6.4.1.2.2 Intel® RST UI/Installer (Experienced Users)

This section instructs how to enable the System Acceleration with Optane™ technology as part of the OS installation. Experienced users who want to have full control of the configuration process are recommended to use this method.

System prep:

1. Ensure the computer HW meets the minimum requirements for this feature (see the Minimum Requirements section above). This includes installing the Intel® Optane™ memory module or Intel® Optane™ NVMe SSD in the proper remapable M.2 connector or PCIe slot.
2. Confirm that the system BIOS has a supported Intel® RST pre-OS UEFI driver integrated:
   a) Enter your system BIOS and locate the Intel® Rapid Storage Technology HII UI
   b) At the top of the UI page confirm that the Intel RST UEFI Driver version is 15.9.3.xxxx or later
3. Set SATA PCH I/O controller to proper mode:
   a) While still in the BIOS find the page that allows you to configure the PCH and SATA mode to "Intel RST Premium…" **(AHCI mode is not supported)**
   b) Enable remapping so that Intel® RST can control the Intel® Optane™ memory module or Intel® Optane™ NVMe SSD

Note: The location of the above functions/settings in the system BIOS on your computer will vary based on the manufacturer of your computer system. If you have trouble locating these settings contact your system manufacturer. Alternately you can follow the instructions for all users in the previous section.

OS installation:	Obtain your copy of the Windows 10 installation media and install the OS using the default installation or any customizations desired.
Verify that OS detects Optane™ module:	1. Complete the OS installation and boot to the Windows desktop 2. Launch Windows* Device Manager 3. Expand 'Disk drives' and confirm that all drives are detected A. Intel® Optane™ model disk (fast media) B. HDD/SSD (slow media) C. If accelerating a data drive, then there should be a 3rd drive for the OS D. Exit Device Manager
Optane™ volume creation:	***WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process*** 1. Obtain the Intel® RST SW/driver installation package and run the executable (SetupRST.exe) 2. Install the defaults and reboot the computer 3. From Windows* desktop, find and launch the Intel® RST UI application. 4. The application will open to the 'Status' page 5. Click the 'Intel® Optane™ Memory' tab 6. Click the "Enable" link to start the enabling process. 7. A pop-up will prompt you to select: A. From drop down, select the NVMe fast media Optane™ device B. From drop down, select the SATA disk to be accelerated (for OS drive acceleration, select the drive labeled 'SYSTEM"; for data drive acceleration select a drive not labeled "SYSTEM") 8. Click OK to start the enable process 9. Depending on the size of your Optane™ module, the following may be seen: A. Progress indicator: Let progress indicator complete to 100% B. Text: Finalizing operation… 10. Click the [Reboot] button to complete the process 11. System reboots into Windows to complete the enabling process.
Verify system is Optane™ Memory accelerated:	From Windows desktop, launch the Windows Device Manager 1. Go to 'Disk Drives' and click to expand

- You should see a single drive labeled "Intel Optane+slow media" (e.g. "Intel Optane+1.0TBHDD")
- If you are accelerating a data drive you will see at least one other drive that is the OS drive

2. Go to Disk Manager
   - The Intel® Optane™ module should not be listed.
3. Another method is to launch the Intel® RST UI or Intel® Optane™ UI, depending on your install, and see the Status page.
   These three items confirm that the Intel® Optane™ NVMe SSD/memory module and the large capacity HDD/SSD have been combined to form the accelerated Intel® Optane™ volume.

6.4.2 System Upgrade (Windows*10 OS Already Installed)

This section covers upgrading those systems that have Windows* 10 already install by installing a new Intel® Optane™ memory module then enabling System Acceleration with Intel® Optane™ Memory.

6.4.2.1 Intel® Optane™ Memory UI/Installer

System starting configuration:

- Windows 10* x64 OS installed on SATA HDD/SSD
- Intel® RST driver installed
- System BIOS has Intel® RST pre-OS version 15.5.0.2875 or later integrated

Optane™ module installation:

1. Power down the system and install the Intel® Optane™ memory device into an M.2 connector located either on the motherboard or PCIe adapter card plugged into a remapable PCIe slot or module connected to CPU PEG port
2. Reboot system to Windows

Verify OS detects Optane™ module:

Note: Unless the following conditions are met, the Optane™ memory module will not be detectable:

a) Intel® RST 15.5 PV or later driver installed (if not installed then the Optane™ module will not be detectable yet; skip to the next section and proceed to install the 15.5 or later SW package with driver)
b) The BIOS has the Intel RST 15.9.3 or later UEFI preOS driver installed

c) The PCIe port with the Optane™ module is remapped

1. Boot to the Windows desktop
2. Launch Windows Device Manager
3. Expand 'Disk drives' and confirm that all drives are detected:
    A. Intel® Optane™ Memory module/Disk
    B. SATA HDD/SSD/SSHD (to be accelerated)
    C. Additional disk for OS if accelerating a data drive

Optane™ volume creation:

The following steps in this process will automatically accomplish all or some of the following tasks depending on the state of your system:

- *Install the required SW needed for Optane™*
- *If in AHCI mode, switch the PCH SATA controller to the proper 'Intel RST…' mode*
- *Enable remapping required for the Optane™ PCIe NVMe device*
- *Installs the Optane™ Memory UI*
- *And finally, with user input, enable System Acceleration with Intel® Optane™ Memory*

If any of the steps above are already in the required state the installer will automatically skip those steps that are not required and complete the installation and configuration process.

WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process

1. Obtain the Intel® Optane™ SW/driver installation package and run the executable (**SetupOptaneMemory.exe**)
2. Install the defaults
3. When this part of the installation process completes the installer will prompt you to 'Restart', click [Finish] to restart the system:
    A. During this reboot the system will accomplish any required system configuration tasks (e.g. BIOS settings like switching from AHCI mode) and may reboot more than once
    B. When complete, the system will reboot back to the Windows desktop
4. There will be a few seconds delay; the installation process continues and completes driver installation
5. A pop-up ("Thank you for installing…") will prompt you to continue to enable Optane™; click [Yes] and the Optane™ UI will launch

	6. To complete the Optane™ enabling process (to combine the Optane™ device and the OS system drive into a single accelerated Optane™ volume), on the 'Setup' page, click the [Enable] button and click the [Yes] button to proceed 7. Once any necessary file migration completes, click the [Restart] button to reboot 8. Reboot to Windows and this completes the configuration process
Verify the system is Optane™ Memory accelerated:	From Windows desktop, launch the Windows Device Manager 1. Go to 'Disk Drives' and click to expand • You should see a single drive labeled "Intel Optane+slow media" (e.g. "Intel Optane+1.0TBHDD") 2. Go to Disk Manager • If you see the Optane™ volume displayed in 'Disk drives', then you should see a single Windows OS disk. The Optane™ module should not be listed. 3. Another method is to launch the Intel® RST UI or Intel® Optane™ UI, depending on your install, and see the Status page. These three items confirm that the Intel® Optane™ NVMe SSD/memory module and the large capacity HDD/SSD have been combined to form the accelerated Intel® Optane™ volume.

6.4.2.2 Intel® RST UI/Installer

This method is for upgrading a system that is in one of the Intel RST Optane™ modes and running the Intel RST driver

System starting configuration:	• Windows 10 OS installed on SATA HDD/SSD • RST driver installed • System BIOS SATA controller set to one of the Intel RST Optane™ modes **Note:** *: The system must meet minimum requirements to support System Acceleration with Intel® Optane™ technology. See the 'Minimum Requirements' section above*
Optane™ module installation:	1. Power down the system and install the Optane™ device into an M.2 connector located either on the motherboard

	or PCIe adapter card plugged into a remapable PCIe slot or a CPU PEG port 2. Enter BIOS and enable remapping on the port where the Optane™ device is plugged into 3. Reboot to Windows
Verify OS detects the Optane™ device:	1. Boot to the Windows desktop 2. Launch Windows Device Manager 3. Expand 'Disk drives' and confirm that newly installed Intel Optane™ device is detected as a drive A. Intel Optane™ Model Disk (fast media) B. HDD/SSD (slow media)
Optane™ volume creation:	***WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process*** 1. Obtain the Intel RST SW/driver installation package and run the executable (SetupRST.exe) 2. Install the defaults and reboot the computer 3. From Windows desktop, find and launch the Intel® RST UI application. 4. The application will open to the 'Status' page 5. Click the 'Optane™ Memory' tab 6. Click the "Enable" link to start the enabling process. 7. A pop-up will prompt you to 'Select a compatible fast drive:' (the Optane™ NVMe device) 8. Click [Yes] button to start the enable process 9. Depending on the size of your Optane™ module, a progress indicator may be displayed; let it complete to 100% 10. Click the [Reboot] button to complete the process 11. System reboots into Windows to complete the enabling process.
Verify the system is Optane™ Memory accelerated:	From Windows desktop, launch the Windows Device Manager 1. Go to 'Disk Drives' and click to expand • You should see a single drive labeled "Intel Optane+slow media" (e.g. "Intel Optane+1.0TBHDD") 2. Go to Disk Manager • If you see the Optane™ volume displayed in 'Disk drives', then you should see a single Windows OS disk. The Optane™ module should not be listed.

	3. Another method is to launch the Intel® RST UI or Intel® Optane™ UI, depending on your install, and see the Status page. These three items confirm that the Intel® Optane™ NVMe SSD/memory module and the large capacity HDD/SSD have been combined to form the accelerated Intel® Optane™ volume.

6.4.3 System Upgrade (Windows* 10 OS Already Installed with BIOS in AHCI mode)

This section covers upgrading those systems that have Windows* 10 already install by installing a new Intel® Optane™ memory module then enabling System Acceleration with Intel® Optane™ Memory.

Note: This upgrade path is only possible via the Intel® Optane™ Memory UI installer program.

6.4.3.1 Intel® Optane™ UI/Installer

System starting configuration:	• Windows* 10 OS installed on SATA HDD/SSD • Intel® RST AHCI driver or default inbox AHCI driver installed • System BIOS SATA controller in AHCI mode
Intel® Optane™ memory module installation:	1. Power down the system and install the Intel® Optane™ memory module or NVMe SSD into an M.2 connector located either on the motherboard or a PCIe adapter card plugged into a remapable PCIe slot or a CPU PEG port 2. Reboot system to Windows
Verify OS detects Intel® Optane™ module:	Note: Skip this section if the Intel® RST AHCI driver is installed 1. Boot to the Windows desktop 2. Launch Windows Device Manager 3. Expand 'Disk drives' and confirm that both drives are detected: A. Intel® Optane™ memory module/disk B. HDD/SSD

Intel® Optane™ volume creation:

The following steps in this process will automatically accomplish some or all of the following tasks depending on the state of your system:

- *Install the required SW needed for Optane™*
- *Switch the PCH SATA controller from AHCI mode to the proper 'Intel RST…' mode*
- *Enable remapping required for PCIe NVMe devices*
- *Installs the Intel® Optane™ Memory UI*
- *And finally prompt to enable System Acceleration with Intel® Optane™ Memory*

If any of the steps above are already in the required state the installer will automatically skip those steps that are not required and complete the installation and configuration process

WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process

1. Obtain the Intel® RST and Optane™ SW/driver installation package and run the executable (SetupOptaneMemory.exe)
2. Install the defaults
3. When this part of the installation process completes the RST installer program will prompt for you to restart your system to continue the process (click default [Finish] to restart the computer).
4. Depending upon the initial configuration of your system:
    A. During this reboot the system will accomplish any required system configuration (BIOS settings) and may reboot more than once
    B. When all necessary configuration processes complete, the system will reboot back to the Windows desktop
5. The installation continues and a pop-up box will prompt you to continue to enable Intel® Optane™ volume; click [Yes] to continue
6. Click [Yes] for User Account Control pop-up
7. The Intel® Optane™ Memory UI will launch, click 'Enable' then click 'Yes' to continue to combine your OS system drive or a data drive and your Intel® Optane™ memory module (or SSD) into a single accelerated Optane™ volume
8. Once the migration process successfully completes you will be prompted to reboot to complete the process; Click 'Restart'

	9. Your computer will reboot to Windows* completing the Intel® Optane™ Memory enabling process
Verify the system is Optane™ Memory accelerated:	From Windows desktop, launch the Windows Device Manager 1. Go to 'Disk Drives' and click to expand • You should see a drive labeled "Intel Optane+slow media" (e.g. "Intel Optane+1.0TBHDD") 2. Go to Disk Manager • The Optane™ module should not be listed. 3. Another method is to launch the Intel® RST UI or Intel® Optane™ UI, depending on your install, and see the Status page. These three items confirm that the Intel® Optane™ NVMe SSD/memory module and the large capacity HDD/SSD have been combined to form the accelerated Intel® Optane™ volume.

6.5 Intel® Optane™ Memory Member Disk Upgrades

This section covers upgrading an Optane-enabled system with a new system disk or a new Intel® Optane™ memory module.

6.5.1 Slow Disk (SATA HDD, SSD, SSHD) Upgrade/Replacement

Disable Optane™:

Before replacing any component of your Optane™ Memory storage solution, you must always disable Optane™ Memory to ensure all data has been migrated from the Optane™ module and synched with the slow disk.

NOTE: You should backup any important data prior to starting this process!

Intel® Optane™ Memory UI installed:

1. From Windows desktop, launch the Intel® Optane™ Memory UI
2. Click on the 'Setup' tab
3. Click Disable to start the disabling process
4. Confirm the action
5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (progress indicator reaches 100%), the [Restart] button will be displayed
6. Click the [Restart] button to reboot the system and complete the disablement process

Intel® RST UI installed:

1. From Windows desktop, launch the RST UI
2. Click the 'Optane™ Memory' tab at the top
3. Click Disable to start the disabling process
4. Confirm the action
5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (progress indicator reaches 100%), the [Reboot] button will be displayed
6. Click the [Reboot] button to reboot the system and complete the disablement process

Reinstall your OS or copy over your current to the new slow disk	***WARNING: In all Intel® Optane Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process*** At this point you can reinstall the operating system and rebuild your system or you can use whatever process you wish to use to clone/copy your current system and apply it to the new slow disk: Replacing an accelerated OS drive: 1. Install the new slow media (SATA HDD, SSD, or SSHD) a. If the slow media is the OS disk, then continue with step 2 b. If the slow media is a data drive 2. Install the Windows 10 OS or apply clone/copy 3. Once the OS has been installed or transferred to the new slow media, re-enable the Intel® Optane™ Memory 4. Complete any customizations or application installs 5. Replacing an accelerated data drive 6. If possible backup all data from original data drive 7. Remove the old data drive and insert the new data drive 8. Re-enable data drive acceleration with Optane™ Memory 9. Copy all required data to the new data drive 10. EOT

6.5.2 Intel® Optane™ Memory Module Upgrade/Replacement

Disable Optane™ Memory:	1. From Windows desktop, launch the Intel® Optane™ Memory UI • *or launch the Intel® RST UI* 2. Click the 'Setup' tab at the left • *or for the Intel® RST UI, click the Intel Optane™ Memory tab at the top* 3. Click Disable to start the disabling process 4. Confirm the action 5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (after the progress indicator reaches 100%), the [Restart] *or [Reboot]* button will be displayed

	6. Click the [Restart] *or [Reboot]* button to reboot the system and complete the disablement process
Replace the Intel® Optane™ memory module:	1. Power down the computer 2. Open the computer and locate the old Optane™ memory module that you wish to replace ***(Note: consult your computer manufacturer for the location of the slot for the Optane™ memory module and instructions to remove and insert an M.2 module)*** 3. Remove the old module and insert the new module. 4. Close the computer and power it on and boot into the Windows OS.
Re-enable Intel® Optane™ Memory:	***WARNING: In all Intel® Optane™ Memory enabling scenarios, all data that is on the fast disk prior to enabling will be deleted during the Optane™ Memory enabling process*** 1. From Windows desktop, launch the Intel® Optane™ Memory UI • *or launch the Intel® RST UI* 2. Click on the 'Setup' tab • *or for the Intel® RST UI, click the Intel Optane™ Memory tab at the top* 3. Click Enable to start the enabling process 4. Select the fast media to use for enabling and continue 5. Once the UI has completed all tasks for enabling Optane™ and any necessary file migration has completed (after the progress indicator reaches 100%), the [Restart] *or [Reboot]* button will be displayed 6. Click the [Restart] *or [Reboot]* button to reboot the system and complete the enablement process

6.6 Managing Intel® Optane™ Memory

This section describe the management functions available in the various user interfaces that are available to the end-user

6.6.1 During the Pre-OS (Pre-Boot)

6.6.1.1 BIOS HII User Interface for Intel® RST

The BIOS HII UI only allows the user to view storage sub-system configuration and the ability to disable the Optane™ configuration

Disable Optane™ configuration:

1. Enter your system BIOS and locate the Intel® Rapid Storage Technology HII UI
2. You will see listed the Optane™ volume and the physical disks that are used in the Optane™ volume
3. Highlight the Optane™ volume and hit <Enter>
    A. Deconcatenate : Use this option to disable Optane™ and return the system to a non-accelerated state.
        i. [X] Checkbox to preserve user data upon deconcatenation; checked is the default.
            - When checked this action will take time to migrate data from the cache device to the HDD/SSD.

            - **WARNING!!** When unchecked there is no time taken to preserve the user data (any data on the drive will be **DELETED**). RST removes the RST metadata from the disk and does not migrate any user or system data to the HDD/SSD
        ii. <No> Decision box to confirm deconcatenation action; Yes or No (No is default)

Note: The location of the above functions in the system BIOS on your computer will vary based on the manufacturer of your computer system. If you have trouble locating these settings contact your system manufacturer.

6.6.2 During the OS Runtime

6.6.2.1 Intel® Optane™ Memory UI

This section describes management capabilities of the Intel® Optane™ Memory UI

6.6.2.1.1 Enabling/Disabling Intel® Optane™ Memory Acceleration

Setup Tab:

Enable Optane™ configuration:

1. From Windows desktop, launch the Intel® Optane™ Memory UI
2. Click on the 'Setup' tab
3. Click Enable to start the enabling process
4. Select the fast media to use for enabling and continue

	5. Once the UI has completed all tasks for enabling Optane™ and any necessary file migration has completed, the [Restart] button will be displayed 6. Click the [Restart] button to reboot the system and complete the enablement process
Setup: ***Disable Optane™ configuration:***	1. From Windows desktop, launch the Intel® Optane™ Memory UI 2. Click on the 'Setup' tab 3. Click Disable to start the disabling process 4. Confirm the action 5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (progress indicator reaches 100%), the [Restart] button will be displayed 6. Click the [Restart] button to reboot the system and complete the disablement process
Statistics Tab:	***Note!*** *This feature only supports Intel® Optane™ devices that have a capacity of 32GB or more.* There is no user interactions available other than to view the status of "Intel® Optane™ Memory Optimization Schedule". The indicators are: A. Last Optimization B. Next Scheduled Optimization
About Tab:	There are no user actions available other than to view the following information: A. The version number of the Intel® Optane™ UI B. A link to Intel support website C. And some application licensing information

6.6.2.2 Intel® RST UI

This section describes Optane™ management capabilities of the Intel® RST UI.

Note that there is no 'User Pinning' functionality in the Intel® RST UI.

Enable Optane™ configuration:	1. From Windows desktop, launch the RST UI 2. Click the 'Optane™ Memory' tab at the top 3. Click Enable to start the enabling process

	4. Select the fast media to use for enabling and continue 5. Once the UI has completed all tasks for enabling Optane™ and any necessary file migration has completed , the [Reboot] button will be displayed 6. Click the [Reboot] button to reboot the system and complete the enablement process
Disable Optane™ configuration:	1. From Windows desktop, launch the RST UI 2. Click the 'Optane™ Memory' tab at the top 3. Click Disable to start the disabling process 4. Confirm the action 5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (progress indicator reaches 100%), the [Reboot] button will be displayed 6. Click the [Reboot] button to reboot the system and complete the disablement process

6.6.3 Roaming Optane™ Volumes

Roaming Optane™ volumes intact between computers can be done safely if certain conditions are met (both member disks must remain paired; do not physically separate):

1. The target computer (where you are moving the volume to) is the same or similar Optane™ capable system as the source computer (where you are moving the volume from)
2. The target computer BIOS meets the minimum requirements to support Optane™
3. The BIOS must have the following properly set prior to booting to the roamed volume:
    a. The SATA PCH must be set to the proper mode to support Optane™ (Intel RST … or Intel RST Premium…) ***WARNING!: If the computer is in AHCI mode and you boot with the Optane volume installed, the Optane™ volume will become corrupted and all data on the Optane™ volume will be inaccessible and have to be reset to non-Optane which will delete all data on the volume. You MUST ensure that the SATA PCH mode is in the proper mode as indicated above!***
    b. The PCIe slot where the Optane™ module will be inserted must be set to remapping enabled
    c. The SATA port for the slow disk must be an internal SATA port
4. Once these prerequisites have been met it is safe to insert the Optane™ volume member disks into the target computer and boot to the Optane™ volume. **It is recommended to always backup important data prior to roaming as a precaution.**

6.7 I/O Error Handling

6.7.1 Host I/O Failures

Fast Media- I/O failure:	1. Action taken by the RST SW: A. The RST driver returns the I/O error to Windows. 2. System Ending State: A. No change: Optane acceleration remains enabled.
Slow Media- I/O failure:	1. Action taken by the RST SW: A. RST driver returns the I/O error to the Windows 2. System Ending State: A. No change: Optane acceleration remains enabled.

6.7.2 RST I/O Failures:

6.7.2.1 During Optane Enabling

READ failure to slow drive:	1. Action taken by the RST SW: A. Write bad ECC to fast drive B. Enabling is halted and rollback to disabling C. Notify user of significant amount of errors and recommend disabling the Optane™ Memory 2. System ending state: A. System Acceleration enabled. 3. Recommended corrective action: 1. Disable Optane™ Memory 2. Replicate system disk to new drive 3. Replace faulty slow drive with new drive 4. Re-enable Optane™ Memory
WRITE failure to fast disk:	1. Action taken by the RST SW: A. Continue in WT cache mode during Optane™ Memory enabling. B. Enabling is halted and rollback to disabling C. Display failure in UI. 2. System ending state: A. Optane™ Memory disabled. 3. Recommended corrective action: A. Power down system B. Replace faulty fast media with new fast media C. Enable Optane™ Memory

6.7.2.2 During Optane™ Memory Disabling

WRITE failure to slow media:	1. Action taken A. Continue disabling Optane™ Memory 2. System ending state: A. Optane™ Memory successfully disabled 3. Recommended corrective action: A. Reboot the system normally B. Let the Intel® RST driver handle any RST specific recovery tasks.
READ failure to fast media:	1. Action taken by the RST SW: A. Write bad ECC to slow drive for LBAs that cannot be read from fast drive B. Continue the Optane™ Memory volume disabling 2. System ending state: A. Optane™ Memory disabled B. System Normal state 3. Recommended corrective action: A. None

6.7.3 Missing Media Failures:

6.7.3.1 Drive Missing at Boot

Fast disk missing:	1. Action taken by the RST SW: A. If system boots and the Optane™ module is not detected by the RST UEFI driver, the RST UEFI driver will protect the current data on the slow media by disabling it and not exposing it to the boot manager. This allows the user the opportunity to locate and reinsert the missing fast media (or try additional reboots) to allow the UEFI driver additional opportunity to detect and pair the slow media and fast media and return the Optane™ volume to a normal state. 2. System ending state: A. System is unbootable. B. Slow media set to 'Disabled' state 3. Recommended corrective action: A. Power down and confirm the location of the missing Optane™ memory module.

	1. Reattach if missing 2. If not missing: a) Check the connector b) Disconnect the slow media and boot to BIOS menu c) Enter the RST HII UI d) Confirm that the Optane™ memory module is displayed (its Status should be 'Offline' in this scenario) e) If the Optane™ memory module is not present, then exit the RST HII UI and enter your BIOS menu where the SATA controller remapping function is located and confirm that the port location of the Optane™ memory module has remapping enabled f) If remapping not enabled, enable it. Save the configuration and power down the computer. 3. Reattach the slow media and power up the computer

Slow disk missing:

1. Action taken by the RST SW:
    A. If system boots and the slow media device is not detected by the RST UEFI driver, the RST UEFI driver will protect the current data on the Optane™ Memory module by disabling it and not exposing it to the boot manager. This allows the user the opportunity to locate and reinsert the missing slow media (or try additional reboots) to allow the UEFI driver additional opportunity to detect and pair the slow media and fast media and return the Optane™ volume to a normal state.
2. System ending state:
    A. System is unbootable.
    B. Fast media set to status Disabled[1]

***Note[1]**: a current errata on file has the status of the Disabled fast media displayed as 'Cache' when viewed in the RST preOS UEFI HII UI. The correct status should be displayed as 'Disabled' and will be fixed in a future release.*

3. Recommended corrective action:
    A. Power down and confirm the location of the missing slow media drive.
        1. Reattach if drive is physically missing
        2. If not missing:
            a) check the connection
            b) Boot into the BIOS and go to the Intel® Rapid Storage UI
            c) In the UI see if either the slow or fast disk is listed on the main page
            d) It appears the slow disk has experienced an uncorrectable failure and may need to be replaced. Go to the Disk Failure section in this chapter for instructions to replace the failed slow disk.

6.7.3.2 Drive Hot Unplugged

Hot unplug fast media disk:

1. Action taken
    A. No RST action taken; System is expected to experience a Windows* bug check (blue screen)
2. System ending state:

	A. Windows* blue screen bug check B. Upon next boot, the Intel® RST driver will evaluate the storage subsystem 3. Recommended corrective action: A. Reinsert the fast media back into the remapped PCIe slot B. Reboot the system normally C. Let the Intel® RST driver handle any RST specific recovery tasks.
Hot unplug slow media disk:	1. Action taken by the RST SW: A. No RST action taken; System is expected to experience a bug check (blue screen) 2. System ending state: A. Blue screen bug check B. Upon next boot, the Intel® RST driver will evaluate the storage subsystem 3. Recommended corrective action: A. Reinsert the slow media to its SATA port B. Reboot the system normally C. Let the Intel® RST driver handle any RST specific recovery tasks.

6.7.4 S.M.A.R.T. Events

S.M.A.R.T. event (event count threshold reached) on fast media:	1. Action taken (once the SMART event count threshold is reached): A. The fast media's block cache (BC) is placed in Write-Through mode (reduced performance) B. Pop-up message with notification of media errors with recommendation to disable Optane™ Memory and correct the problem (e.g. replace the fast media) 2. System ending state: A. Optane™ Memory remains enabled. B. Optane™ Memory placed in reduced performance mode (WT cache mode) 3. Recommended corrective action: A. Disable Intel® Optane™ Memory B. Power down the system C. Replace the Intel® Optane™ memory module D. Reboot and enable Intel® Optane™ Memory
S.M.A.R.T. event (event count threshold reached) on slow media:	1. Action taken (once the SMART event count threshold is reached):

	A. The fast media's block cache (BC) is placed in Write-Through mode (reduced performance) B. Pop-up tray icon notification of media errors with recommendation to disable Optane™ Memory and correct the problem (e.g. replace the fast media) 2. System ending state: A. Intel® Optane™ Memory remains enabled. B. Intel® Optane™ Memory placed in reduced performance mode (WT cache mode) 3. Recommended corrective action: A. Disable Optane™ Memory B. Power down the system C. Recover (or whatever duplication method you choose) your faulty system drive to a new drive and replace the faulty drive 4. Reboot and enable Intel® Optane™ Memory

6.7.5 Fast Media: Block Cache Errors

RST Metadata Read failure during boot:	1. Action taken by the RST SW: A. Optane™ volume is placed 'offline'. 2. System ending state: A. Optane™ volume offline. B. System unbootable 3. Recommended corrective action: A. Retry boot B. If system unbootable after several tries then the metadata is corrupted and unrecoverable
RST Metadata Write failure during shutdown	1. Action taken by the RST SW: A. Treat the same as a dirty shutdown B. RST evaluate storage subsystem at next boot 2. System ending state: A. System is bootable. B. System in normal state 3. Recommended corrective action: A. Continue power down/restart as normal

Read failure of cache frame with dirty data exceeds the preset threshold value (fast media failure could be imminent)	1. Action taken by the RST SW: A. The fast media's block cache (BC) is placed in Write Through mode (reduced performance) B. Tray icon notification of media Read errors with recommendation to disable Optane™ Memory and correct the problem (e.g. replace the fast media). C. RST evaluate storage subsystem at next boot 2. System ending state: A. BC in WT cache mode in reduced performance state. B. Optane™ Memory remains enabled 3. Recommended corrective action: A. Disable Optane™ Memory B. Power down the system C. Replace the Optane™ memory module D. Reboot and enable Optane™ Memory

6.8 Intel® Optane™ Memory Member Disk Failures Requiring Replacement of Disk

6.8.1 Slow Disk (SATA HDD, SSD, SSHD) Replacement

If your system encounters an unrecoverable catastrophic failure of the slow media (the OS system disk becomes inoperable), there is no recovery/repair available. Use the following recommended steps to get the system back up and running:

Data recovery:	No Intel® RST data recovery tools are available for a drive that has experienced mechanical or electrical failure and is considered 'inoperable'. There may be third party tools available for recovering data in this situation. Intel has no recommendations regarding third party data recovery tools.

Disassociate the Optane™ memory module.	Disassociation of the Optane™ module is required in order to be able to reuse it to enable Optane™ with the replacement slow media. **Note: Disassociation removes the Optane™ configuration information from the Optane™ module and deletes all data (data which is unusable cached data from the inoperable OS disk that is being replaced). Once this disassociation is completed the module can again be used to enable Optane™.** 1. Boot into your system BIOS 2. Enter the Intel® Rapid Storage Technology pre-OS UI ***(Note: consult your computer manufacturer for the location of this UI in your system BIOS)*** 3. The slow media (Intel® Optane™ memory module) should be displayed on the main page, A. Highlight it and hit <Enter> to take you to the "Reset to non-Optane" page B. Under 'Disk Actions' highlight the action 'Reset to non-Optane' and hit <Enter> C. Highlight 'Yes' and hit <Enter> 4. The disk is now available to be used to enable System Acceleration with Intel® Optane™ Memory
Reinstall your OS or recover it from a previous backup	At this point you can reinstall the operating system and rebuild your system: 1. Install the new slow media (SATA HDD, SSD, or SSHD) 2. Install the Windows 10 OS 3. Re-enable Intel® Optane™ Memory 4. Complete any customizations or application installs

6.8.2 Intel® Optane™ Memory Module Replacement

This section documents support for Intel® Optane™ memory module upgrades or replacements for systems that already have Intel® Optane™ Memory enabled.

Whether you are upgrading the memory module or replacing a faulty memory module, the following steps will apply:

Disable Optane™ Memory:	1. From Windows desktop, launch the Intel® Optane™ Memory UI • *or launch the Intel® RST UI* 2. Click the 'Setup' tab at the left • *or for the Intel® RST UI, click the Intel Optane™ Memory tab at the top* 3. Click Disable to start the disabling process 4. Confirm the action 5. Once the UI has completed all tasks for disabling Optane™ and any necessary file migration has completed (after the progress indicator reaches 100%), the [Restart] *or [Reboot]* button will be displayed 6. Click the [Restart] *or [Reboot]* button to reboot the system and complete the disablement process
Replace the Intel® Optane™ memory module:	1. Power down the computer 2. Open the computer and locate the old or faulty Optane™ memory module that you wish to replace **Note:** ***Consult your computer manufacturer for the location of the slot for the Optane™ memory module and instructions to remove and insert an M.2 storage module*** 3. Remove the old or faulty module and insert the new module. 4. Close the computer and power it on and boot into the Windows OS.
Re-enable Intel® Optane™ Memory:	1. From Windows desktop, launch the Intel® Optane™ Memory UI • *or launch the Intel® RST UI* 2. Click on the 'Setup' tab • *or for the Intel® RST UI, click the Intel Optane™ Memory tab at the top* 3. Click Enable to start the enabling process 4. Select the fast media to use for enabling and continue 5. Once the UI has completed all tasks for enabling Optane™ and any necessary file migration has completed (after the progress indicator reaches 100%), the [Restart] *or [Reboot]* button will be displayed 6. Click the [Restart] *or [Reboot]* button to reboot the system and complete the enablement process

6.9 Windows Recovery Environment Support

This section documents support for system recovery when Intel® Optane™ Memory is enabled on the system.

6.9.1 Setting Up the Windows Recovery Environment (WinRE)

All WinRE recovery methods are supported provided the following steps are taken prior to installing the Operating system:

6.9.1.1 OEM Factory Built Systems

Prior to building the system that the reference PC image will be used for imaging systems as they are built on the factory assembly line, update the .wim files located in the Windows OS distribution media or ISO by injecting the RST 15.5.0.1051 or later driver files. Using the built-in Windows DISM tool, mount and add the RST 15.5.0.1051 driver into the following image files:

1. Boot.wim
2. Install.wim
3. WinRE.wim (note that the WinRE image file is located within the Install.wim image and will have to be mounted while the Install.wim image is mounted)

Systems built with these file having the RST driver injected into them will be setup to support Intel® Optane™ Memory enabled systems in the event of a recovery condition that causes to system to boot to the WinRE menu.

6.9.1.2 Post-Factory Built Systems

For systems built outside of the factory using the default Windows setup, the in-box storage driver is used when the Windows Recovery Environment is configured during the Windows setup. The in-box driver does not support Intel® Optane™ Memory volumes. Thus when the system boots to the WinRE environment with the in-box storage driver, the Optane™ Memory volume will not be detectable. None of the recovery methods will be successful.

However, prior to experience a recovery condition, the WinRE.wim file can be updated with the RST 15.5 driver. Then if a recovery condition happens, the WinRE will be able to detect the Optane™ Memory volume. Use the following procedure to update the WinRE on a live system (while the Windows* OS is running):

Note: ***Not all recovery conditions will be successful 100% of the time, not even on non-Optane™ Memory enabled systems.***

Open a command prompt on your desktop, and run the command:

`>\reagentc.exe /info` (the output printed to screen will tell you where to find the WinRE.wim image file; the output is as follows):

Windows RE Location: \\?\GLOBALROOT\device\harddisk0\partition2\Recovery\WindowsRE

- ***disk0*** : this indicates that the Winre.wim image file is on disk 0 (use diskpart list disk)
- ***partition2*** : this indicates that WinRE.wim is on partition 2 of disk 0 (diskpart select disk 0, select partition 2)
- ***Recovery\WindowsRE*** : indicates the directory where the image file is located

So, on the second partition of the harddisk0 (also known as the "C:" drive, according to "Diskpart"), you will find a hidden "recovery" directory, with subdirectory "WindowsRE". Within here is the "winre.wim" file.

Now to inject the drivers into the image file, first, place all the drivers you wish to inject into an easily accessible directory (such as c:\ temp\drivers, in our example), and then run the following commands:

```
>\diskpart
>\select disk 0
>\select partition 2
>\assign letter=R
>\exit
>\md c:\temp\mount\winre
>\md c:\temp\drivers
    "copy RST 15.5 driver files to the c:\temp\drivers directory"
>\dism /mount-wim /WimFile:R:\recovery\WindowsRE\winre.wim /index:1
   /mountdir:c:\temp\mount\winre
>\dism /image:c:\temp\mount\winre /add-driver /driver:C:\temp\drivers /recurse
>\dism /unmount-wim /mountdir:c:\temp\mount\winre /commit
```

6.9.2 WinRE Menu Items

6.9.2.1 Troubleshoot → Advanced Options

Invoke the Windows Recovery Environment menu. Main page "Choose an option:", select 'Troubleshoot', then select 'Advanced options':

System Restore:	Before you can use this option, you must have created a 'restore point' while in Windows Control Panel. From Control Panel, select 'Recovery' and follow defaults to 'Configure System Restore'. Once 'restore point' is created you can use this option: 1. Click on the option; your system will reboot 2. Click on the account you wish to use 3. Enter the password and click [Continue] 4. Follow the default prompts 5. Select a backup and follow the prompts to complete the **restore**

System Image Recovery:	Before you can use this option, you must have created a 'backup' while in Windows Control Panel. From Control Panel, select 'Backup and Restore (Windows 7)', the select 'Setup backup', and in the left navigation click on 'Create a system image' and follow defaults. Once an image is created you can use this option: 1. Click on the option; your system will reboot 2. Click on the account you wish to use 3. Enter the password and click [Continue] 4. Follow the default prompts 5. Select a backup and follow the prompts to complete the **recovery**

6.9.2.2 Troubleshoot → Recover from a Drive

Invoke the Windows Recovery Environment from the bootable 'Recovery Drive':

Recovery Disk:	**Warning! Use this method as a last resort. The 'Fully clean the drive' option of this process will reformat your entire HDD and reinstall a fresh OS. Personal data is not reserved with this method of Window Recovery.** Before you can use this option, you must have created a 'Recovery Disk' while in Windows Control Panel. Connect a USB thumb drive, minimum 32GB (you may need larger based on your system). From Control Panel, select 'Recovery' and follow defaults to 'Create a recovery disk'. Once the bootable 'Recovery disk' is created you can use this option: 1. Insert the USB Recovery disk and boot your system. 2. Select the USB drive as the boot device and boot to the USB key 3. Click on your Keyboard Layout 4. Click on 'Troubleshoot' 5. Click on 'Recover from a drive' 6. Click on the option you wish to use a. Option 1: Just remove my files b. Option 2: Fully clean the drive 7. Follow the default prompts to complete the recovery process

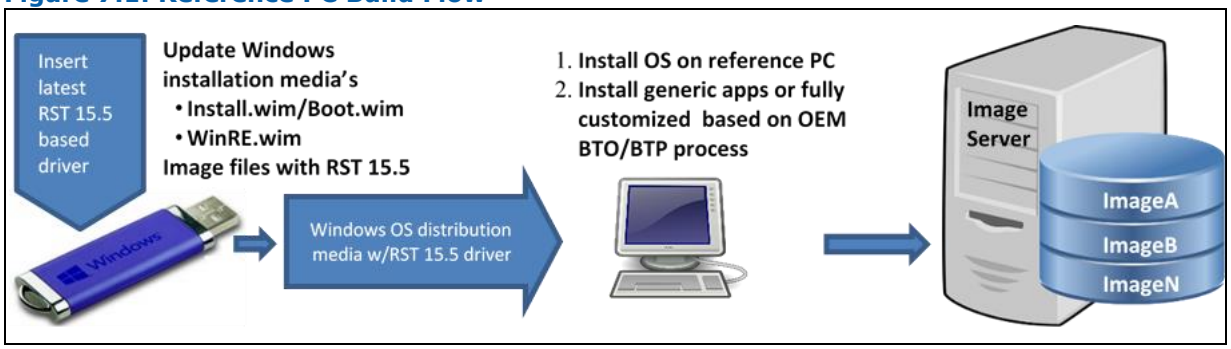
7 OEM HVM Factory Support for Intel® Optane™ System Acceleration

As there are many possible OEM factory processes, it is beyond the scope of this document to cover all possible OEM factory flows. There are two objectives for this chapter. One is to provide OEMs with an example of a possible flow based on what RST believes will provide the factory the fastest build process that is based upon the Microsoft provided manufacturing tools and processes. The other objective is to provide OEMs with all the requirements and considerations they need to take into account when building systems configured with Intel® Optane™ Memory.

7.1 Intel® Optane™ Pre-Factory Recommendations for the 'ReferencePC'

This section is meant to provide guidance on building the reference PC ('ReferencePC') in the pre-factory-assembly environment. This is where the OEM does their normal build process of installing the Windows OS and any applications required on a PC (ReferencePC) with the HW that represents the ordered PC that will be going through the factory assembly process (the CustomerPC). The OEM/ completes the build of the ReferencePC, extracts and stores the "gold" image of the install to a server ('Image Server') where the image can be accessed by the factory floor. At the proper point during the factory assembly/build process, the image is downloaded from the 'Image Server' onto the 'CustomerPC' and applied to the storage media (HDD, SSD, Optane™ drive, etc.). Below are recommended tasks to ensure that the final customer OS image will properly support System Acceleration with Optane™ Memory. This process can be used for BTO (Build-to-Order) or BTP (Build-to-Plan) HVM.

Figure 7.1: Reference PC Build Flow



7.1.1 Optane™ Recommendations for the 'ReferencePC'

7.1.1.1 Windows installation media update

1. Download the latest Intel® RST SW/driver package that supports Intel® Optane™ Memory
2. Obtain the desired Windows 10 OS distribution media from Microsoft
3. Follow the Microsoft process to add drivers to an image file (.wim). Add the Intel® RST driver to the following image files in the Windows 10 distribution media:
    - Boot.wim
    - Install.wim

- WinRE.wim

4. (Optional if you already have a WinPE bootable disk with the Intel® RST 15.5 or later driver installed). Create bootable WinPE USB with the Intel® RST 15.5 or later driver installed and copy the RSTCLI64.exe command line utility to a folder on the disk.

7.1.1.2 Windows OS Installation

1. Boot to the WinPE thumb drive and create the Optane™ volume that the OS will be installed to using the following command:

rstcli64 --OptaneMemory --enable --fast-drive [fastdrive_ id] --drive-to-accel [slowdrive_ id] --file-cache-offset

Example: *rstcli64 --OptaneMemory --enable --fast-drive 0-3-0-0 --drive-to-accel 0-0-1-0]* --file-cache-offset [LBA][1]

[LBA][1] (see important note at the end of this section)

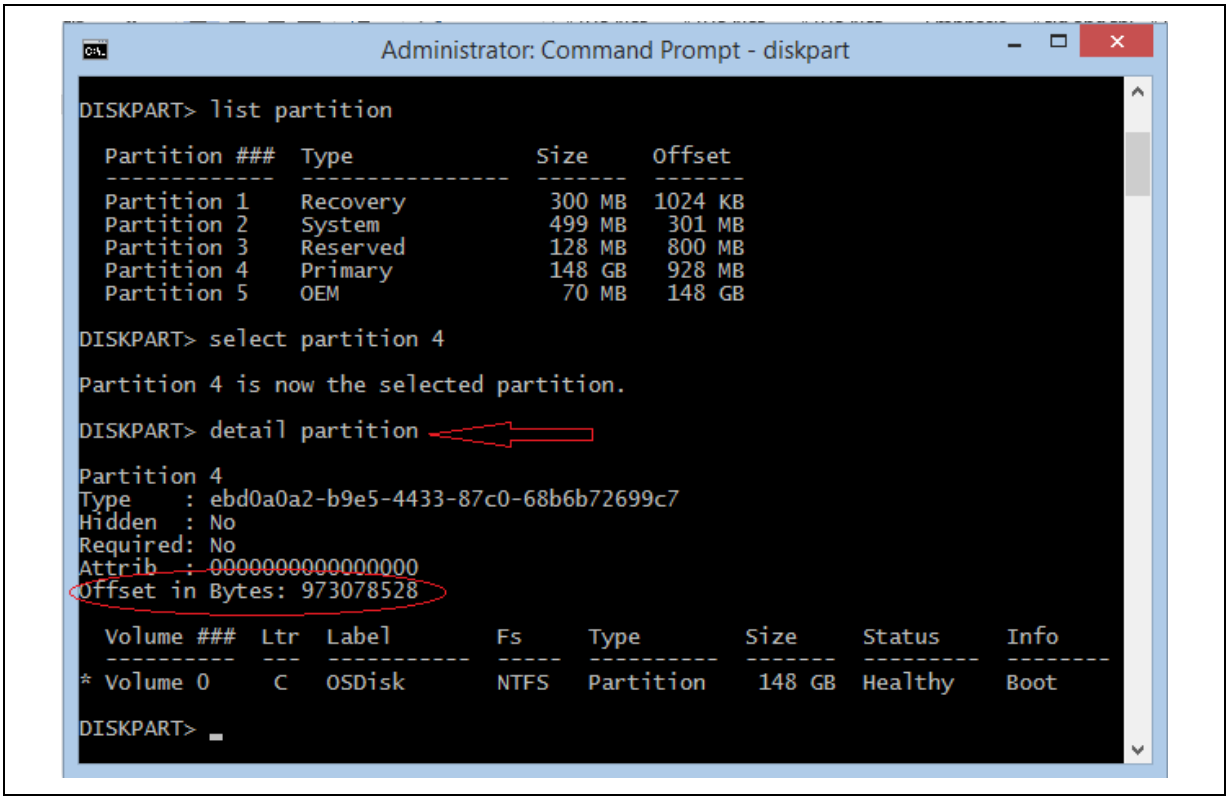
Note: If the fast PCIe NVMe Optane™ drive and the slow media to be accelerated are the only RST controlled drives in the system, then you can use the following simplified command:

rstcli64 --OptaneMemory --enable

The ***fast-drive*** and the ***drive-to-accel*** will be auto-detected if they are the only two drives in the system (the USB thumb drive is not controlled by RST so it will not be detected as a third drive)

2. Reboot and insert the updated Windows OS distribution media
3. Run the Windows installation setup program and complete the OS installation.
4. Complete installation and configuration of all apps and settings that are to be included in the 'Reference PC Image' (ReferencePC.wim)

Figure 7.2: Example of Partition Offset LBA in Windows Diskpart



Note: *This parameter is very important when enabling Optane™ Memory using Optane™ devices that are 32GB or larger when you using the RSTCLI64.exe tool (or RcfgSata.efi tool). 32GB or larger SKUs have a file cache region in addition to the block cache. This parameter designates the LBA of the slow disk that aligns with the beginning of the file cache to begin. This LBA is at the beginning of the Windows* primary (C:\) partition. You must determine the offset of the Windows primary partition. This number is usually given in Bytes if you use the 'Detail' command in 'Diskpart' (or MB if you use the 'list partition' command). In Diskpart select the disk (or Optane™ volume) containing the OS. Then select the Windows* partition, then type the following command:* ***'detail partition'****. The partition's offset will be given in 'Bytes'. Divide that number by 512 (an LBA is 512B) to determine the LBA of the offset. Example: (using info in Figure 6.2) if the offset is given as 973078528 bytes, then the offset LBA would be calculated as follows:*

973078528 ÷ 512 = 1900544 LBA. Thus the command to enable Optane™ Memory would be:

rstcli64 --OptaneMemory --enable --fast-drive 0-3-0-0 --drive-to-accel 0-0-1-0 --file-cache-offset 1900544

Or alternately using the 'list partition' command it can be calculated as follows:

(928 x 1024 x 1024) ÷ 512 = 1900544 LBA.

If this parameter is incorrect, the first time the system is booted into Windows after the install (e.g. during OOBE), in the background the RST diver will rebuild the Optane™ volume to place the file cache at the proper offset. This will have a slight negative impact to the system performance during the volume rebuild period.

Note: : Manufacturers that boot the customer computer that is being built into the Windows 10 OS during their factory build process after enabling Optane™ on the build system, can omit the file-

cache-offset parameter. This is because when you boot the customer computer that is being built into the Windows 10 OS the first time, the Intel® RST SW automatically repairs the file cache offset in the background if it is not specified or if it is specified incorrectly during the Optane™ enabling step. Although it is recommended that the offset parameter be specified correctly, it can be omitted with little impact to the Optane™ configuration in this manufacturing scenario.

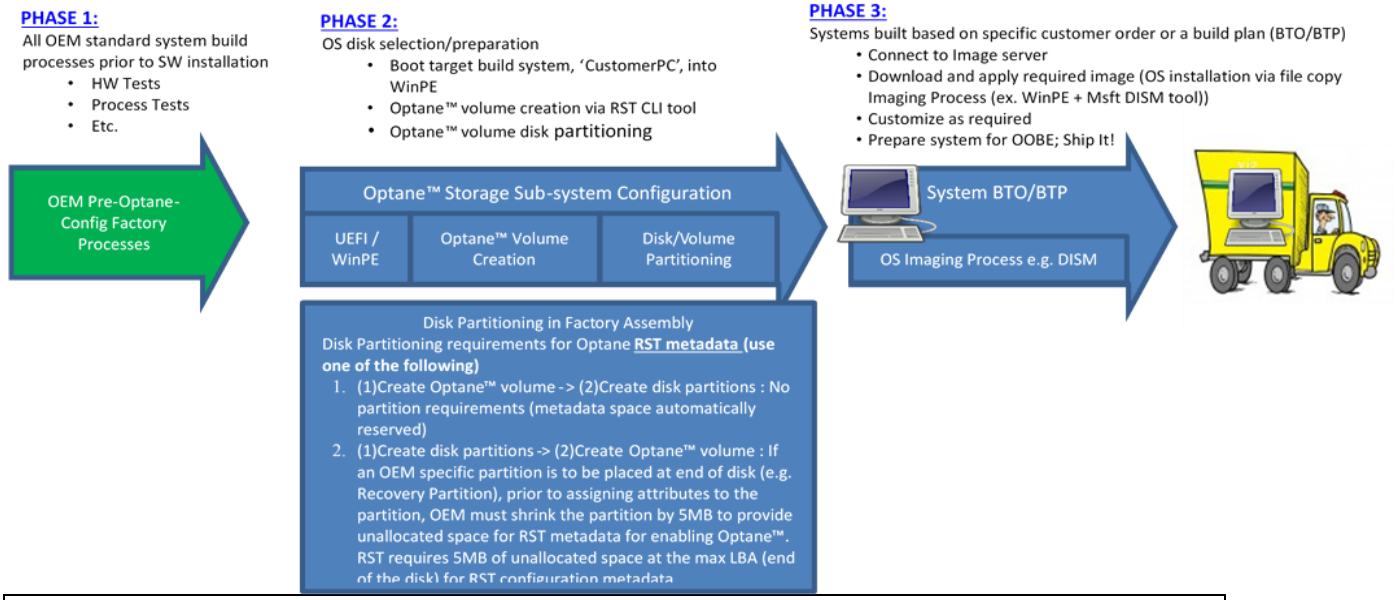
7.1.1.3 Extract 'Reference PC' "Gold" Image and Store for Later Use in the Factory Build of 'Customer PC'

1. Run Microsoft 'DISM" Imaging tool to capture gold 'Reference PC Image' (ReferencePC.wim)
2. Store the gold image on a source that will be available to the 'Customer PC' during the factory assembly build process (e.g. on a server, "Image Server", that is accessible via the network in the factory environment)
3. This completes the pre-factory 'Reference PC'/'ReferencePC image' build process

7.2 Intel® Optane™ Factory-Enabled System Build

This build process allows the OEM to ship the system to end-user fully configured and optimally accelerated at the time of the end-user OOBE. This allows for the best OOBE as the cache regions have already been populated with data when shipped from the factory.

Figure 7.3: Factory Build Flow



7.2.1 Requirements/Considerations

1. Ensure that there are no Windows system partitions on the Optane™ memory device. The Optane™ enabling process will delete all non-Windows OS partitions but will fail to enable if a Windows OS partition is detected. **(Note: all data on the Optane™ memory module will be deleted during the enable process)**
2. System must be in either Intel® RST, or Intel® RST Premium mode with the Intel® RST 15.5 OS driver installed and the Intel® RST 15.5 UEFI driver integrated in the system BIOS

7.2.2 PHASE 1: OEM Pre-Optane™-Config Factory Processes

This is the initial part of the OEM factory build process for the 'CustomerPC' build. There are no Optane™ requirements for this portion of the build process. These are processes specific to the OEM that the OEM would perform on all system configurations, not just Optane™. Any processes in this phase that require testing the final storage configuration would require that the system be in UEFI BIOS configuration, with the SATA controller in one of the Intel RST modes and the Optane™ device remapped.

7.2.3 PHASE 2: Optane™ Storage Sub-System Configuration

During this phase of the 'CustomerPC' build process, there are some Optane™-specific considerations:

1. Boot 'CustomerPC' to WinPE environment (ensure RST 15.5 driver is loaded and RSTCLI64 tool is on the WinPE boot environment or connected thumb drive).
2. Use the Microsoft 'Diskpart' tool to 'clean' the slow media prior to enabling Optane™. ***(This ensures fastest Optane™ volume creation with no data to migrate. This is more important with any Optane™ module that is larger than the 16GB SKU)***
3. As there is only the Optane™ device and the slow media attached to the system, use the following simplified command to enable Optane™:

   rstcli64 --OptaneMemory --enable --file-cache-offset xx

   For Optane™ devices larger than the 16GB SKU, ensure that this value for 'xx' is correctly calculated. See section Windows OS installation above for details on calculating this value.

   (The fast-drive and the drive-to-accel will be auto-detected if they are the only two drives in the system controlled by RST)

Note: Optional process (although not recommended) is to skip this step until a later point either in this phase or Phase 3 of the build process. If this step is skipped then in the next step the factory must reserve 5MB of unallocated space at the end of the slow media (max LBA) for RST metadata.

4. Partition the Optane™ volume (or slow media disk) using the recommended Microsoft* partitioning for Windows* 10. The following figure is the default partition layout for Windows* 10 UEFI/GPT based computers. **UEFI/GPT is required for Optane™**

Recovery – Windows Recovery Environment (WinRE) partition is where the Windows recovery tools are stored. This partition has a minimum size requirement of 300MB and must be large enough to place the WinRE.wim image and must be formatted NTFS. If the size is less than 500MB, there must be at least 50MB of free space, otherwise 320MB of free space is required for larger than 500MB partitions.

Note: Make sure the WinRE.wim image file has the RST 15.5 driver added.

a. System – ESP (EFI System Partition) must be a minimum of 100MB (this is the partition that the system boots to) and must be FAT32 formatted
b. MSR – (Microsoft Reserved Partition), beginning in Win10 the size is 16MB (receives no partition ID)
c. Windows – this is the OS partition and must be a minimum of 20GB (64-bit OS versions) and must be formatted NTFS.
d. OEM Recovery Partition -OEM-specific partitions (Utility partition) - OEMs that build systems with Recovery Image partitions usually place these partitions at the end of the disk. This partition is where OEM can copy the full Install.wim, this is recommended for Optane™ system, ***(must have the RST 15.5 driver added)*** as the recovery image. **If step 3 above was skipped then you are partitioning the slow media disk not the Optane™ volume. Prior to assigning attributes to the partition you must shrink the Recovery partition by 5MB to allow for RST metadata at the max LBA. This will allow for Intel® Optane™ Memory to be enabled later.** If step 3 was not skipped then you are partitioning the Optane™ Volume and the metadata space has already been reserved by the RST driver.

5. This completes phase 2 of the build flow (assuming you used the recommended flow and did not skip step 3), the 'CustomerPC' state is:
   a. Optane™ volume already enabled
   b. Optane™ volume properly partitioned for Optane™
   c. OEM Recovery partition large enough to store the 'ReferencePC.wim' image as Install.wim.

7.2.4 PHASE 3: OS and SW Download and Installation and OOBE Preparation

In this final phase of the factory build process, the factory process downloads the ReferencePC.wim (Install.wim), the WinRE.wim to the 'CustomerPC' and install the images. Also any additional custom applications not included in the 'ReferencePC.wim gold image can be downloaded and installed during this phase. Once all installations have completed, the factory process then prepares the 'CustomerPC' for the OOBE. The Optane™ recommended process should be similar to the following:

1. While still in the WinPE environment (no reboot is required once you start Phase 2), connect to the 'ImageServer' and download the ReferencePC.wim image to the OEM Recovery partition (rename it Install.wim)
2. Download the WinRE.wim to the WinRE partition
3. Apply the image to the Windows partition using the Microsoft* DISM tool
4. Use BCDBOOT to configure the system partition (ESP)
5. Register the location of the Recovery image (Install.wim) by using the REAgentC command in Windows
6. Register the location of the WinRE tools (WinRE.wim) by Using the REAgentC command
7. This completes Optane™ setup requirements
8. Continue with any additional configuration and OOBE preparation
9. This completes phase 3; 'CustomerPC' is ready to be shipped.

7.3 Intel® Optane™ Post-Factory-Enabled Builds

This build process allows the OEM to setup and fully configure their systems with all the Optane™ components installed but without the Optane™ volume enabled. The Optane™ acceleration enabling portion is to be accomplished during the end-user initial boot of the system OOBE. Via scripting, RSTCLI64 will enable the Optane™ volume during the initial customer boot. This is not the recommended build process as the initial OOBE is not accelerated.

7.3.1 Requirements/Considerations

1. 5MB of unallocated space must be left at the end of the HDD/SSD slow disk. OEM should not place any unmovable or protected partitions within this space
2. If required to place partition at the end of the disk, OEM must leave an offset of (Max LBA – 5MB) for the RST metadata (this can be done by shrinking the partition by 5MB prior to assigning attributes/GUID type to the partition)
3. Ensure no Windows OS partitions are on the Optane™ device
4. System must be in Intel RST or Intel RST Premium mode with RST 15.5.0.1051 or later driver installed and RST 15.5.0.2875 or later UEFI driver integrated into the system BIOS

7.3.2 Example Build Flow

1. OEM/System Integrator begins their build process
2. OEM/System Integrator can run any diagnostic and test programs as required on the system
3. Partition the SATA HDD/SSD as required for the system build
    a) Ensure no partition is located in LBA offset of Max LBA – 5MB. This space must be unallocated
4. Apply Windows image to the HDD/SSD prepared for OOBE
    a) The Intel® RST RSTCLI64 tool must be inserted in the Windows image with script to enable Optane™
    b) Inject RST 15.5 driver in all .wim files (boot.wim, install.wim, winre.wim)
5. Complete the build process and prepare the system for OOBE shipment
    a) Post-factory the OOBE runs script using RSTCLI64 to enable System Acceleration with Intel® Optane™ technology

7.4 OEM Application Pinning in the Factory Flow

This section covers OEM options for enabling application pinning in the factory process. OEM Pinning process is of type 'Soft Pinning'. File/application pinning falls in one of two types, a) Hard pinned files and b) Soft pinned files:

Hard Pinned files: These files are files pinned by the end user. They are pinned into the file cache region of the Intel® Optane™ memory module and can only be removed from the cache by the end user 'Unpinning' them. They are not influenced by the Intel® RST caching policies. So even if they are never used they will never be automatically evicted from the cache by the Intel® RST caching policies.

Soft Pinned files: These files are files pinned by OEM/ISV. They are also pinned into the file cache region of the cache. Soft pinned files are influenced by the Intel® RST caching policies. As such they can be automatically evicted from the file cache if they become stale/cold from non-use or inactivity.

7.4.1 Manufacturing Considerations for Pinning

Process to accomplish pinning:

JSON Configuration Files:
JSON configuration files (e.g. "OEM_App-A.json") will be used in the factory environment for OEMs to define what files associated with an application will be pinned into the file cache as specified by the OEM. Each JSON file will define an application to be pinned and include a list of files that are important[1] to the functionality of that application which should be pinned along with the app's executable file.

Note1: ***Please contact your Application SW vendor for suggestion/guideline on which files are deemed as important enough to include in pinning to enhance the performance of that application.***

JSON File Naming:
Each JSON file must be named according to the application's executable file name. If file name is OEMApp01.exe, then the JSON file for that application must be named OEMApp01.json.

JSON File Naming Example

Application to be Pinned Name	Application Executable	JSON File Name
Chrome	Chrome.exe	chrome.json
Windows Photo Viewer	ImagingDevices.exe	ImagingDevices.JSON
Windows Media Player	WmPlayer.exe	wmplayer.JSON

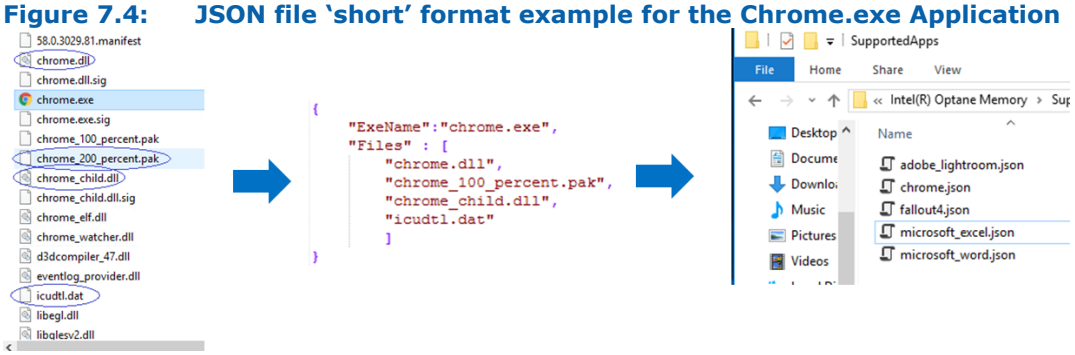
JSON File Location:
JSON files must be placed in a specific location of the computer's directory structure. OEM must place the files in the following directory:

%CommonProgramFiles%\Intel\Intel Optane Memory\SupportedApps
E.g. most commonly C:\Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps

***Note**: Normally this directory is created by the Intel® RST installers (SetupRST.exe and SetupOptaneMemory.exe). However, depending upon specific OEM factory build flow, this directory may have to be manually created by the OEM.*

JSON File Format:
The following screenshot shows one of the formats of the JSON file. The following 'short' example is for those applications where all of the supporting files to be pinned are located in the same directory as the application's executable file.

Figure 7.4: JSON file 'short' format example for the Chrome.exe Application



Note: Intel will provide JSON file schema for manufacturers to use in creating their JSON files.

7.4.2 Using the JSON files in the Factory Flow

It is recommended that the JSON files be integrated into the factory flow process post customer specific application install prior to setting the computer to the OOBE configuration.

Typical BTO flow:

1. OEM-specific pre-SW installation HW verification
2. Enable/create Optane™ volume, partition disk, and apply gold OS image (pre-built)
3. Install BTO customer-specific applications
4. Set computer to OOBE-ready

Example BTO w/pinned applications flow:

1. OEM-specific pre-SW installation HW verification
2. Enable/create Optane™ volume, partition disk, and apply gold OS image (pre-built)
3. Install BTO customer-specific applications in the following order:
    a. Install and configure custom applications
    b. Copy the required JSON files into the correct folder in the file system's directory tree (a JSON file for each application to be pinned is required)
    c. Integrate the Intel® RST OEM Pinning process into the manufacturing build process at this point to setup the pre-Pinning task (this prepares the computer to
        i. Install BTO customer-specific applications
        ii. Pin customer-specific applications
4. Prep the computer for OOBE and power down the computer in OOBE-ready state

7.4.3 OEM Pinning Methods

There are several possible methods available to accomplish the pinning process in the factory. The method selected to be used by a specific computer manufacturer/builder is determined by their build process and their customers' requirements. More than one method might be used in build process.

7.4.3.1 Pinning via Unattend.xml file

Unattend.xml file is generated by the Windows installation process. Pinning can be accomplished during the OS installation by modifying the Unattend.xml file. Figure 7.5 below shows an example of the format required to be followed to modify the Unattend.xml to include OEM Application Pinning during the OS installation.

Note: *If the factory build process does not require any generic applications to be pinned (these would be applications that would ship with every computer regardless of specific customer order that the computer manufacturer/builder would deem advantageous to pin), then modifying the Unattend.xml file may not be feasible.*

Figure 7.5 Example of 'Pinning' section of the Unattend.xml

```
<SynchronousCommand wcm:action="add">
    <CommandLine>rundll32.exe C:\Windows\IAStorAfsService\Optane.dll,OptanePinOemApplication path_to_OEM_app.exe</CommandLine>
    <Order>23</Order>
    <Description>Pin OEM application</Description>
</SynchronousCommand>
```

- Ensure that this section is placed near the end of the Unattend.xml file after the customer applications are installed
- Replace OEM_app.exe with the application's executable (e.g. path_to_Chrome.exe)
- Replace <Description> text with actual application (e.g. "Pin Chrome")
- Order of each app to be pinned is irrelevant; they can be pinned in any order

Pinning via the Unattend.xml is used when OEM wishes to define and pin applications during the OS Installation/Setup. OEM may choose this pinning option when there are default applications that will be pinned on all builds regardless of any order customization.

Also can be used when final build image "Gold Image" is the image that all builds will receive. E.g. an IT build that every system will receive.

7.4.3.2 Pinning via Windows Registry modification of the Install.wim

OEM may choose to use this method similar to the reasons for choosing the Unattend.xml method. OEM will modify the Install.wim to include the JSON files and the registry modifications to pin the application on first login by end-user. The build flow is as follows:

1. Mount the Install.wim of the gold image
2. Create directory for JSON files
   (%MountDir%\Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps)
3. Copy JSON files to SupportedApps directory
4. Load the registry of the install.wim image
5. Modify the registry of install.wim; use .reg file or edit directly through script to add registry entry for each application to be pinned
6. Unload the install.wim registry
7. Unmount and commit the changes to the Install image

Example of using a batch file and a reg file to accomplish the above flow for Excel, Internet Explorer, Windows Photo Viewer, and Power Point:

Batch file:

```
Dism /Mount-Image /ImageFile:E:\WindowsPart2.wim /index:1 /MountDir:C:\Temp\Mount\Windows

md C:\Temp\Mount\Windows\Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps

xcopy C:\Temp\JSON\*.* C:\Temp\Mount\Windows\Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps

reg load HKLM\MyKey2 C:\Temp\Mount\Windows\Windows\System32\config\software

C:\Temp\Mount\OEMPinRunOncePinned.reg

reg unload HKLM\MyKey2

Dism /unmount-image /MountDir:C:\Temp\Mount\Windows /commit
```

Reg file (e.g. OEMPinRunOncePinned.reg):

```
Windows Registry Editor Version 5.00


[HKEY_LOCAL_MACHINE\MyKey2\Software\Microsoft\Windows\CurrentVersion\RunOnce]

"Excel"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication C:\\Program Files\\Microsoft Office\\Office15\\excel.exe"

"Iexplore"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication C:\\Program Files\\Internet Explorer\\iexplore.exe"

"ImagingDevices"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication C:\\Program Files\\Windows Photo Viewer\\ImagingDevices.exe"

"PowerPnt"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication C:\\Program Files\\Microsoft Office\\Office15\\powerpnt.exe"
```

Once this has been done the OEM applications will be pinned upon first boot after the image has been applied to the build system.

7.4.3.3 Pinning via Windows Registry modification

OEM may choose to use this method when the applications to be pinned are installed after the general OS build (common for BTO factory flow).

Example build flow is as follows:

1. Complete installation of all applications that are to be pinned
2. Create the directory for JSON files
   (C:\Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps)
3. Create .reg file (see example below: *OEMPinRunOnce.reg*) or use script to add registry keys
4. Create JSON files for each of the applications to be pinned

5. Copy JSON files to the SupportedApps directory
6. Modify the registry; use .reg file or edit directly through script to add registry entry for each application to be pinned
7. Prepare the computer for OOBE
8. Shutdown the computer; *the computer is now setup to run the pinning process on first time user boots into the OS.*

Example of using a batch file and a reg file to accomplish the above flow for Excel, Internet Explorer, Windows Photo Viewer, and Power Point:

Batch file (**AppPnForOOBE.bat**):

```
Rem You must first create the directory where the App pnning requires the JSON files

Rem be located in order to pin the files associated with the applications

md C:\"Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps"

xcopy C:\Temp\JSON\*.* C:\"Program Files\Common Files\Intel\Intel Optane Memory\SupportedApps\"

C:\Temp\OEMPinRunOnce.reg
```

Reg file (e.g. **OEMPinRunOnce.reg**):

```
Windows Registry Editor Version 5.00


[HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\RunOnce]

"Excel"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication C:\\Program Files\\Microsoft Office\\Office15\\excel.exe"

"Iexplore"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication
C:\\Program Files\\Internet Explorer\\iexplore.exe"

"ImagingDevices"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication
C:\\Program Files\\Windows Photo Viewer\\ImagingDevices.exe"

"PowerPnt"="C:\\windows\\system32\\rundll32.exe C:\\Windows\\IAStorAfsService\\Optane.dll,OptanePinOemApplication
C:\\Program Files\\Microsoft Office\\Office15\\powerpnt.exe"
```

Once this has been done the OEM applications to be pinned will be pinned upon first boot of the computer by the end-sure.

§ §

8 Intel Rapid Storage Technology for PCIe NVMe Storage Devices

8.1 OEM System BIOS Requirements

System BIOS	System BIOS must enable remapping for the device to be supported by RST(consult the Intel BIOS Writers Guide/Specification for your platform)

8.2 General Requirements

Hardware	Only enabled on specific SKUs of the Skylake platforms (see 1st section above in this chapter). SATA/PCIe Remapping must be enabled on the platform. (Refer to the External Design Specification for the respective platforms: Doc #545659: *Skylake Platform Controller Hub (SKL PCH) External Design Specification – Volume 1 of 2* • Tables 1-2, 1-3 • Chapter 3 • Chapter 25, Table 25-2 Doc #546717: *Skylake H Platform Controller Hub (SKL PCH-H) External Design Specification – Volume 1 of 2* • Tables 1-2, 1-4, 1-5, and 1-7 • Chapter 3 • Chapter 25, Table 25-3).
Operating System	All x64 bit supported Operating Systems for this release
Supported Devices	NVMe PCIe Devices
SATA Mode	RAID Mode Only
Specification Support	Will adhere to the NVMe Specification The driver will implement MSI-X vectors for PCH AHCI devices and remapped PCIe SSDs

8.3 Feature Limitations

Intel® Rapid Storage Technology for PCIe NVMe Storage Devices has the following feature limitations:

- No support for:
  - Legacy AHCI DEVSLP
  - RTD3

  - Hot Plug
  - InstantGo*
- Supports: maximum of 3 ports can be remapped using x2 or x4 lanes
- If used in a RAID volume, all member devices must be on the same bus type

8.4 PCIe NVMe Device Usage Model

Intel® Rapid Storage Technology NVMe PCIe Device storage is limited to the following usages:

- As a single Pass-Thru Bootable device with PCH SATA controller in RAID mode.
- As a cache device for Intel® Smart Response Technology with the PCH SATA controller in RAID mode.
- As a member disk in a RAID volume with the PCH SATA controller in RAID mode (all member devices have to be on the same bus type).
- As a spare disk for a RAID volume (has to be on the same bus type as the RAID member devices)

8.5 Intel® RST for PCIe NVMe Storage Use cases

Prerequisite:

- A remapping-enabled PCH chipset.
- The platform BIOS must implement the remapping logic and turn it on.
- The PCIe Storage device must be attached to a remapable PCIe slot or PCIe M.2 connector.
  - Only **x2** and **x4** lane support for Skylake and later platforms
  - Only **x1** and **x2** lane support for pre-Skylake platforms that support re-mapping
- The PCIe Storage device must be NVMe-controller based.
- System must be in RAID mode.
- The platform must include the Intel® Rapid Storage Technology UEFI driver*.

If any of the above conditions are not met, the PCIe NVMe SSD will not be recognized by the RST driver.

*No Legacy OROM support

Use Cases:

Use Case	Description	Boot Support
Pass-through Device: Using NVMe PCIe Storage device as a pass-through device	When all the prerequisites referred above are met, the PCIe Storage device can be used as a pass-through device. It can be a boot device or a data device. Configurations: • Up to 3 pass-through disks PCH SKU dependent)	RAID

SRT Cache Device: Turn on SRT, using NVMe PCIe Storage device to accelerate SATA HDD, RAID volume	When all the prerequisites referred above are met and all conditions to enable SRT are met, SRT can be turned on to use the PCIe Storage device as the cache device to accelerate the SATA HDD. *Refer to section on SRT	RAID
Extra Space Volume: Turn on SRT, using NVMe PCIe Storage device to accelerate SATA HDD	When all the prerequisites referred above are met and all conditions to enable SRT are met, SRT can be turned on to use the PCIe Storage device as the cache device to accelerate the SATA HDD. Any additional space on the PCIe device not being used for cache can be used disk volume. *Refer to section on SRT	RAID
RAID Volume: Creation of RAID volumes using NVMe PCIe storage devices	When all the prerequisites referred above are met the PCIe devices can be used to create RAID volumes. Only supported RAID levels are 0, 1, and 5. The actual RAID level support will be dependent upon your platform's specific SKU and configuration; consult your platform documentation Configurations: • 3-disk RAID volume (RAID 0 or 5) • 2-disk RAID volume (RAID 0 or 1) + Spare • 2-disk RAID volume (RAID 0 or 1) + Single Disk	RAID

8.6 Intel® Rapid Storage Technology UEFI Compliance Utility for PCIe Storage

Beginning with Intel® RST UEFI version 13.0, the RcmpSata.efi utility includes the ability to test Intel® RST PCIe conformance for OEMs and ODMs in the UEFI shell (RcmpSata.efi) or DOS environment (RcmpSata.exe).

**RcmpSata utility is also available in earlier releases for Legacy OROM and UEFI compliance testing in the Pre-OS environment (see section 3.5.2.4).*

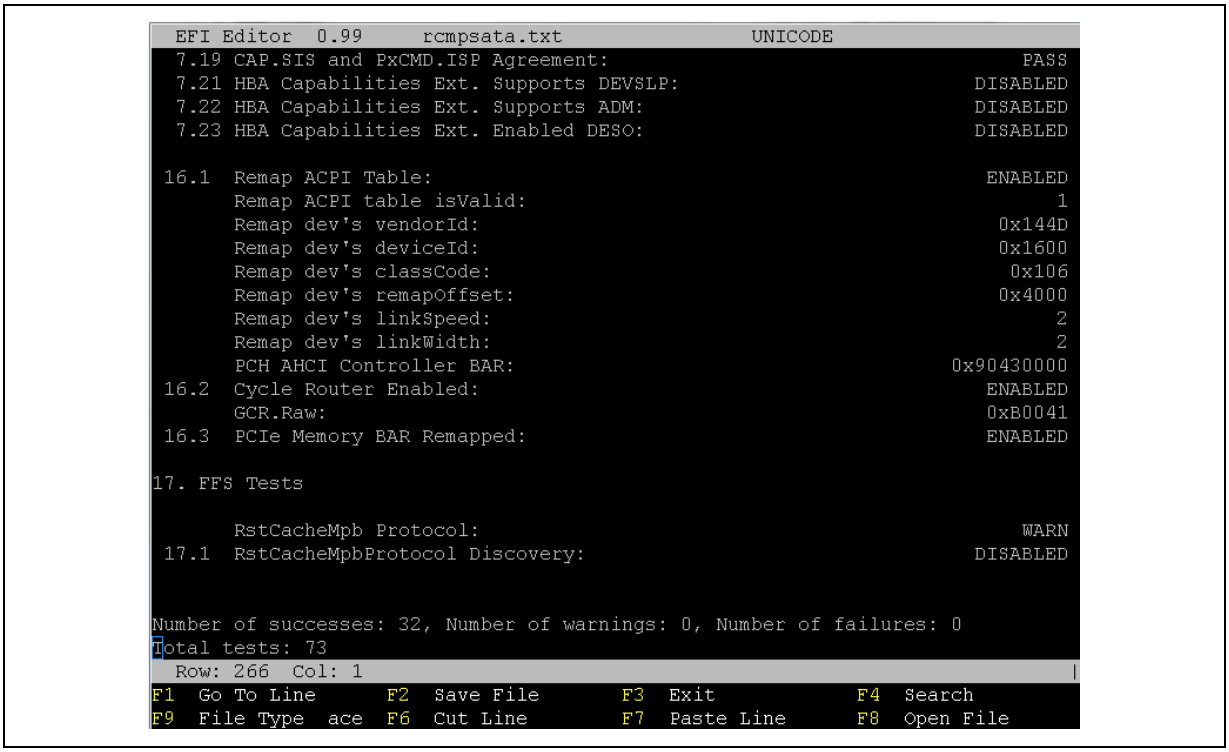
With the RcmpSata.efi utility downloaded to a Fat32 formatted USB drive attached to the platform, the following syntax can be used in the UEFI shell to download compliance data to a text file for viewing in a text editor (where '#' is the file system number of the USB drive shown when booting to the UEFI shell):

Fs#:> rcmpsata.efi > rcmpsata.txt

You may scroll through the text file in the UEFI shell by typing the following command:

Fs#:>edit rcmpsata.txt

The final test results are displayed at the end. Test Section 16 will confirm "remap" for PCIe is enabled/disabled for debugging issues in the Pre-OS environment.



§ §

9 Intel Rapid Storage Technology for PCIe AHCI Storage Devices

Beginning with the Intel® RST 13.0 Release version, PCIe storage devices are supported on the following SKUs:

PCH SKU	PCH Family	Segment
H97	LPT-H	Refresh Desktop
Z97	LPT-H	Refresh Desktop

9.1 OEM System BIOS Requirements

System BIOS	System BIOS must enable remapping for the device to be supported by RST(consult the Intel BIOS Writers Guide for your platform)

9.2 General Requirements

Hardware	Only enabled on specific SKUs of the Haswell Refresh platforms. Remapping must be enabled on the platform. (Refer the Product Design guide for the respective platforms).
Operating System	All x64 bit supported Operating Systems for this release
Supported Devices	AHCI PCIe Devices
SATA Mode	RAID Only

9.3 Warnings

9.3.1 Features Limitations

Intel® Rapid Storage Technology 13.1 for PCIe Storage Devices will not support the following features:

- Legacy AHCI DEVSLP
- L1.2 Support for PCIe devices is not supported until 13.5 release
- RTD3

- Hot Plug
- InstantGo*

9.3.2 PCIe Device Usage Model

Intel® Rapid Storage Technology PCIe Device storage is limited to the following usages:
- As a single Pass-Thru Bootable device with controller in RAID mode.
- As a cache device for Intel® Smart Response Technology with controller in RAID mode.

9.4 Intel Rapid Storage Technology for PCIe Storage Use cases

Prerequisite:
- The chipset must have the remapping hardware.
- The BIOS must implement the remapping logic and turn it on.
- The PCIe Storage device must be attached to remapable PCIe slot or PCIe M.2 connector.
- The PCIe Storage device must be AHCI-controller based.
- System must be in RAID mode.
- The system BIOS must include the Intel® Rapid Storage Technology UEFI driver*.

If any of the above conditions are not met, the PCIe SSD will not be recognized by the RST driver.
*No Legacy OROM support
Reference documents: For detailed use case information refer to RST13.0 Use case document ID 539242 published on CDI

Use Cases:

Use Case	Description	Boot Support
Using PCIe Storage device as a pass-through device	When all the prerequisites referred above are met, the PCIe Storage device can be used as a pass-through device. It can be a boot device or a data device.	RAID
Turn on SRT, using PCIe Storage device to accelerate SATA HDD	When all the prerequisites referred above are met and all conditions to enable SRT are met, SRT can be turned on to use the PCIe Storage device as the cache device to accelerate the SATA HDD. *Refer to section on SRT	RAID

9.5 Intel® Customer Reference Board BIOS Settings

The following settings can be used on Intel® Customer Reference Boards (CRB) that support remapping. Consult with your BIOS vendor for compatible settings for Intel® RST PCIe storage.

BIOS Settings using the EDISTO BEACH FAB 4 SKU-2 SDIO WIFI card with PCIe SSD connected

1. Enter into Bios set-up and choose the below settings:
    a. Intel Advanced Menu->PCI Subsystem Settings->Install Ext OpRom Before BIOS = "Ext PCIe Both Storage and Other OpRom"
    b. Intel Advanced Menu -> SA Configuration->X2APICOPTOUT = Disabled
    c. Intel Advanced Menu -> SA Configuration-> VT-d = disabled
    d. Intel Advanced Menu-> PCH-IO Configuration->PCI Express Configuration->PCI Express Root Port 6->PCIe Speed = Gen1
    e. Intel Advanced Menu->PCH-IO Configuration->SATA Configuration->SATA Mode Selection = RAID
    f. Intel Advanced Menu->PCH-IO Configuration-> SATA Configuration ->PCIe Nand Configuration = Enabled
    g. Intel Advanced Menu->PCH-IO Configuration->PCI Express Configuration->PCIe Nand Port Selection = Port 6
    h. Intel Advanced Menu->PCH-IO Configuration->SATA Configuration->EFI RAID = Enabled
    i. Boot Maintenance Manger Menu->boot configuration Menu->CSM control = Always On
2. Restart the SUT to reflect the setting in BIOS
3. Enter into Bios and navigate to TVP Device manager-> Check for the PCIe NAND SSD

9.6 Intel® Rapid Storage Technology UEFI Compliance Utility for PCIe Storage

Beginning with Intel® RST UEFI version 13.0, the RcmpSata.efi utility includes the ability to test Intel® RST PCIe conformance for OEMs and ODMs in the UEFI shell (RcmpSata.efi) or DOS environment (RcmpSata.exe).

**RcmpSata utility is also available in earlier releases for Legacy OROM and UEFI compliance testing in the Pre-OS environment (see section 3.5.2.4).*

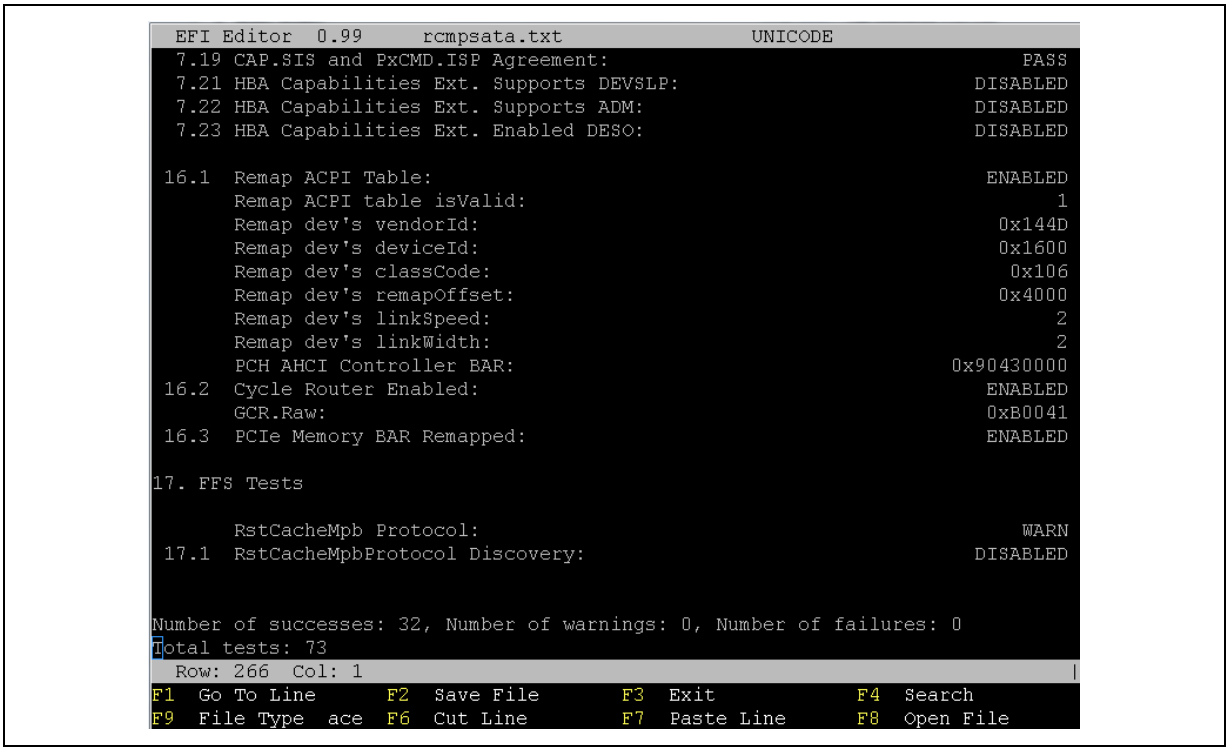
With the RcmpSata.efi utility downloaded to a Fat32 formatted USB drive attached to the platform, the following syntax can be used in the UEFI shell to download compliance data to a text file for viewing in a text editor (where '#' is the file system number of the USB drive shown when booting to the UEFI shell):

Fs#:> rcmpsata.efi > rcmpsata.txt

You may scroll through the text file in the UEFI shell by typing the following command:

Fs#:>edit rcmpsata.txt

The final test results are displayed at the end. Test Section 16 will confirm "remap" for PCIe is enabled/disabled for debugging issues in the Pre-OS environment.


EFI Editor 0.99 rcmpsata.txt UNICODE
7.19 CAP.SIS and PxCMD.ISP Agreement: PASS
7.21 HBA Capabilities Ext. Supports DEVSLP: DISABLED
7.22 HBA Capabilities Ext. Supports ADM: DISABLED
7.23 HBA Capabilities Ext. Enabled DESO: DISABLED
16.1 Remap ACPI Table: ENABLED
Remap ACPI table isValid: 1
Remap dev's vendorId: 0x144D
Remap dev's deviceId: 0x1600
Remap dev's classCode: 0x106
Remap dev's remapOffset: 0x4000
Remap dev's linkSpeed: 2
Remap dev's linkWidth: 2
PCH AHCI Controller BAR: 0x90430000
16.2 Cycle Router Enabled: ENABLED
GCR.Raw: 0xB0041
16.3 PCIe Memory BAR Remapped: ENABLED
17. FFS Tests
RstCacheMpb Protocol: WARN
17.1 RstCacheMpbProtocol Discovery: DISABLED
Number of successes: 32, Number of warnings: 0, Number of failures: 0
Total tests: 73
Row: 266 Col: 1
F1 Go To Line F2 Save File F3 Exit F4 Search
F9 File Type ace F6 Cut Line F7 Paste Line F8 Open File

§ §

10 How to Enable the Platform for Intel® RST Support of BIOS Fast Boot

Beginning with the Intel® RST 12.0 Release version, Intel® RST implements pre-OS UEFI driver and Windows runtime driver support for the platform BIOS Fast Boot specification.

10.1.1 OEM System BIOS Vendors' Requirements

This section covers what the OEM's system BIOS vendor must provide in order for the Intel® RST Fast Boot BIOS implementation to be enabled on the platform.

System BIOS

The Intel® RST UEFI driver requires the following BIOS components:

- 2KB of non-volatile UEFI variable storage with access from runtime and as a boot service
- Access to UEFI Hand-off Block Hand-off Info Table (PHIT HOB) to determine boot mode
  - BOOT_WITH_FULL_CONFIGURATION = Fast-Boot disabled
  - BOOT_WITH_MINIMAL_CONFIGURATION = Fast-Boot enabled
  - BOOT_ON_S4_RESUME = Fast-Boot enabled

10.1.2 Supported System Configurations

This section covers the system configurations that are required for the Intel® RST BIOS Fast Boot implementation to support the platform's BIOS Fast Boot specification.

System BIOS

SATA controller must be set to RAID mode

HW Configuration

Any of the following configurations are supported:

1. Platform configured with a single pass-through SSD or PCIe storage device (system boot drive)
2. Platform configured with a single SSD/mSATA + 2 HHD's as a RAID 1 (system boot drive), Accelerated with Intel® Smart Response Technology*
3. Platform configured with a single SSD/mSATA + single HDD (system boot drive), Accelerated with Intel® Smart Response Technology*

Note: ATAPI devices do not affect BIOS Fast Boot specification

*Fast Boot will not be supported on configurations where the PCIe Storage device is used as the cache disk.

Operating System

- Windows 8.1 64
- Windows 8 64
- Windows 7 64

§ §

11 Creating a RAID Volume

RAID volumes can be created three different ways. The method most widely used by end-users is to use the Intel Rapid Storage Technology UI in Windows*. The second method to create a RAID volume is to use the Intel Rapid Storage Technology option ROM user interface (or the Intel® RST pre-OS UEFI HII UI). The third way, used by OEMs only, is using the pre-OS RCfgSata or Windows (including WinPE) RSTCLI 32/64 utilities.

11.1 Minimum Requirements

1. BIOS: Requires a BIOS with the Intel® RST preOS UEFI driver or OptionROM integrated
2. Number of disks required for the desired RAID level
    a. RAID 0- 2 disks minimum (6 max)
    b. RAID 1- 2 disks only
    c. RAID 5- 3 disks minimum (6 max)
    d. RAID 10- 4 disks only
3. 5MB unallocated space at the end (max LBA) of the disk for RST metadata

11.2 Feature Limitations

Table 10-1 Disk, volume, and array maximum limitations

Configuration	Maximum
Number of arrays per system	4
Number of volumes per system	4
Number of volumes per array	2
Number of disks per volume/array	6
Number of disks per system	8

11.3 Using Intel® Rapid Storage Technology UI

1. Run the Intel Rapid Storage Technology UI from the following Start menu link within Windows:
    Pre-Windows* 8
    Start→Programs→Intel Control Center (optional) ->Intel® Rapid Storage Technology→Intel Rapid Storage Technology UI
    Windows* 8 and Newer

The UI is not added to the Start Window upon installation and must be added manually
File Explorer->Local Disk(C)->Program Files->Intel->Intel® Rapid Storage Technology->IAStorUI.exe (right click and pin to start)

2. Based on the available hardware and your computer's configuration, you may be able to create a volume by selecting the 'easy to use' options such as 'Protect data' under 'Status', or by selecting a volume type under 'Create'. Based on the number of non-RAID disks available to you and the size of the disks the user will only be able to see the possible volume creation options... (e.g. if you have only two disks ...you can only see options to create RAID 0, RAID1 and Recovery(Intel® RRT) ; if you have three disks, you can only see options for creating RAID 0, RAID 1, RAID5 and Recovery)

Note: To create a volume the user must be in admin mode and the system must be in RAID Ready mode with two or more hard disks connected to it

3. Instructions to create a volume by selecting volume type under 'Create'

    a. After selecting the volume type to create, click on 'Next'

    b. Now configure the volume by providing the volume name, selecting the hard disks to be part of the volume and strip size if applicable

    NOTE: When configuring a volume, the application will only list the disks that meet the min requirements to be part of the volume. Based on the first disk selected or the order of selection, some disks may become grayed out if one or more requirements are not met. Changing the order of selection generally helps re-enable disks that were grayed out. For Ex: If the first selection is a system disk, only disks that are of equal or greater size will be presented for selection and other remains grayed out. For more information on disk requirements refer 'creating a volume' under help file in the UI.

    c. Once the disks are selected for volume creation, the user will presented with option, if you want preserve data on which selected disk. Click on 'Next' and select the 'Create Volume' button.

4. After the RAID volume is created, you will be shown a dialog box stating that the RAID volume was successfully created and you will you will need to use Windows Disk Management or other third-party software to create a partition within the RAID volume and format the partition. Click OK to close this dialog box.

5. After formatting the partition, you may begin to copy files to, or install software on, the RAID volume.

11.4 Using Intel® Rapid Storage Technology Legacy Option ROM User Interface

1. Upon re-boot, you will see the option ROM status message on the screen – press CTRL-I to enter the Intel Rapid Storage Technology option ROM user interface.

2. In the Main Menu, select option #1 'Create RAID Volume'. Enter the name you want to use for the RAID volume, then press Enter.

3. Select the RAID level by using the arrow keys, then press Enter.

4. Press Enter to select the disks to be used by the array that the volume will be created on. Press Enter when done.

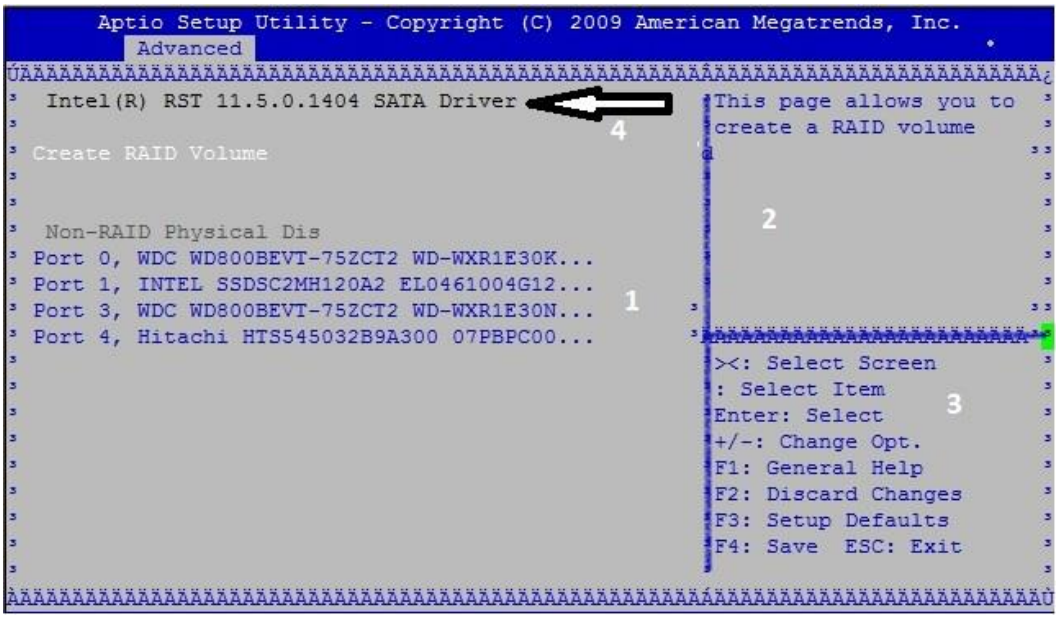
5. Select the strip size (128 KB is the default for RAID 0) by using the arrow keys, then press Enter when done.
6. Enter the size for the RAID volume in gigabytes. The default value will be the maximum size. If you specify a smaller size, you will be able to create a second volume in the remaining space using the same procedure.
7. After this is done, exit the Option ROM user interface.

11.5 Using Intel® Rapid Storage Technology UEFI User Interface

Note: This section is OEM dependent. Where/how the OEM chooses to implement the UEFI UI is based on OEM preference

1. Upon re-boot, launch the Intel® RST UEFI user interface (HII compliant)

Figure 4



The UEFI UI is divided into three main sections:

1. Section 1 is the main section.
    a. It displays RAID configuration and status information
    b. It displays RST UEFI driver version (see arrow #4 in figure above)

c. It displays physical devices enumerated by the RST UEFI driver that are not part of the RAID volume

2. Section 2 gives a brief description of current page of the UI

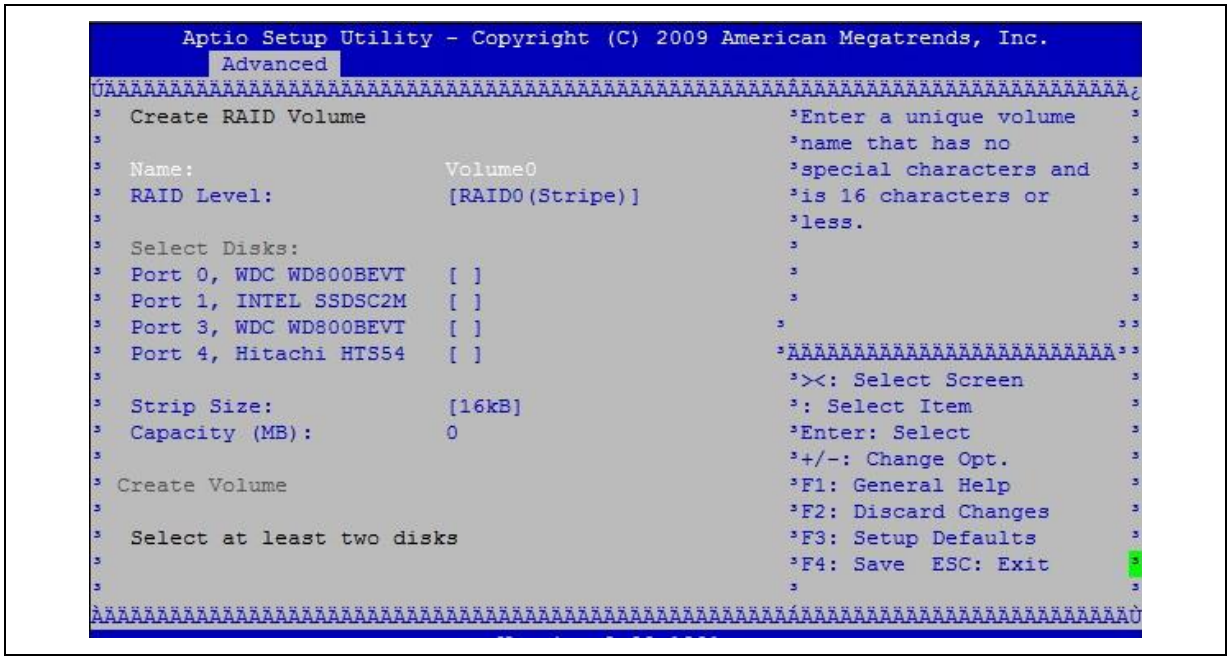
3. Section 3 gives information on how to navigate within the current page of the UEFI UI.

Note: ***This section is not implemented by the RST UEFI driver and is specific to the BIOS that was used for documentation purposes.***

2. In the Main Menu, select 'Create RAID Volume'

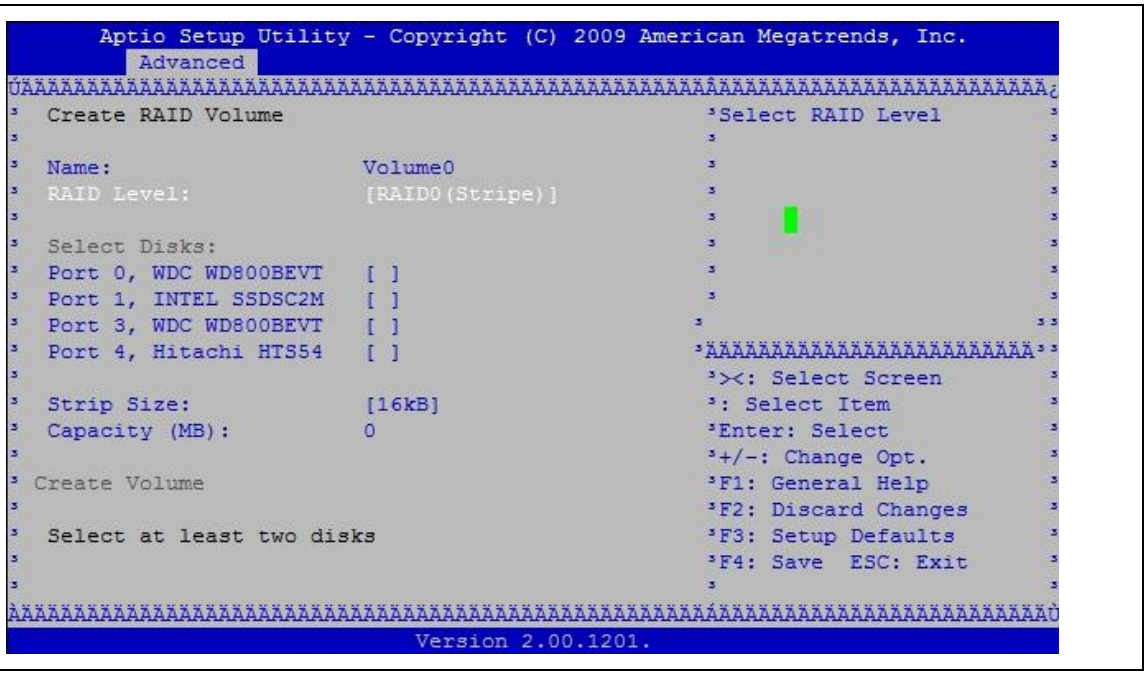
   a. Enter the name you want to use for the RAID volume, then press <Enter>.

Figure 5



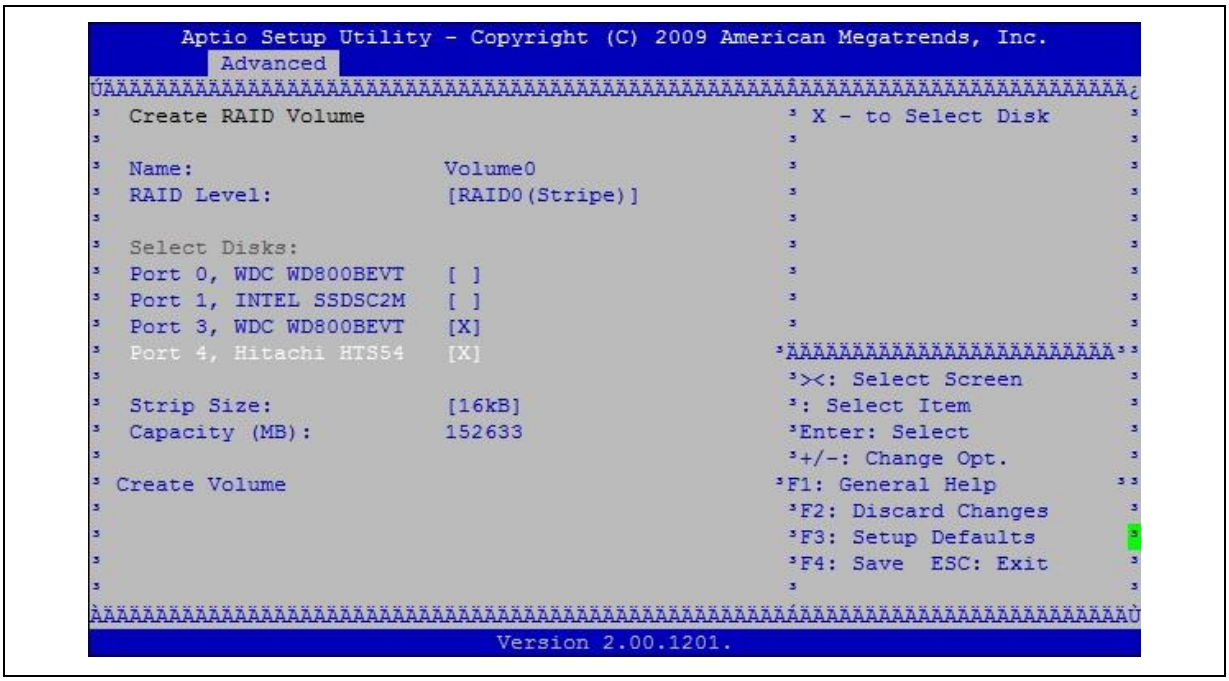
   b. Scroll down to 'RAID Level" and press <Enter> to select a RAID level

Figure 6



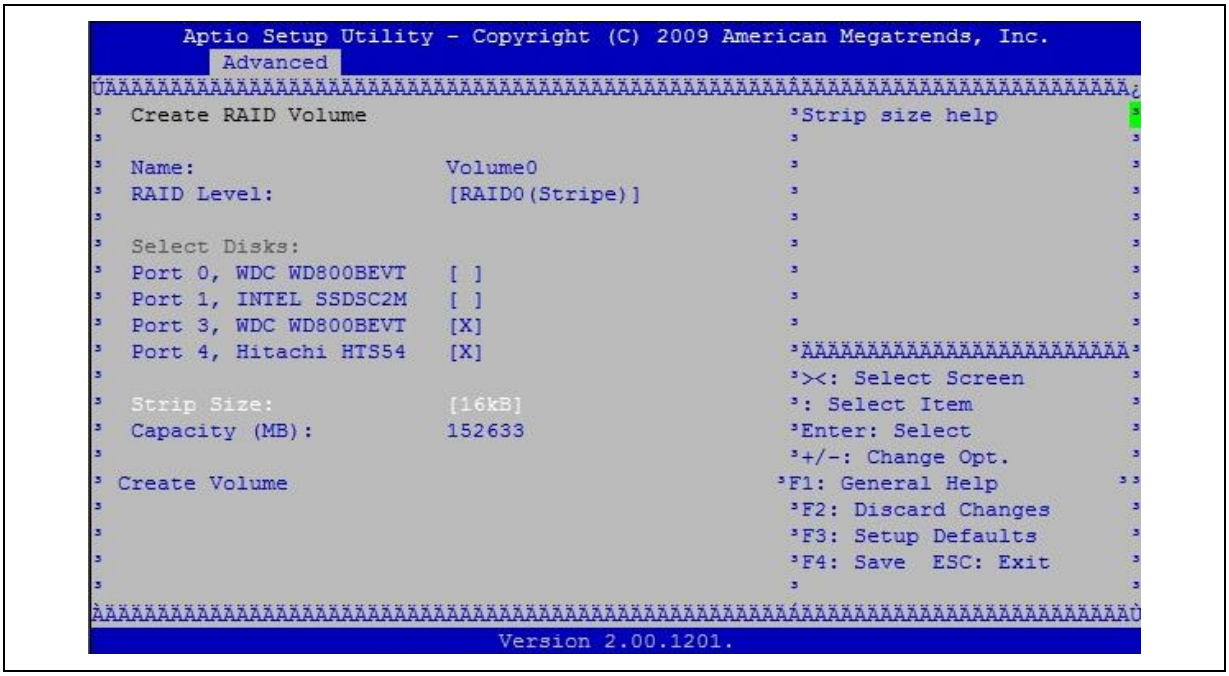
c. Scroll down to 'Select Disks' and at each disk that you wish to include in the RAID volume press <space bar>

Figure 7



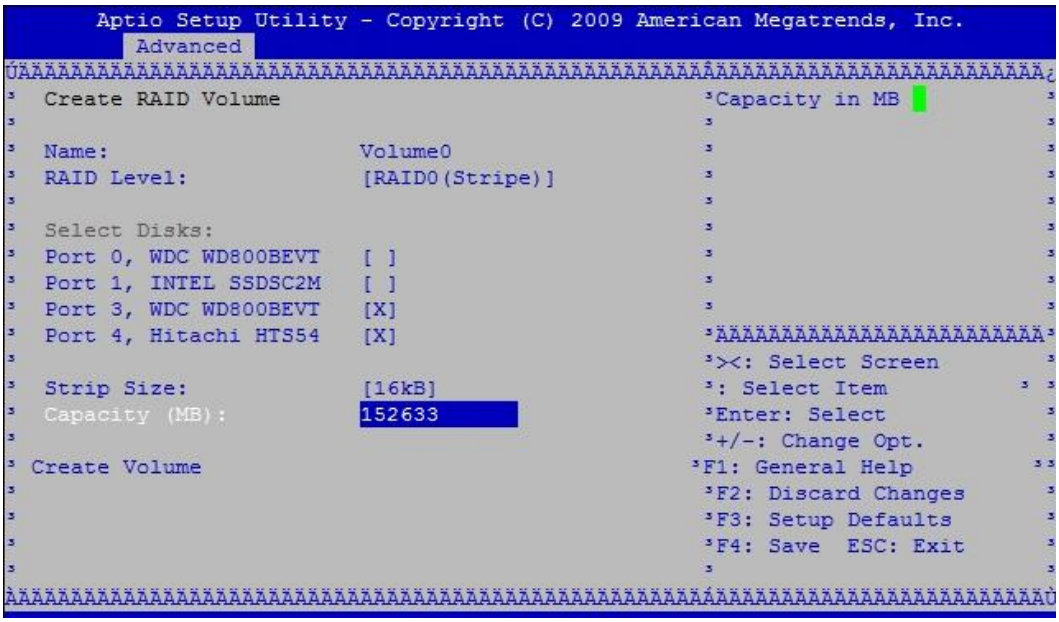
d. Next scroll down to 'Strip Size' and press <enter> to select a Strip size or continue if you wish to use the default strip size

Figure 8



e. Next scroll down to 'Capacity (MB)' where the maximum capacity is selected and displayed in MB. To select a smaller capacity for the RAID volume, type in the size in MB that you wish to use

Figure 9



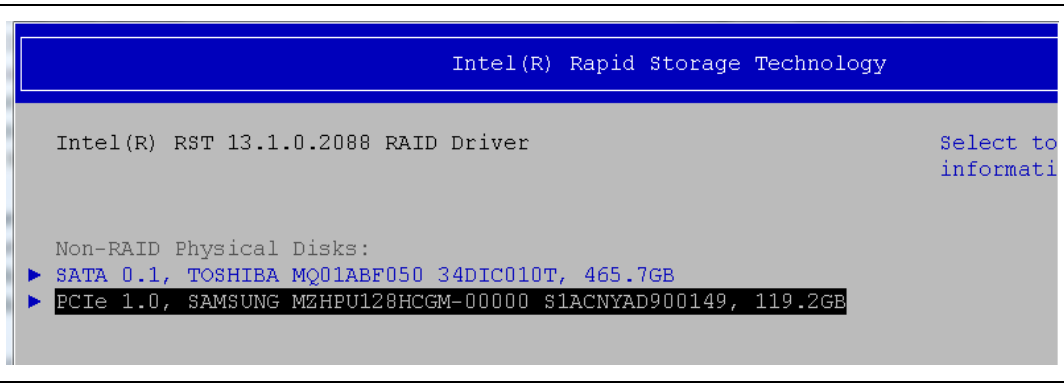
f. Next scroll down to 'Create Volume' and press <Enter>

3. After this is done, exit the Intel® RST UEFI UI.

Note: ***The "Create Volume" action will only be enabled if the RAID volume options selected will result in a valid configuration.***

4. Changes in HII, Beginning with Intel® RST UEFI 13.0, for PCIe Devices include new labeling for Devices and multiple controller management ability.
   - Device Ids numbering scheme = <Device Type><Controller ID>"."<Device ID>
   - Example below: "PCIe 1.0"

Figure 10



11.6 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows*)

Note: rstcli and rstcli64 can be used interchangeably below.
Run "rcfgsata.exe in DOS environment (or rcfgsata.efi from UEFI shell) or "rstcli.exe (or rstcli64.exe)" (Windows environment) with the following command line flags to create a RAID volume.
With PCIe Storage, the command line utilities will require a controller ID to be specified when creating RAID volumes:

SATA SSD, HDD, SSHD Controller ID = 0
PCIe AHCI SSD Controller ID = 1
PCIe NVMe SSD Controller ID = 2

The following command line will instruct the utility to create a RAID 0 volume named "OEMRAID0" on hard drives attached to the SATA Controller (Controller #0) on Port 0 and 1 with a strip size of 128 KB and a size of 120 GB:

```
C:\>rcfgsata.exe (or rcfgsata.efi) /C OEMRAID0 /DS 0.0 0.1 /SS 128 /L 0 /S 120
    C:\>rstcli.exe –C –l 0 –n OEMRAID 0-0-0-0 0-1-0-0 –s 128 –z 120
```

The following command will create a RAID volume using all of the default values. It will create a RAID 0 volume with a strip size of 128 KB on the two hard drives in the system. The volume will be the maximum size allowable.

```
    C:\>rcfgsata.exe /C OEMRAID0 (requires that only two disks can be attached to the
    system)
```

The following command line will instruct the utility to create a RAID 0 volume named "PCIeRAID0" on 1 PCIe AHCI SSD (Controller #1) and 1 PCIe NVMe SSD (Controller #2) attached to the system on remapped Port 0 and Port 2 with a strip size of 128 KB and a size of 120 GB:

```
C:\>rcfgsata.exe (or rcfgsata.efi) /C OEMRAID0 /DS 1.0 2.2 /SS 128 /L 0 /S 120
    C:\>rstcli.exe –C –l 0 –n OEMRAID 0-0-0-0 0-1-0-0 –s 128 –z 120
```

The following command line will display usage for all support command line parameters:

```
C:\>rcfgsata.exe (or rcfgsata.efi) /?
    C:\>rstcli.exe --help
```

Note: Selecting the strip size is only applicable for RAID 0, RAID 5, RAID 10 levels. Strip size is not applicable for RAID 1.

§ §

12 Deleting a RAID Volume

RAID volumes can be deleted in three different ways. The method most widely used by end-users is the Windows user interface utility. The second method is to use the Intel Rapid Storage Technology Option ROM user interface. The third way, used by OEMs only, uses the RAID Configuration utility.

12.1 Using the Windows* User Interface Utility

1. Run the Intel Rapid Storage Technology UI from the following Start menu link:

   Start→All Programs→Intel® Rapid Storage Technology →Intel Rapid Storage Technology UI

2. Under 'Status' or 'Manage' Click on the volume you want to delete. The user will be presented with the volume properties on the left.
3. Click on 'Delete volume'
4. Review the warning message, and click 'Yes' to delete the volume.
5. The 'Status' page refreshes and displays the resulting available space in the storage system view. You can now use it to create a new volume.

12.2 Using the Option ROM User Interface

1. Upon re-boot, you will see the Intel Rapid Storage Technology option ROM status message on the screen – press CTRL-I to enter the option ROM user interface.
2. Within this UI, select option #2 'Delete RAID volume'.
3. You should be presented with another screen listing the existing RAID volume.
4. Select the RAID volume you wish to delete using the up and down arrow keys.
5. Press the Delete key to delete the RAID volume
6. Press Y to confirm.

Note: Option #3 'Reset Hard Drives to Non-RAID' in the option ROM user interface may also be used to delete a RAID volume. This resets one or more drives to non-RAID status, by deleting all metadata on the hard drives. This has the effect of deleting any RAID volumes present. This function is provided for re-setting the hard drives when there is a mismatch in RAID volume information on the hard drives. The option #2 'Delete RAID Volume' on the contrary, will allow deleting a volume at a time, while retaining the existing RAID array metadata (for instance Matrix RAID).

12.3 Using the Intel® Rapid Storage Technology UEFI User Interface

Note: This section is OEM dependent. Where/how the OEM chooses to implement the UEFI UI is based on OEM preference. Use the following example for Intel CRB.

1. Upon re-boot, enter the system BIOS and select the Intel® Rapid Storage Technology menu for the UEFI user interface
2. In the Main Menu, go to the 'RAID Volumes' section, highlight the volume to be deleted and press <Enter>
    a. Select 'Delete', then press <Enter>.
    b. At the dialogue box press <Enter> to confirm the deletion of the volume

Note: All data on the volume will be lost!)

3. After this is done, exit the Intel® RST UEFI UI.

12.4 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows*)

Run "rcfgsata.exe in DOS environment (or rcfgsata.efi in UEFI shell)" or "rstcli.exe/rstcli64.exe" (Windows environment) with the following command line flag to delete a RAID volume. The following command line will instruct the utility to delete a RAID 0 volume named "OEMRAID0"

```
C:\>rcfgsata.exe /D OEMRAID0
C:\>rstcli.exe --manage --delete OEMRAID0
```

The following command line will display usage for all support command line parameters:

```
C:\>rcfgsata.exe(rcfgsata.efi) /?

C:\>rstcli.exe --help
```

§ §

13 Common RAID Setup Procedures

13.1 Build a SATA RAID 0, 1, 5 or 10 System

This is the most common setup. This configuration will have the operating system striped for RAID 0, or mirrored for RAID 1, or striped with parity for RAID 5, or mirrored and striped across two or up to four drives for RAID 10. All RAID member drives must be from the same BUS PROTOCOL GROUP. To prepare for this, you must have the Intel RAID driver on a floppy drive (USB). See the procedure for creating this floppy (USB) further down in this document.

1. Assemble the system using a motherboard that supports Intel Rapid Storage Technology and attach the drives depending on the RAID level that will be built.

2. Enter System BIOS Setup and ensure that RAID mode is enabled. This setting may be different for each motherboard manufacturer. Consult the manufacturer's user manual if necessary. When done, exit Setup.

13.1.1 Using the Legacy OROM User Interface

1. Upon re-boot if your system is using a legacy OROM, you will see the Option ROM status message on the screen – press CTRL-I to enter the Intel Rapid Storage Technology Option ROM user interface.

2. Within this UI, select option '1. Create RAID Volume'. When 'Create RAID Volume' menu is displayed, fill the following items:

    a. Name: Enter a volume name, and press Enter to proceed to next menu item,

    b. RAID Level: select RAID level (0, 1, 5, 10), and press Enter to proceed to next menu item;

    c. Disks: press Enter on 'Select Disks' to select the hard drives to be used for your configuration.

    d. Within the 'SELECT DISKS' window, choose the hard drives and press Enter to return to the 'MAIN MENU'.

    e. Strip Size: Applicable for RAID levels 0, 5, and 10 only. You may choose the default size or another supported size in the list and press Enter to proceed to the next item.

    f. Capacity: The default size would be the maximum allowable size summation of all the drives in your configuration. You may decrease this volume size to a lower value. If you specified a lower capacity size volume, the remaining space could be utilized for creating another RAID volume. Press Enter to proceed to the next item.

    g. Create Volume: Press Enter to create a volume.

h. Press 'Y' to confirm the creation of volume.

3. After this is done, exit the Intel Rapid Storage Technology option ROM user interface by pressing the Esc key or Option #4.

4. Begin OS setup by booting from the Windows OS installation CD.

5. Installation procedures as follows: Use the 'load driver' mechanism when prompted. Insert a USB key with the Intel® RST driver and browse to the directory on the USB key where the driver that you wish to install is located. Select the driver INF file. If correct the proper Intel controller for your system will be shown. Continue the driver install.

6. Finish the Windows installation and install all other necessary drivers.

7. Install the Intel Rapid Storage Technology software package obtained from the Intel VIP website. This will add the *Intel* Rapid Storage Technology UI that can be used to manage the RAID configuration.

13.1.2 Using the UEFI HII User Interface

1. Upon re-boot if your system is using the RST UEFI Driver and the HII protocol is in the system BIOS, you will see the Intel® Rapid Storage Technology option within the BIOS setup menu.

2. Select this menu. Choose the 'Create RAID Volume'. When 'Create RAID Volume' menu is displayed, fill the following items:

    a. Name: Enter a volume name, and press Enter to proceed to next menu item,

    b. RAID Level: select RAID level (0, 1, 5, 10), and press Enter to proceed to next menu item;

    c. Disks: press space bar to 'Select Disks' to select the devices to be used for your configuration.

    d. Within the 'SELECT DISKS' window, choose the devices and press Enter to return to the 'MAIN MENU'.

    e. Strip Size: Applicable for RAID levels 0, 5, and 10 only. You may choose the default size or another supported size in the list and press Enter to proceed to the next item.

    f. Capacity: The default size would be the maximum allowable size summation of all the drives in your configuration. You may decrease this volume size to a lower value. If you specified a lower capacity size volume, the remaining space could be utilized for creating another RAID volume. Press Enter to proceed to the next item.

    g. Create Volume: Press Enter to create a volume.

    h. Press 'Y' to confirm the creation of volume.

3. After this is done, exit the Intel Rapid Storage Technology menu HII user interface by pressing to save changes and the Esc key.

4. Begin OS setup by rebooting from the Windows OS installation CD.

5. Installation procedures as follows: Use the 'load driver' mechanism when prompted. Insert a USB key with the Intel® RST driver and browse to the directory on the USB key where the

driver that you wish to install is located. Select the driver INF file. If correct the proper Intel controller for your system will be shown. Continue the driver install.

6. Finish the Windows installation and install all other necessary drivers.

7. Install the Intel Rapid Storage Technology software package obtained from the Intel VIP website. This will add the *Intel* Rapid Storage Technology UI that can be used to manage the RAID configuration.

13.2 Build a "RAID Ready" System

The following steps outline how to build an Intel "RAID Ready" system with OS installed on a single SATA hard drive. A "RAID Ready" system can be upgraded to RAID 0, RAID 1, RAID5 or RAID 10 at a later time using the RAID migration feature built into Intel Rapid Storage Technology. Intel® RST enables you to install additional SATA hard drives, and then migrate to a RAID level volume without re-installing the operating system.

1. Assemble the system using a motherboard that supports Intel Rapid Storage Technology with Intel Rapid Storage Technology OROM integrated into the BIOS and attach one SATA hard drive.

2. Enter System BIOS Setup; ensure that RAID mode is enabled. This setting may be different for each motherboard manufacturer. Consult your manufacturer's user manual if necessary. When done, exit Setup.

3. Begin Windows Setup by booting from the Windows OS installation CD.

4. Installation procedures as follows: Use the 'load driver' mechanism when prompted. Insert a USB key with the Intel® RST driver and browse to the directory on the USB key where the driver that you wish to install is located. Select the driver INF file. If correct the proper Intel controller for your system will be shown. Continue the driver install:

5. Finish the Windows installation and install all other necessary drivers.

6. Install the Intel Rapid Storage Technology software package obtained from the Intel VIP website. This will add the *Intel* Rapid Storage Technology UI that can be used to manage the RAID configuration.

13.3 Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System

If you have an existing "RAID Ready" system as defined in section Build a SATA "RAID Ready" System, then you can use the following steps to migrate from a single-drive non-RAID configuration to a two drive RAID 0 or RAID 1 configuration. The resulting configuration will be identical to that created by the procedure in section Build a SATA RAID0, 1, 5 or 10 System. To prepare for this, you will need to install another drive with a capacity equal to or greater than the capacity of the drive being used as the source hard drive and also belong to the same BUS PROTOCOL GROUP as the source drive.

1. Note the port number of the source hard drive already in the system; you will use this to select hard drive for preserving data for the migration.

2. Install the second drive on the system.

3. Boot Windows, then install the Intel Rapid Storage Technology software, if not already installed, using the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel Rapid Storage Technology UI and start menu links.

4. Open the Intel Rapid Storage Technology UI from the Start Menu and select the volume type under Create from the Actions menu. Click on 'Next'

5. Under the configure options provide the volume name , select disks

6. When the disks are selected, the user will be presented the option to select the disk on which to preserve the data. Here the user need to select the right disk on the which the data needs to preserved and migrated

7. After the migration is complete, reboot the system. If you migrated to a RAID 0 volume, use Disk Management from within Windows in order to partition and format the empty space created when the two hard drive capacities are combined. You may also use third-party software to extend any existing partitions within the RAID volume.

13.4 Migrate an Existing Data Hard Drive to a RAID 0 or RAID 1 Volume

If you are booting from a parallel ATA (PATA*) drive that contains the operating system, you may use the Intel Rapid Storage Technology to create a RAID 0 or RAID 1 volume on two SATA drives. Also, if you have a single SATA hard drive that contains program or personal data, you may use the migration feature to use this hard drive as the source hard drive for a migration. After the migration is completed, you will have a two hard drive RAID 0 volume where data is striped or a two hard drive RAID 1 volume where the data is mirrored across the two SATA hard drives. To do this, the PCH I/O RAID Controller must be enabled in the BIOS and you must have the Intel Rapid Storage Technology software installed.

Begin with a system where you are booting from a PATA hard drive. Make sure the PCH I/O RAID controller is enabled and the Intel Rapid Storage Technology is installed. Then do the following:

1. Note the serial number of the SATA hard drive that is already installed. You will use this to select it as the source hard drive when initiating the migration.

2. Physically attach the second SATA hard drive to the available SATA port.

3. Boot to Windows, install the Rapid Storage Technology software, if not already installed, using the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel Rapid Storage Technology UI and start menu links.

4. Open the Intel Rapid Storage Technology UI from the Start Menu.

5. Follow steps 4 to 7 in section Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System

13.5 Migrating From one RAID Level to Another

RAID level migration allows an existing RAID configuration to be migrated to another RAID configuration. The following migrations are possible.
NOTE: Not all migrations are supported on all chipsets. The support varies depending on the chipset

and the ports supported on the chipset (For supported migrations for each chipset Intel Rapid Storage Technology product requirements document):

Change Type from	To
2-disk recovery volume	2-disk RAID 1
2-disk RAID 1	2-disk recovery volume
2-disk RAID 1	2-disk RAID 0 3, 4, 5 or 6-disk RAID 0 3, 4, 5 or 6-disk RAID 5
2-disk RAID 0	3, 4, 5 or 6-disk RAID 5
3-disk RAID 0	4, 5 or 6-disk RAID 5
4-disk RAID 0	5 or 6-disk RAID 5
4-disk RAID 10	4, 5 or 6-disk RAID 5

Note: In order for the migration options to be accessible, the minimum required SATA hard drives for the RAID level have to be met.

Follow the procedure illustrated below:

1. Start the Intel Rapid Storage Technology UI application:

   Start Menu ->All Programs -> Intel Rapid Storage Technology -> Intel Rapid Storage Technology UI

2. Under 'Status' or 'Manage', in the storage system view, click the array or volume to which you want to modify. The volume properties now display on the left.
3. Click 'Change type'.
4. In the 'Change Volume Type' dialog, type a new name if you want to change the default name.
5. Select the new volume type, and then click 'OK'.
6. The 'Manage' page refreshes and reports the new volume type.
7. After the migration starts, you can view the migration progress under status.
8. When the Status field indicates volume as ‘Normal’, the migration is complete.

13.6 Create a RAID Volume on Intel® SATA Controller While Booting to Different Controller

This configuration is for users who would like to use a RAID 0 volume as a high performance data hard drive or use the data redundancy properties of RAID 1. Starting with a configuration where the system is booting to a Windows, with installation on a different disk controller, the user can add two SATA hard drives and create a RAID volume on them.

1. Physically install two SATA hard drives to the system.

2. Enter System BIOS Setup; ensure that RAID mode is enabled. This setting may be different for each motherboard manufacturer. Consult your manufacturer's user manual if necessary. When done, exit Setup.

3. Boot to Windows; install the Intel Rapid Storage Technology software, if not already installed, use the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel Rapid Storage Technology UI and Start menu links.

4. Use the Intel Rapid Storage Technology UI to create a RAID 0 volume on two SATA drives according to the procedure in section Using Intel® Rapid Storage Technology UI of this document.

5. After the RAID volume is created, you will need to use Windows Disk Management or other third-party software to create a partition within the RAID volume and format the partition. At this point, you may begin to copy files to, or install software on, the RAID volume.

13.7 Build a RAID 0 or RAID 1 System in an Automated Factory Environment

This is a two-part process. First, create the master image of the Windows installation; you will load these on the system before they are delivered to the customer. The second part is to apply this image to a system that has two SATA hard drives installed with a RAID 0 or RAID 1 volume. This procedure will apply the image to the RAID volume so that the system may boot from it and the operating system will be fully striped by the RAID 0 volume or mirrored by the RAID 1 volume. The same procedure, and master image, could be applied to a single SATA hard drive to create a "RAID Ready" system.

13.7.1 Part 1: Create Master Image

1. Build a RAID 0 or RAID 1 System as described in section Build a SATA RAID 0, 1, 5 or 10 System of this document.

2. Install the Intel Rapid Storage Technology software from the CD-ROM included with your motherboard or after downloading it from the Internet. This will add the Intel Rapid Storage Technology UI that can be used to manage the RAID configuration in Windows*.

3. Use third-party software to create an image of the RAID volume as if it were a physical hard drive or create an image of the partition within the RAID volume containing the operating system, program and data files.

4. Store it in a place where it can be accessed by systems on the assembly line.

13.7.2 Part 2: Apply Master Image

1. Assemble the system using a motherboard that supports Rapid Storage Technology and attach two SATA hard drives.

2. Enter System BIOS Setup; ensure that RAID mode is enabled. This setting may be different for each motherboard manufacturer. Consult your manufacturer's user manual if necessary. When done, exit Setup.

3. If the system has CSM on, and can boot to a DOS environment, use the Intel RAID Configuration utility (RCfgSata.exe). Else if CSM is off, or not present, boot to the UEFI shell and use the

RcfgSata.efi utility to create a RAID volume. The following command line will instruct the utility to create a RAID 0 volume named "OEMRAID0" on hard drives on Port 0 and 1 with a strip size of 128 KB and a size of 120GB (rcfgsata.efi can replace rcfgsata.exe if using the UEFI shell environment):

'/DS' for device selection will distinguish the different controllers for device selection: <Controller><Port>

0.0 = SATA device on port #0

1.3 = PCIe AHCI device remapped to port #3

2.6 = PCIe NVMe device remapped to port #6

Create RAID 0 using 1 PCIe AHCI SSD on port 3, and 1 PCIe NVMe SSD on port 6:

C:\>rcfgsata.exe /C:OEMRAID0 /DS:1.3 2.6 /SS:128 /L:0 /S:120.

The following command line will display all supported command line parameters and their usage: *C:\>RCfgSata.efi /?*

4. The system does not need to be rebooted before moving on to the next step. If there are no PATA hard drives in the system, the RAID volume created will become the boot device upon reboot.

5. Use third-party software to apply the image created in Part 1 to the RAID volume you created in Part 2.

§ §

14 RAID Volume Data Verification and Repair Feature

This feature is available starting with Intel® Matrix Storage Manager 6.1.

14.1 Verify and Repair Volume Feature

The RAID volume verification feature identifies any inconsistencies or bad data on a RAID 0, RAID 1, RAID 5, or RAID 10 volume and reports the number of inconsistencies or number of blocks with media errors found during RAID volume data verification.

When the verification process is complete, a dialog will appear that displays the number of verification errors, verification errors repaired and blocks with media errors that were found.

Follow the below steps to start RAID volume data verification

1. Under 'Status' or 'Manage' click on the RAID volume you want to perform the verify operation under 'storage system view'. The volume properties now display on the left.

2. Click on 'Advanced' and then Click on 'Verify'

3. For RAID 0 the verification process starts once you click 'verify'. For RAID1, 5, 10, Recovery volumes, a dialog box with check box option to repair the errors found automatically during the verification process is present. If the user wants to perform repair you can select this box and then click 'verify'.

4. The verification progress is shown under 'status'

5. When the verification process is complete and the volume status is set to normal, now you can click on the volume under 'status' or 'manage'. Under the volume properties to the left under 'Advanced' you can view the number of verification errors, verification errors repaired and blocks with media errors that were found.

14.2 Verify and Repair Scheduler

The Verify and Repair feature includes a scheduler for the Verify and Repair (V and R) operation. To enable the scheduler take the following steps:

Pre-conditions: UI installed, at least 1 RAID volume on the system that is initialized, in normal state, and a valid RAID type (RRT, R0**, R1, R5, R10) ***RAID 0 volumes can only do a Verify; they cannot be repaired***

1. Login to Windows and launch the Intel® RST UI and click on the 'Preferences' tab at the top of the UI

2. From the 'Preferences' page, select the 'Scheduler' button on the left navigation pane to display the

3. Check mark the 'Enable scheduler' checkbox

4. Select 'Recurrence' schedule: Once (default), Daily, Weekly, or Monthly

5. Select the 'Start Date'; day for the scheduler to begin/run the V and R operation



6. Select the ‘Time’ of the scheduled runs on a 24 hour clock
7. Select the ‘Recur every’ schedule: choices will vary depending upon what is selected for ‘Recurrence’ (this step is not applicable for Recurrence of once)
8. Select whether or not to Automatically Repair Errors encountered during the Verify operation
9. Click ‘Apply Changes’ to enable

§ §

15 Intel® Rapid Recover Technology

This technology utilizes RAID 1 functionality to copy data from a designated Master drive to a designated Recovery drive with the following limitations:

- **The size of the Master drive must be less than or equal to the size of the Recovery drive.**
- **The size of the Master drive is limited to less than or equal to (<=) 1.3125TB in capacity.**

When a Recovery volume is created, complete capacity of the Master drive will be used as the Master volume. Only one Recovery Volume can exist on a system. There are 2 methods of updating the data on the Master to the Recovery drive. They are:

- Continuous Update Policy
- On Request Update Policy

When using the continuous update policy, changes made to the data on the master drive while the recovery drive is not available are automatically copied to the recovery drive becomes available. When using the Update on request policy, the master drive data can be restored to a previous state by copying the data on the recovery drive back to the master drive.

Some of the advantages of Intel® Rapid Recover Technology are:

- More control over how data is copied between master and recovery drives
- Fast volume updates (only changes to the master drive since the last update are copied to the recovery drive)
- Member hard drive data can be viewed in Windows* Explorer
- Better power management on mobile systems by spinning down the Recovery drive when in On Request Update Policy mode or when the Recovery drive goes offline when in Continuous Update Policy mode.

Applications: Critical data protection for mobile systems; fast restoration of the master drive to a previous or default state.

15.1 Creating a Recovery Volume Through the RAID Option ROM

A Recovery volume consists of two disks – a primary disk and a recovery disk.

A Recovery volume can be created through the RAID Option ROM or through Intel® Rapid Storage Technology UI application.

Follow the below steps to create a Recovery volume through the OROM

1. Enter the OROM by pressing the <Ctrl> and <I> keys early during system POST.
2. Under the ‘Create RAID’ volume option, select the option to create a Recovery volume.
3. Select the Primary disk and the Recovery disk.

Note: The Primary disk size must be less than or equal to the Recovery disk size.

OROM Recovery menu provides the following options

1. Enable Only Recovery Disk
2. Enable Only Master Disk

15.2 Creating a Recovery Volume Using Intel® Rapid Storage Technology UEFI User Interface

Follow the below steps to create a Recovery volume through the UEFI UI

i. Enter the BIOS Setup Menu and select Intel® Rapid Storage Technology menu.

ii. Select 'Create RAID Volume'.

iii. Select the RAID Level [Recovery].

iv. Select Name and type in the name of the volume.

v. Highlight each drive and press <space> bar to select either R or M depending on which disk will be Recovery or Master.

vi. Highlight Synchronization, press <Enter> and select Mode of 'On Request' or 'Continuous'

vii. Highlight 'Create Volume' and press <Enter>

viii. Volume created will be displayed on Main Page.

15.3 Creating a Recovery Volume Through Intel® RST UI

To create a Recovery volume through the Rapid Storage Technology UI, the system needs to be configured in RAID mode with 2 drives. Boot the system and open the Rapid Storage Technology UI application.

Follow the below steps to create a Recovery Volume

1. Under Create select the volume type as 'Recovery' and click 'Next'
2. Under the 'Configure Volume' you can change the default volume name if you want, then select the 'master' disk and then the 'recovery' disk. Now change the 'update' mode if needed to 'On Request'. The default selection is 'continuous'.
3. Once all the above selections are made, click 'Next'
4. Under 'Confirm' review the selected configuration. If you are not ok with the configuration click 'back' or click 'create volume' if you are fine with the configuration.

5. Now you will see a dialog box with warning message and read the warning message before clicking ‘ok’ to make sure you are erasing data on the right disk.

6. Once you click ‘ok’ the volume creation starts and progress of the volume creation can be viewed under status. Once the status is set to ‘normal’ the volume creation is completed.

7. The system will synchronize the Primary with the Recovery disk once after the creation of the Recovery volume.

15.4 Changing Recovery Volume Modes

When you have a recovery volume on your system in ‘continuous mode’ or ‘on request’ mode and you need to change the mode of the recovery volume, follow the below steps

1. Open Intel® Rapid Storage Technology UI.

2. Under ‘Manage’ or ‘Status’ click on the recovery volume under the storage system view on right where you need to change the update mode. The volume properties now display on the left view

3. Click ‘change mode’ and then click ‘yes’ to confirm.

4. The page refreshes and the volume properties report the new update mode. NOTE: Disabling the continuous update policy requires the end-user to request updates manually. Only changes since the last update process are copied. The recovery volume will remain in On Request Policy until the end-user enables continuous updates.

15.5 Update Recovery Volume in On Request Update Policy

When the recovery volume is ‘on request’ mode on your system and you need to synchronize the data between both the master and recovery disk, follow the below instructions

1. Open Intel® Rapid Storage Technology UI.

2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.

3. Click ‘Update data’.

4. A dialog box is shown stating that the only changes since the last update will be copied. Select the check box if you don’t want this confirmation message to display each time you request an update. Click ‘Yes’ to confirm.

5. The progress of update process can be viewed under ‘status’ or ‘manage’.

15.6 Access Recovery Drive Files

When data recovery to the master disk of a recovery volume is required, you can use ‘access the recovery disk files’ option. This action is only available if a recovery volume is present, in a normal state, and in on request update mode. Follow the below instructions to access the recovery drive file when you have a recovery volume in ‘on request’ mode on your system (If the recovery drive is not in continuous mode, use the instructions in section 8.3 to change the mode)

1. Open the ‘Intel Rapid Storage Technology UI’.

2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.

3. Click on 'Access recovery disk files'.

4. Now you can view recovery disk files using Windows Explorer*.

Note: The recovery drive can only be accessible in read only mode and data updates are not available in that state

15.7 Hide Recovery Drive Files

This action is only present when the recovery drive is on request mode and the recovery drive files are accessible. Follow the below instructions to hide the recovery drive files

1. Open the 'Intel Rapid Storage Technology UI'.

2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.

3. Click 'Hide recovery disk files'.

4. Now the recovery drive files are no longer accessible in Windows Explorer.

5. The page refreshes and data updates on the volume are now available.

15.8 Scenarios of Recovering Data

Scenario 1:
What happens if the Recovery drive that is part of the Intel® Rapid Recover Technology volume fails or gets stolen?

Solution:
When a Recovery drive that is part of an Intel® Rapid Recover Technology volume fails, follow the below steps to set up a new disk as the Recovery drive.

1. Shut down the system.
2. Remove the failed Recovery disk and insert a new hard drive. The size of the new drive must be greater than or equal to the Master drive.
3. Boot to the Master drive and open Intel Rapid Storage Technology UI.
4. Under 'Status' or 'Manage', in the storage system view, click the recovery volume to be rebuilt. The volume properties now display on the left.
5. Click on 'rebuild to another disk'
6. Now a dialog box is shown requesting you to select one of the non-RAID disks to rebuild the volume.
7. Once the disk selection is complete, click 'rebuild'
8. Now you can view the progress of the build under 'status' or 'manage'

Scenario 2:

What happens if the Master Drive fails and/or the user would like to do a reverse synchronization to a new Master Drive?

Solution:
If the Recovery volume was in Continuous update policy when the Master drive crashed, then the system will continue to function off of the Recovery drive.

If the Recovery volume was in Update on Request policy, then a Master drive failure may result in a BSOD.

In either case, follow the below steps to create a new Master drive using the Recovery Drive.

1. Shut down the system.
2. Remove the old Master disk and connect a new Hard Disk Drive to be designated as the new Master disk. **Note:** The size of the new Master drive should be less than or equal to the Recovery disk.
3. Power on the system. It will automatically boot from the Recovery drive. After the operating system is running, select the Intel® Rapid Storage Technology UI from the Start Menu.
4. Under 'Status' or 'Manage', in the storage system view, click the recovery volume to be rebuilt. The volume properties now display on the left.
5. Click on 'rebuild to another disk'.
6. Now a dialog box is shown requesting you to select one of the non-RAID disks to rebuild the volume.
7. Once the disk selection is complete, click 'rebuild'.
8. Now you can view the progress of the build under 'status' or 'manage'.

Scenario 3:
What is the expected behavior if a power failure occurs (and no battery supply available) in the middle of migration for each of the below?

- Creating a recovery volume (migration)
- Updating a recovery volume (Copy some files from Master drive to Recovery drive)
- Verify and Repair a recovery volume
- Recovering a recovery volume (copy from a Recovery drive to a Master Drive

Solution:

In each case, upon the next reboot, the migration, or Verifying a Recovery Volume, or Verify and Repair a Recovery Volume or Recovering a Recovery Volume operation would continue normally starting from where it had been interrupted by the power failure.

In the case where the Recovery volume was getting updated or was being recovered, if it were a fast synchronization, then if writes had been in progress while the power was lost, then it would result in a dirty shutdown. As a result, the fast synchronization would degenerate to a slow synchronization or a complete update.

Note: If the system is running is on battery, the volume will not synchronize if it is in continuous update policy. If the volume is in Update on Request policy, then the synchronization will be successful.
Additional comments: need to call out that an on update volume should first be updated before the recovery disk is valid.

Scenario 4:
Once a system is configured with Intel Raid Recover Technology, a user would like to revert the Master Drive Data to a Previous State.

Solution:
If the recovery volume is set to the on request update policy, you can revert master drive data to the state it was in at the end of the last volume update process. This is especially useful when a virus is detected on the master drive or guests use your system.

1. Restart the system. During the system startup, press Ctrl-I to enter the user interface of the Intel® Rapid Storage Technology option ROM.
2. In the 'MAIN MENU' select 'Recovery Volume Options'.
3. In the 'Recovery Volume Options' menu, select 'Enable Only Recovery Disk' to boot from the recovery drive.
4. Exit the option ROM and start up Windows*.
5. After the operating system is running, select the Intel® Rapid Storage Technology UI from the Start Menu.
6. Under 'Status' or 'Manage', in the storage system view, click the recovery volume to be recovered. The volume properties now display on the left.
7. Click on ‘recover data’ and then click ‘ok’ on the dialog box.
8. Now you can view the progress of the recovery under ‘status’ or ‘manage’.
9. Once the recovery of the volume is completed, you can reboot to the master drive.

15.9 System Running from Recovery Drive

The “System Running from Recovery Drive “ is an existing feature in the current UI but is documented here for the sole purpose of providing Validation and Localization with the flow of expected behavior for test pass preparation.

15.10Drive Offline or Missing

System Configuration	2 hard drives: recovery drive connected, master drive offline or missing
Product Condition	Recovery volume created with recovery drive normal and master drive offline or missing

- Access UI OROM – Note that the master drive is designated as an offline disk or master drive missing
- Select option 4 Recovery Volume Options

Figure 11



Then Select Option 2 Enable Only Recovery Disk.

§ §

16 Pre-OS Installation of Intel® Rapid Storage Technology Driver

The Intel® Rapid Storage Technology driver can be loaded before installing the Windows OS on a RAID volume or when in AHCI mode. All later Windows OS releases do not require that the Intel® RST driver be installed and loaded prior to the OS installation. On those OS versions the Intel® RST driver can be loaded post OS installation. The Intel® Rapid Storage Technology AHCI driver can be installed over Window's native AHCI driver.

16.1 Pre-OS Driver Installation Using the "Load Driver" Method

1. During the Operating system installation, after selecting the location to install Windows click on 'Load Driver' button to install the Intel® RST storage AHCI/RAID driver.
2. When prompted, insert the media with the Intel® RST driver files and press Enter.
3. You can find the media and browse to the folder where the files are located.
4. Follow the steps to load the driver and return the installation.
5. Continue the installation.

§ §

17 Determining the Version of the RAID Driver

There are two accurate ways to do this. The first is to use the Intel Rapid Storage Technology UI. The second alternate method is to locate the driver (iaStorA.sys) itself and view its properties.

17.1 Using Intel® Rapid Storage Technology User Interface (UI)

Use this method if the Intel® RST UI is installed on the system; if not use the alternate method.

1. Run the Intel Rapid Storage Technology UI from the following Start Menu path:
2. Start→All Programs→Intel® Rapid Storage Technology →Intel Rapid Storage Technology UI
3. Click on the top menu button 'help' to launch the 'Help' window. In the 'help' window click the top menu button 'System Report^
4. If not already expanded, click on 'Intel® Rapid Storage Technology' link to expand the item. Under it you can view the driver version in the following format: WW.XX.YY.ZZZZ
5. This is the current version of the user interface utility installed on your system. The WW.XX.YY portion is the product release number; the ZZZZ portion is the build number. E.g. 10.5.1.1001.

17.2 Using Intel® RST File Properties (Alternate)

1. Locate the file "iaStorA.sys" within the following path:

   <System Root>\Windows\System32\Drivers
2. Right Click on "iaStorA.sys" and select Properties
3. Select the "Details" tab (for Windows 7; may vary for other OS versions)
4. At the top of this tab, there should be a parameter called "File version". Next to it is the version of the driver currently installed on your system. It should have the same format and version as the one you obtained using the Intel Rapid Storage Technology UI

17.3 Determining the Version of the Option ROM

There are two ways to determine the version of the Intel Rapid Storage Technology option ROM integrated into the system BIOS. Use the following procedure to determine the version.

17.3.1 Using Intel® Rapid Storage Technology UI

1. Follow the procedure illustrated in section 11.1
2. Look for the parameter RAID Option ROM version.



17.3.2 Using Intel® RST Option ROM User Interface

1. Early in system boot-up, during post, or when you see the "Intel® RAID for Serial ATA" status screen output, type CTRL-I. This will open the Option ROM user interface.
2. The following banner will be displayed:
3. Intel® Rapid Storage Technology option ROM w.x.y.zzzz Intel® SATA Controller
4. w.x.y.zzzz is the version of the Option ROM currently installed on your system. The w.x.y portion is the product release number; the zzzz portion is the build number.

17.3.3 Using the EFI Shell

If the UEFI Driver is enabled the following command can be issued from the EFI shell:

Shell:>Drivers

The Intel® RST UEFI driver will be shown along with version, where xx.x.x.xxxx will be replaced with the actual UEFI OROM Version i.e.:

"CD 0000000B B - - 1 2 Intel® RST xx.x.x.xxxx SATA Driver"

§ §

18 Un-Installation

Uninstalling the RAID driver could potentially cause an end-user to lose access to important data within a RAID volume. This is because the driver can only provide functionality for the Intel® SATA RAID controller. Therefore, Intel does not provide a way to permanently remove the driver from the system. However, disabling the Intel® SATA RAID Controller causes the operating system to not use the RAID driver.

The uninstallation application that is included with the Intel Rapid Storage Technology software can remove all components except the RAID driver (i.e. it removes the UI application, Start Menu links, Control Panel Applet, etc.).

Use the following procedures to remove the Intel Rapid Storage Technology software or to disable the SATA RAID controller:

18.1 Un-Installing Intel® RST Software (except the RAID Driver)

1. Run the Uninstall program from the following start menu link:
2. Start→All Programs→Intel® Rapid Storage Technology →Uninstall
3. The first dialog box that appears gives you the option of un-installing all components of the Intel Rapid Storage Technology software except the RAID driver. Click 'OK' to do so.
4. The next dialog box is a confirmation that you would like to un-install all components of the software except the RAID driver. Click 'Yes' to confirm.
5. All components of the software will be un-installed except the RAID driver. You should no longer see any Start menu links to the UI application or a control panel applet for Intel Rapid Storage Technology. However, the RAID configuration should still function normally.

18.2 Disabling the RAID Driver by Disabling the RAID Controller

WARNING: If you use this method and your computer's operating system is installed to a disk attached to the Intel® SATA RAID Controller, you will no longer be able to boot into that operating system!

1. Enter System BIOS Setup and disable RAID Mode. This setting may be different for each motherboard manufacturer. Consult your manufacturer's user manual if necessary. When done, exit Setup.
2. Reboot the system (The OS must have been installed on a disk not attached to the Intel® SATA RAID controller). You should no longer see the RAID Option ROM status screen during boot, and you should no longer see the Intel® SATA RAID Controller in Device Manager.
3. At this point, Windows will no longer be using the RAID driver and you will not have Intel RAID functionality. **All data contained in existing RAID volumes will no longer be accessible.** To re-enable Intel RAID functionality, re-enter System BIOS Setup and re-enable RAID mode.

Note: To Uninstall: End-users can use this same procedure to disable the Intel® SATA RAID Controller if necessary. In fact, the uninstall program used in section 12.1 of this document will

display a text file with a similar procedure. Run the Uninstall Program, click ‘Cancel’ when presented with the first dialog box, then click ‘Yes’ at the second dialog box to read the text document containing the procedure.

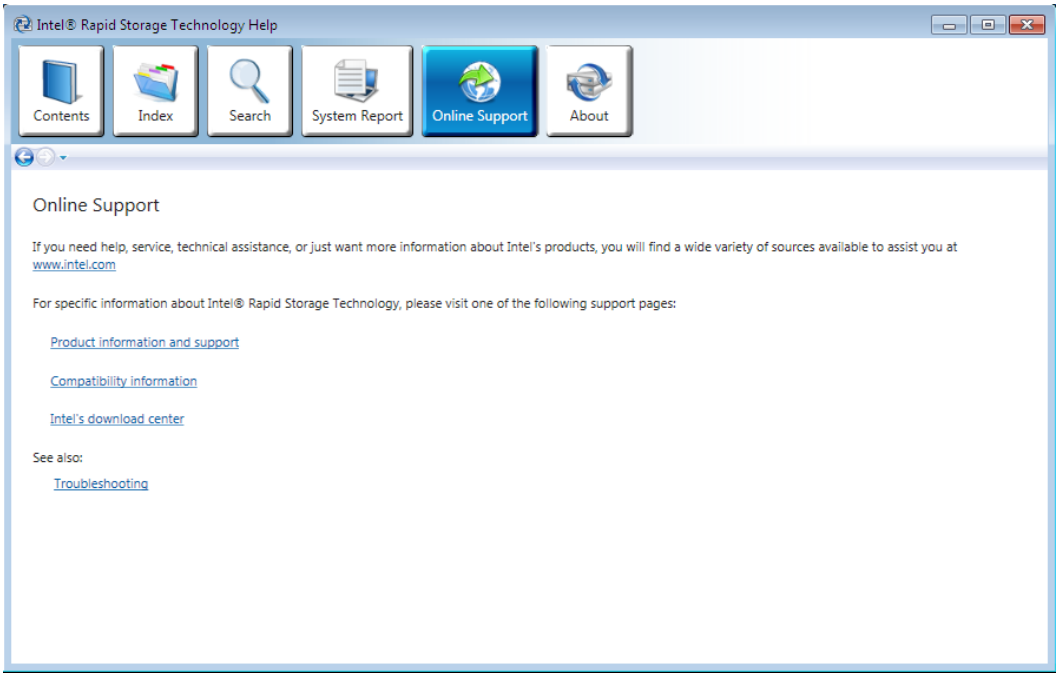
§ §

19 Registry Customizations

Note: Windows registry changes require reboot to take effect.

After installation of the Intel Rapid Storage Technology, the registry will contain keys to allow customization of several features. Customize Support URLs in Rapid Storage Technology UI

The Rapid Storage Technology UI [Help] Menu, Submenu [Online Support] when selected will display a pop-up window with the support URLs as shown in the figure below:



Product information and support : (http://www.intel.com/p/en_US/support/highlights/chpsts/imsm)

Compatibility information : (http://www.intel.com/support/chipsets/imsm/sb/CS-020680.htm)

Intel's download center : (http://downloadcenter.intel.com/SearchResult.aspx?lang=eng&ProductFamily=Chipsets&ProductLine=Chipset+Software&ProductProduct=Intel%c2%ae+Rapid+Storage+Technology+(Intel%c2%ae+RST)&ProdId=2101&LineId=1090&FamilyId=40)

19.1 Zero Power ODD Settings

Beginning with the Intel® RST 10.0 release and the Intel® Mobile Express Chipset SATA AHCI and the Intel® Desktop/Workstation/Server Express Chipset SATA AHCI controllers (PCH), the product supports the zero-power ODD feature (also referred to as ZPODD). It is intended to allow an unused ODD to be powered off, and then powered on only upon receipt of demand requests or when the ODD eject button is pressed. This goal is achieved by utilizing ACPI** methods to change the power condition of the drive when several platform (HW) conditions exist.

Note: **: This feature is not supported on Windows XP and older operating systems.

Associated with this feature are two registry keys located at

[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters]

1. ZPODD enable/disable

```
[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device]

"OddZeroPowerEnable"=dword(0, 1)
```

This key determines a platform's eligibility for the feature. When the value is zero then this feature will be disabled. When the value is non-zero or not present the feature will be enabled. Default value will be enabled (**1**).

2. ODD idle timeout

```
[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device]

"SecondsToOddZeroPower"=dword:(30, 300)
```

This key determines the idle timeout value. When the value is zero then this feature will be disabled. The value is the number of seconds the ODD must be idle (defined as a period of time in which no non-GESN commands are received; minimum value is 30 and maximum value is 300) before the ODD will be powered off. The default value is **60.** If the registry value is set to a value outside this range then the default value of 60 seconds will be used.

19.2 E-mail Notification UI Visible Enable/Disable

By default the e-mail notification feature is visible in the UI under the [Preferences] button. The registry can be modified to allow or disallow the end-user from being able to configure the system to allow email notification. The registry key by default is not populated in the registry. In order to remove the functionality from the UI the registry key has to be created using the following settings:

Open the registry editor and add the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Intel\IRST

Create a new DWORD (32) value as follows:

DisableEmail

Registry Value	Description
DisableEmail	DWORD(32) = 0 (default): By default or when this value is created and cleared to 0, the UI [Preferences] page displays an item in the left column of the page that allows the user to setup the desired e-mail preferences for the system.
	DWORD(32) = 1: When this value is created and set to 1, the UI will not display a menu item in the 'Preferences' page for the end-user to setup email notification on the system. The feature is disabled.

19.3 Disabling Maximized Mode Option for Intel® SRT

OEMs have the ability to disable the Accelerate Maximized mode option and limit the Intel® Smart Response Technology to Enhanced mode selection only.

The registry key by default is not populated in the registry. In order to remove the functionality from the UI the registry key has to be created using the following settings:

Open the registry editor and add the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Intel\IRST

Create a new DWORD (32) value as follows:

DisablePerformanceMode

Registry Value	Description
DisablePerformanceMode	DWORD(32) = 0 (default): By default or when this value is created and cleared to 0, the UI will allow the end-user to Accelerate disks/volumes in both Maximized and Enhanced mode.
	DWORD(32) = 1: When this value is created and set to 1, the UI will not display the option to Accelerate to Maximized mode and will not allow the user to change from Enhanced mode to Maximized mode

19.4 Rebuild On Hot Insert

When a redundant RAID volume is in the 'Degraded' state and a hot insert event of a new disk is detected by the RST driver, the driver will automatically begin a rebuild of the degraded RAID volume to the new disk. The new disk must meet all the requirements to be an array member of the degraded RAID array.

The registry key by default is not populated in the registry. In order to enable the functionality in the UI the registry key has to be created using the following settings:
Open the registry editor and add the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Create a new DWORD (32) value as follows:

RebuildOnHotInsert

Registry Value	Description
RebuildOnHotInsert	DWORD(32) = 0 (default): By default or when this value is created and cleared to 0, this feature is disabled.

	DWORD(32) = 1: When this value is created and set to 1, this feature is enabled and when all the system conditions are met, the driver will begin an auto-rebuild upon hot insertion of a supported disk.

19.5 SATA Asynchronous Notification

The location of the key to control SATA Asynchronous Notification functionality has been added.

The registry key by default is not populated in the registry, but 'AN' is enabled by default. In order to change the functionality in the driver, the registry key has to be created using the following setting:
Open the registry editor and add the following key:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Create a new DWORD (32) value as follows:

Controller0PhyXANEnable

Where 'X' represents the SATA port on which 'AN' is to be disabled or enabled.

Registry Value	Description
Controller0PhyXANEnable	DWORD(32) = 0 When set to 0, this feature is disabled.
	DWORD(32) = 1 (default) When set to 1, this feature is enabled.

19.6 Runtime D3 (RTD3)

RTD3 is part of Microsoft Power Management Framework* (PMF). The following registry setting is MSFT specific. A device may be "spun down" after being in Idle a designated time while the system is running (RTD3). When a value less than 60 seconds is chosen, it will default to 20 minutes. When using "MAXULONG" value, the registry key will be ignored and devices will not "spin down", which effectively "disables" RTD3.

KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Enum\SCSI\<Device ID>\<instance ID>\Device Parameters\Storport

Refre MSFT on how to set this registry key: http://support.microsoft.com/kb/241679/en-us

The REG_DWORD value is as follows:

MinimumIdleTimeoutInMS

Registry Key: Name	Description
MinimumIdleTimeoutInMS	Type: REG_DWORD(32) = {0 to MAXULONG-1}

	This value specifies the minimum amount of time the power framework must wait to power down a logical unit once it is at idle.
	REG_DWORD(32) = *MAXULONG*: The Default: MAXULONG, indicating unset. If the miniport provides no timeout value, the actual default value is 5 minutes, or, 5 * 60 * 1000.

19.7 Hybrid Hinting

Hybrid Hinting is enabled automatically without the need for user intervention. In the case that Hybrid Hinting needs to be disabled, a registry key can be created for that purpose. On Windows 8.1 platforms, a registry key is automatically added without user intervention.

19.7.1 Instructions to Disable Hybrid Hinting

Enter the registry editor and navigate to the following path:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Create a new DWORD (32) value as follows:

HybridHintDisabled

Registry Value	Description
HybridHintDisabled	DWORD(32) = 0 Leaves Hybrid Hinting enabled
	DWORD(32) = 1 Disables Hybrid Hinting

19.7.2 Hybrid Hint Reset

As of Win* 8.1, MSFT introduced hinting in their inbox driver. This leaves the hybrid threshold in states that causes decreased performance when the RST driver is installed AFTER Win* 8.1 OS is installed using MSFT inbox driver. To prevent this scenario, RST has provided the following registry key when an SSHD is detected on the system and Hybrid Hinting minimum requirements are met:

Name: **HybridHintReset**

Location:
KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Note: Upon the first reboot after the driver installation is complete, the registry key is written and the hybrid log reset by the driver. Once the registry key is written, it will remain throughout reboots and OS upgrades. If deleted manually, it will be rewritten automatically by the driver upon the next reboot.

19.7.3 Disable Hybrid Hinting During Hibernation

On systems using SSHDs with small NAND size and when hybrid shutdown is enabled in Microsoft OS, it is recommended to disable hybrid hinting during hibernation. This key can be added to

prevent longer shutdown time in some cases. In these cases, OEMs/ODMs may use this key to base their evaluation results.

The following registry key can be added to disable hybrid hinting during hibernation.

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Create a new DWORD (32) value as follows:

HiberFileHintDisable

Registry Value	Description
HiberFileHintDisable	DWORD(32) = 0 Enable hybrid hinting during hibernate (Default)
	DWORD(32) = 1 Disable hybrid hinting during hibernate

§ §

20 Power Savings with Intel® Rapid Storage Technology

20.1 Link Power Management (LPM)

Intel® Rapid Storage Technology implements the Link power management (LPM) feature described by the Serial ATA specification to overcome the power demand of a high-speed serial interface, SATA and providing the capability of SATA at the minimum power cost. LPM, when used in conjunction with a SATA hard drive that supports this feature, enables lower power consumption. LPM was initially enabled by default on mobile platforms starting with ICH6M with Intel® Matrix Storage Manager. Starting with ICH9R this feature has also been supported on desktop platforms with Intel® Matrix Storage Manager 7.5 release but not enabled by default.

Beginning with the Intel® Rapid Storage Technology 10.0 release, LPM support is enabled by default on both mobile and desktop platforms. OEM's who wish to modify the default settings for LPM on their platforms can follow the instructions in the section titled Instructions to disable/enable LPM.

20.1.1 Instructions to Disable/Enable LPM

After the system is setup with the OS and Intel® Rapid Storage Technology is installed, follow the below instructions to modify the default LPM support.

Note: Beginning with the Intel® Rapid Storage Technology 10.0 release, the registry keys are no longer populated in the Windows registry by default. The RST driver does not require the registry keys to be present to support the default settings.

1. 1. Go to Start->Run
2. 2. Type in RegEdit and press the Enter Key.
3. 3. Go to the below mentioned location to insert or configure the registry keys for LPM.

Note: *OEM's need to configure the LPM settings per SATA port. Ports are numbered starting with zero (refer to the desired platform EDS for the number of ports supported on your platform).*

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\iaStorA\Parameters\Device

4. Now add the following registry keys under the registry location mentioned in step3, if they are not available (These registry keys are not available by default, they can be added by using automated scripts, .reg files, executable utilities, etc.). If you find the below registry keys already available, you can modify the values for desired support. Values are modified on a port by port basis so modify all ports that you wish the changes to be supported on. **

Per-port Setting:

Replace the '**X**' with the SATA port number to independently control HIPM/DIPM per port.

Host Initiated Power Management:
DWORD: ***Controller0PhyXHIPM***
Value: 0 = disable, **1 = enable (default)**
-- (Old key, DWORD: LPM)

Device Initiated Power Management:
DWORD: ***Controller0PhyXDIPM***
Value: 0 = disable, **1 = enable (default)**
-- (Old key, DWORD: DIPM)

Configure HIPM to use partial or slumber when the drive is in a D0 ACPI device state
DWORD: ***Controller0PhyXLPMState***
Value: **0 = Partial (default)**, 1 = Slumber
-- (Old key, DWORD: LPMState)

Configure HIPM to use partial or slumber when the drive is in a D3 ACPI device state (Device receives a start_stop_unit request: e.g. HDD idle spin-down).
DWORD: ***Controller0PhyXLPMDstate***
Value: 0 = Partial, **1 = Slumber (default)**
-- (Old key, DWORD: LPMDState)

Controller-wide Setting:

This allows auto partial to slumber to be enabled. Actual setting of APS is controlled by the values below:
Auto Partial to Slumber:
DWORD: ***EnableAPS***
Value: 0 = disable, **1 = enable (default)**

***Warning:** If you edit the registry incorrectly, you can cause serious problems that may require you to reinstall your operating system. Intel does not guarantee that problems that are caused by editing the Registry incorrectly can be resolved.*

20.1.2 LPM Updates in 14.0 (APS, SIPM)

20.1.2.1 APS

The RST driver acknowledges the APS configuration as it is exposed by the system BIOS.

20.1.2.2 Software Initiated Power Management (SIPM)

By default RST DRIVER enables SIPM for each active SATA LINK.

20.1.2.2.1 SIPM Initiated PM Enabling Rules

The RST driver enables SIPM only for SIPM capable SATA ports, and ASP will be set to PARTIAL, and host SATA storage controller will report HIPM capability.

- The RST DRIVER enables SIPM only when ASP is set to PARTIAL
- By default the RST DRIVER enables SIPM on SIPM capable SATA ports, excluding eSATA ports and hot-pluggable SATA ports.
- SIPM for Hot-Pluggable SATA Ports: By default the RST DRIVER enables SIPM on SIPM capable SATA ports, excluding eSATA ports and hot-pluggable SATA ports
- SIPM for eSATA Ports: The RST DRIVER keeps SIPM disabled on SATA ports which are defined by the platform BIOS as eSATA ports
- SIPM Capable SATA Ports: The RST DRIVER enables SIPM only for SATA ports which connected devices are reporting HIPM capability
- SIPM will be enabled when the host is reporting HIPM capability

- SIPM will only be enabled for SATA devices which have APS disabled

20.1.2.2.2 SIPM Timeout Configuration

The RST DRIVER reads the value of SIPM timeout from the Windows registry. In case the Windows registry key is not present the RST DRIVER will use default values for the timeout. The SIPM timeout will be configurable for each SATA device independently

- The RST driver allows configuring the SIPM timeout for each SATA device independently
- The RST driver sets the SIPM timeout to 0 (zero) in order to disable SIPM for the given SATA port
- The RST driver reads the SIPM timeout value from the Windows registry
- The RST driver uses default values for SIPM timeout when the Windows registry key is not present. The amount of time to elapse before the RST driver puts the link into SLUMBER is as follows:

Windows OS Version	SIPM Default Timeout (milliseconds)
Windows* 8 and later	25
Windows* 7 and earlier	100

20.2 Runtime D3 (RTD3)

Intel® RST now supports Microsoft® Runtime Power Management (RTPM) framework. RTPM allows Windows® to manage the ACTIVE/IDLE power states for discrete target devices, such as hard disks, in an effort to reduce power consumption for mobile platforms.

RTD3 refers to the ability to completely remove power from devices (D3cold) during long idle periods, while the system remains in S0.

RTD3 adheres to the following requirements:

1. OS is Windows* 8 or newer.
2. The SATA AHCI controller must be in AHCI mode.
3. Only HDD's, SSD's, and SSHD's attached to the AHCI controller on ports that are RTD3 capable are managed by RTPM.
4. Devices connected to hot-pluggable ports cannot be managed under RTPM.
5. Devices connected to external ports cannot be managed under RTPM.
6. Devices connected to ports with a mechanical presence switch cannot be managed under RTPM.
7. Devices connected to a Remapped PCIe port cannot be managed under RTPM.

Note: For system's that support RTD3, but are shipped with RST in 'RAID Ready' mode, it is suggested that RTD3 be disabled in the BIOS. This is to prevent end users from creating RAID volumes while RTD3 is enabled.

20.2.1 Adapter RTD3 Support

Current releases of Intel® Rapid Storage Technology do not support Adapter RTD3. This means that only the RTD3 supported devices will be placed into D3hot/D3cold. The Intel® RST driver does not support placing the SATA controller into a D3hot / D3cold state.

For additional information on configuration and usage of this feature, refer to its documentation CDI/IBP document # 516865.

20.3 New for RST 14.0 Release

20.3.1 RTD3 Support - RAID HDD/SSD/SSHDs Unit Support

The RST DRIVER will register SATA RAID ARRAY to Windows Runtime Power Management framework as a single FUNCTIONAL HARDWARE UNIT.

20.3.2 RTD3 Support - RAID w Mixed RTD3 Capable/Non-Capable Ports

RST Driver allows the creation of RAID volumes that utilize both RTD3 capable ports and non-RTD3 capable ports in the same volume. When the RST Driver receives RTD3 events for the RAID Volume, it will RTD3 the ports that are capable of RTD3.

20.3.3 RTD3 Support - RRT

The RST DRIVER shall NOT support RTD3 of SATA ports that are part of an RRT RAID Volume. Ports utilized in an RRT Volume will NOT have power removed or restored by RTD3.

20.3.4 RTD3 Support - SRT

The RST DRIVER will opt-out from RTD3 requests for SRT CACHE DEVICEs.

20.3.5 RTD3 Support - Hot Spares

The RST Driver does NOT support RTD3 of RST RAID Global Hot Spares.

20.3.6 RTD3 Support – Migrations and Rebuilds

The RST Driver will prevent the Windows OS from entering RTD3 for RAID Volumes that are currently under-going internal RAID rebuild or migration operations.
Once the operation is complete, the RST Driver will allow RTD3 activity for the RAID Volume to continue.

20.4 DEVSLP

Beginning with Intel® Rapid Storage Technology 12.5 release, support for a new SATA link power state was introduced, device sleep (DEVSLP). DEVSLP is a fourth and lowest link power state, coming after Slumber. The link will enter DEVSLP when all current IO have completed, the link is in the Slumber state and the DEVSLP idle timer has expired granting permission for the SATA

controller to assert the DEVSLP signal. BIOS is responsible for configuring and enabling DEVSLP. The driver configurable settings for DEVSLP may be found in the section titled SATA Device Sleep (DEVSLP) Settings. For additional information on configuration and usage of this feature, refer to the document CDI/IBP document # 516865.

Intel® RST supports DEVSLP for reduced power during long idle periods such as when the system is in InstantGo*. When DEVSLP is enabled, the Intel® RST driver will support InstantGo* when requested by the OS on pass through devices on the SATA ports that support DEVSLP. InstantGo* is only supported on Windows* 8.1 and newer.

The following recommendation for Intel® Rapid Storage Technology DEVSLP Idle Timeout is taken from Section 4.6 of the Intel "Ultrabook™ Storage Power Management Recommendations White Papers" CDI/IBL #528428

When the system is in InstantGo*, the DEVSLP idle timeout should be set to maximize power savings. Because the I/O pattern while in InstantGo* is not deterministic, the DEVSLP idle time cannot be set to an arbitrarily low value, or else the power consumed by entering and exiting DEVLSP may be (on the average) greater than the power consumed by remaining in the next higher power state (Slumber). Instead, the idle timeout must be set to a value that delivers the best average power consumption. To achieve the best average power across a variety of configurations, the DEVSLP idle timeout should be set to equal the DEVSLP transition energy recoup time. The DEVSLP recoup time is the time in the next higher power state (Slumber) that consumes that same amount of power as entering and exiting DEVLSP. Using the recoup time ensures that (on the average) the device is not placed in DEVSLP before the energy consumed by the transition can be recouped.

Recoup time for a device to enter and then exit the next lowest power state relative to the device remaining in and then exiting from its current power state is calculated as follows:

$$recoup\ time = \frac{(next\ state\ entry\ energy + next\ state\ exit\ energy - current\ state\ exit\ energy)}{(current\ instate\ power - next\ state\ instate\ power)}$$

…where entry and exit transition energy is calculated as follows:

$$state\ transition\ energy\ =\ transition\ time\ *\ average\ power\ during\ transition$$

In the following example, DEVSLP recoup time is calculated for a hypothetical device with the following characteristics:

Table 20-1. DEVSLP Recoup Time

State/State Transition	In-State or Transition Time (S)	In-State or Transition Average Power (W)	In-State or Transition Energy (J)
Slumber	(1)	0.05	
Slumber resume transition	0.001	1.3	0.0013
DEVSLP entry transition	0.000	0	0.0000
DEVSLP	(2)	0.005	
DEVSLP resume transition	0.100	1.3	0.1300

1. Maximum time in Slumber is determined by the DEVSLP recoup time
2. Maximum time in DEVSLP is determined by the RTD3 recoup time (see next section)

Following the general case above, DEVSLP recoup time is calculated as follows:

$$DEVSLP\ recoup\ tim = \frac{DEVSLP\ entry\ energy + DEVSLP\ resume\ energy - Slumber\ resume\ energy}{Slumber\ instate\ power - DEVSLP\ instate\ power}$$

Substituting values from the table above:

$$DEVSLP\ recoup\ time = \frac{0\ J + .13\ J - .0013\ J}{.05\ W - .005\ W} = 2.86\ seconds$$

A device with these characteristics may stay in Slumber for 2.86 seconds and use the same power as would be consumed by transitioning to DEVSLP. Therefore, the recoup time is defined as 2.86 seconds, and the DEVSLP idle time-out when in InstantGo* should be set to 2.86 seconds.

For additional information on configuration and usage of these parameters, refer to the document CDI/IBP document # 516865

20.4.1 DEVSLP Registry Key Setting

Per-port Settings:

This allows the OEM to customize the DEVSLP Idle Time Out value for InstantGo* enabled systems. When the OS enters InstantGo*, the driver configures the SATA controller to enter DEVSLP sooner to save power. This setting will temporarily override the BIOS configured value for the duration of the InstantGo* period. Upon exiting InstantGo*, the driver will restore the value the BIOS originally programmed:
Replace the '**X**' with the SATA port number to independently control the DEVSLP timeout value while in InstantGo*.

DEVSLP Device Idle Time Out Small in milliseconds:

DWORD: ***Controller0PhyXDevSlpDITOsmall***
Value: 0x1– 0x3FF0, default **0x0BB8** (values need to be entered as hex, equivalent decimal values are 1 to 16368 ms, default 3000 ms)

20.5 DevSleep Tool

Beginning with Intel® RST versions 13.5 and newer, OEMs can pre-configure their platforms with the DevSleep tool. This tool is a Windows/WinPE or Dos-based command line interface utility for configuring DEVSLP registry keys for fine tuning Low Power DEVSLP enabled devices.

At boot time, the RST DRIVER shall read the configured registry key values and use them as overrides on a per-device basis.

Each of the following registry keys have the characteristic:
String Value: "<product id> <timeout>"
Example: "xx740ADFD?00NLR1* 3000"
<product id>: The First part is the device product id to match
Special characters can be used in product id:
'?' matches any single character
'*' matches any length of characters
"* 3000": applies a 3000ms value to all disks.
< timeout >: The value of the DEVSLP time variable.

- CsDeviceSleepIdleTimeoutInMS
- DeviceSleepIdleTimeoutInMS
- DeviceSleepExitTimeoutInMS
- MinimumDeviceSleepAssertionTimeInMS

Multiple "product id timeout" pairs can be placed in the registry key. Each pair is separated by a null delimiter.

The RST DRIVER will first look for the per-device registry key and if a device match is found it shall use the value indicated. If per-device registry key does not contain a match for any attached device, then the driver shall use the per-port specific registry key if present. This requirement shall not modify the behavior of any per port registry key.

Per-Device Key	Per-Device String Match	Port Key	Results
Present	Found	Present	Use Per-Device
Present	Found	Missing	Use Per-Device
Present	Missing	Present	Use Port
Present	Missing	Missing	Use Default
Missing	Missing	Present	Use Port
Missing	Missing	Missing	Use Default

20.5.1 CsDeviceSleepIdleTimeoutInMS

Path: HKLM\System\CurrentControlSet\Services\iaStorA\Parameters\Device
Key:
CsDeviceSleepIdleTimeoutInMS

This registry key is the Device Sleep idle timeout (DITOActual) to use when the system is in connected standby. DITOActual = (DITO * (DM+1)

Total DevSlp Idle Timeout is the total amount of time in ms that the host bus adapter will wait after the port is idle before raising the DevSlp signal, max=16368.

If this registry key is not present, the RST DRIVER shall check for the per-port specific registry key "DevSlpDITOsmall". This registry key will not take precedence over the registry setting of "DevSlpDITOsmall" if already present.

20.5.2 DeviceSleepIdleTimeoutInMS

Path: HKLM\System\CurrentControlSet\Services\iaStorA\Parameters\Device
Key:
DeviceSleepIdleTimeoutInMS

This registry key is the Device Sleep idle timeout (DITOActual) to use when the system is ***not*** in Connected Standby (CS). Note: this registry shall apply to both CS and non-CS platforms. The < timeout> is the value of the DEVSLP idle timeout to use when the system is ***not*** in Connected Standby, in milliseconds (decimal value).

20.5.3 DeviceSleepExitTimeoutInMS

Path: HKLM\System\CurrentControlSet\Services\iaStorA\Parameters\Device
Key:
DeviceSleepExitTimeoutInMS

This registry key is the Device Sleep Exit timeout (PxDEVSLP.DETO). The < timeout > value is the DEVSLP exit timeout in milliseconds (decimal value).

20.5.4 MinimumDeviceSleepAssertionTimeInMS

Path: HKLM\System\CurrentControlSet\Services\iaStorA\Parameters\Device
Key:
MinimumDeviceSleepAssertionTimeInMS

This registry key is the minimum amount of time, in ms, that the HBA must assert the DEVSLP signal before it may be de-asserted; Minimum Device Sleep Assertion time (PxDEVSLP.MDAT). The nominal value is 10ms and the minimum is 1ms depending on device identification information.

The < timeout > value is the minimum DEVSLP assertion time in milliseconds (decimal value).

20.5.5 DevSleep Tool Usage

USAGE:
Create Options:
Create Usage:
Creates the new registry keys and populates them with default values
create [--key x] [--inline]
Create Examples:
-C
--create
--create --inline
--create --key CsDeviceSleepIdleTimeoutInMS
--create --key DeviceSleepExitTimeoutInMS --inline
--create --help
Export Options:
Export Usage:
Exports the Dev Sleep registry keys to a distributable .reg file
-export
Export Examples:
E
-export
--export --help
List Options:
List Usage:
Lists all devices and values in the registry
--list [--key x]
List Examples:
-L
--list
--list --key MinimumDeviceSleepAssertionTimeInMS
--list --help
Modify Options:
Modify Usage:
Modifies the reg key
--modify --index z --value y [--key x] [Product ID]
Modify Examples:
-M --index 3 --value 10
--modify --index 0 --value 3 --key DeviceSleepIdleTimeoutInMS
-M --index 1 --value 7 --key CsDeviceSleepIdleTimeoutInMS newproductid

--modify --index 1 --value 7 productid with spaces for all reg keys
--modify --help

Add Options:

Add Usage:
Adds a new registry key
--add --value y Product ID
Add Examples:
-A --value 10 newproductid
--add --value 3 --key DeviceSleepIdleTimeoutInMS newproductid
-A --value 7 productid with spaces for all reg keys
--add --help

Import Options:

Import Usage:
Imports the Dev Sleep registry key to a specified OS image
--import --driveLetter c

Import Examples:
-I --driveLetter C
--import --driveLetter C
--import --help

Delete Options:

Delete Usage:
Deletes an existing registry key
--delete --index z [--key x]

Delete Examples:
-D
--delete --index 3
--delete --index 1 --key MinimumDeviceSleepAssertionTimeInMS
--delete –help

20.6 L1.2 Support

Starting with RST version 13.5, L1.2 will be supported for PCIe SSD devices. L1.2 is the lowest, non-zero power state that uses a 2 step resume process for better responsiveness and power savings. RST will support L1.2 when it is enabled on the platform and when L1.2 is enabled on the PCIe Storage device. Because this feature is controlled primarily by the system BIOS, refer to your BIOS reference guide to for correct system setup.

20.7 InstantGo* Device Notification Support

Note: ***Intel® Rapid Storage Technology support for InstantGo* is limited to platforms that support this feature. InstantGo* is a Microsoft feature and for more details, consult Microsoft.***

Beginning with Intel® Rapid Storage Technology 13.0, the Intel® RST driver also supports InstantGo* Notification on all InstantGo* Notification capable devices connected to the AHCI controller. The Intel® RST driver will notify devices when the system is entering/exiting InstantGo*. This allows supported devices to change policy and be more aggressive in internal power management and power savings.

20.7.1 Requirements

- Devices must support the Advanced Power Management feature (APM) defined in the ATA* standard (ACS-3) and report support for APM levels.*
    - IDENTIFY DEVICE data Word 159 set to 0xA5A5 in the Vendor Specific area.**
- Platform Hardware and Devices must support the DEVSLP Feature.
    - Device supports the Device Sleep feature (per ATA IDENTIFY DEVICE command) IDENTIFY DEVICE data word 78 bit 8 is set to '1'.(i.e., resume from Device Sleep using COMWAKE).
- Devices must support DevSleep_to_ReducedPwrState (as indicated in Identify Device data).
- Supported on pass-thru devices (non-RAID member).

Intel recommendations for IHVs to support devices' INSTANTGO Notification requirements and APM levels can be found in "Ultrabook Storage Power Management Recommendations White Paper" on CDI/IBL #528428.
**InstantGo* Notification feature is currently supported on Intel Dalecrest (530) devices.

20.7.2 Detail Description

Intel® Rapid Storage Technology driver customizes InstantGo* Device Notification on a per device basis when the sytem is entering / exiting InstantGo*.
Once the Intel® RST driver receives notification from the OS that the system is going to enter or exit InstantGo*, Intel® RST driver will notify the device by using the APM (Advanced Power Management) mechanism defined in the ACS (ATA Command Set).
The Intel® RST driver will use a SET FEATURE command to send a hint of power/performance balance to the device. The hint value is 01 – FEh.
Values:

- FEh - max performance at the expense of power
- 01h - max power savings at the expense of performance
- 80h – defined for HDDs as max power savings, but cannot spin down, lower than 80h and the device is allowed to self spin down
- All other values are vendor specific

Intel® RST uses 10h as the default value for entering InstantGo*, and 80h is the default value Intel® RST uses for exiting InstantGo*. Values are customizable by using registry keys.

Note: Additionally, a device should be prepared for a power loss after any completed STANDBY IMMEDIATE command (see ACS-3) regardless of the link power state entered between the completion of a STANDBY IMMEDIATE command and the loss of power. This includes Active, Partial, Slumber and DevSleep link power states.

20.7.3 Registry Settings

InstantGo* Device Notification is enabled by default. The following registry keys can be used to disable InstantGo* Device Notification and/or to change default values on a per device basis.

Open the registry editor and navigate to this path:
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\ Device
Add the following key to disable InstantGo* Notification for device on port '***X***':

DWORD: ***Controller0PhyXCsDeviceNotification***

Value: 0x0– 0x01, default **0x01** (1= enabled / 0=disabled)

The following Registry key will allow customizable APM levels to set the device to when the system enters and/or exits InstantGo*. This will only be done if InstantGo* Notification is enabled.

Open the registry editor and navigate to this path:

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStorA\Parameters\Device

Add the following registry key to customize Entry into InstantGo* notification :

DWORD: ***Controller0PhyXEnterCSApmLevel***

Values: **As below *(default is 10h)***

Add the following registry key to customize Exit from InstantGo* notification:

DWORD: ***Controller0PhyXExitCSApmLevel***

Values: **As below *(default is 80h)***

VALUES: **APM levels interpreted as follows:**

- FEh – Maximum performance mode
- 80h – Minimum power management without standby (e.g. balance between power savings and performance)
- 10h – High performance bursts, quick to low power (e.g. Windows* 8.1 InstantGo*)
- 01h – Minimum power consumption with standby

The requested APM levels persist across resets, but not power cycles. Intel® RST 13.0, and newer, supports this capability and will set the appropriate APM levels after power is applied to the devices (either power on or after resume from RTD3).

InstantGo is a MSFT feature and was formerly referred to as "Connected Standby".

20.8 New in 14.5 Release

20.8.1 Connected Standby Power State Support for SSHD

The RST DRIVER allows CONNECTED STANDBY power state for SSHD under following conditions:

- Platform supports CONNECTED STANDBY power state,
- SSHD is used as BOOT VOLUME,
- SSHD is used in pass-through mode.

20.8.1.1 Hybrid Type Device Criteria

The RST DRIVER classifies a drive as HYBRID TYPE when the drive reports HYBRID HINTING, PUIS and APM capabilities.

20.8.1.2 Spin-Down on Entering Connected Standby

The RST DRIVER uses APM to spin down of the SSHD rotational media when entering CONNECTED STANDBY power state.

20.8.1.3 Use Hybrid Hinting when in Connected Standby

The RST DRIVER uses HYBRID HINTING to redirect I/O requests to SSHD NVCACHE when in CONNECTED STANDBY power state.

20.8.1.4 Use APM to Control Spindle

The RST DRIVER uses APM to control the spindle. The APM will automatically spin-up the rotational media part of the drive for time of a "read cache miss" event.

20.8.1.5 Restore APM on Exiting from Connected Standby

The RST DRIVER restores APM to its original value upon exiting CONNECTED STANDBY power state.

20.8.1.6 PUIS Capabilities for Connected Standby Power State

The RST DRIVER enables PUIS prior entering CONNECTED STANDBY to prevent spin-up of the rotational media part after RTD3 transition requests from the operating system.

20.8.2 CONNECTED STANDBY Power Model

The RST DRIVER utilizes the CONNECTED STANDBY power model for low-power capable platforms.

20.8.3 Adaptive D3 Idle Timeout

The RST DRIVER enables the adaptive D3 idle timeout feature for SATA rotating media.

20.8.3.1 Adaptive D3 Idle Timeout for Hard Disk Drives

The RST DRIVER enables the adaptive D3 idle timeout feature for SATA HDDs.

20.8.3.2 Adaptive D3 Idle Timeout for RAID

The RST DRIVER enables the adaptive D3 idle timeout feature for homogeneous RAID volumes comprise SATA rotating media.

Adaptive D3 Idle Timeout for Matrix RAID

The RST DRIVER disables the adaptive D3 idle timeout feature for MATRIX RAID volumes.

Adaptive D3 Idle Timeout for mixed RAID

The RST DRIVER always disables the adaptive D3 idle timeout feature for MIXED RAID volumes.

20.8.3.3 Adaptive D3 Idle Timeout for Hybrid Hard Drives

The RST DRIVER always enables the adaptive D3 idle timeout feature for SSHDs (HHD).

20.8.3.4 Adaptive D3 Idle Timeout for Non-Rotating Media

The RST DRIVER always disables the adaptive D3 idle timeout feature for all types of non-rotating media.

Adaptive D3 Idle Timeout for PCIe Remapped SSDs

The RST DRIVER always disables the adaptive D3 idle timeout feature for PCIe Remapped SSDs.

Adaptive D3 Idle Timeout for SATA SSDs

The RST DRIVER always disables the adaptive D3 idle timeout feature for SATA SSDs.

20.8.3.5 Adaptive D3 Idle Timeout for ATAPI Devices

The RST DRIVER always disables the adaptive D3 idle timeout for ATAPI devices.

Adaptive D3 Idle Timeout for Tape Devices

The RST DRIVER always disables the adaptive D3 idle timeout feature for tape drives.

Adaptive D3 Idle Timeout for ODDs

The RST DRIVER always disables the adaptive D3 idle timeout feature for ODDs.

Adaptive D3 Idle Timeout for ZPODDs

The RST DRIVER always disables the adaptive D3 idle timeout feature for ZPODDs.

20.8.4 Connected Standby Power Model Support for SRT

The RST DRIVER enables entering into CONNECTED STANDBY power state when a CACHE VOLUME is configured.

20.8.4.1 Allowed Configuration

Minimum 16GB of cache and SRT must be configured in MAXIMIZED MODE.

20.8.4.2 Storage Device Criteria

The CACHE DEVICE and the BACKING STORAGE both must be SATA storage devices.

20.8.4.3 Entering Connected Standby Power State

The RST DRIVER enters into CONNECTED STANDBY power state on explicit notification received from Windows operating system.

Allow Entering Connected Standby for Cache Device

The RST DRIVER allows entering into CONNECTED STANDBY power state under following conditions:

- Platform is supporting CONNECTED STANDBY power state,
- CACHE DEVICE is used to accelerate a BOOT VOLUME.

SRT Enhanced Mode Constraint on MODERN STANDBY Platforms

The RST DRIVER will not D3 the BACKING STORAGE when the CACHE VOLUME is configured in ENHANCED MODE.

20.8.4.4 Exiting Connected Standby for Backing Storage

The RST DRIVER will D0 the BACKING STORAGE upon exiting the CONNECTED STANDBY power state.

20.8.4.5 Waking up Backing Storage

The RST DRIVER prevents the BACKING STORAGE from entering D0 by directing I/O requests to CACHE DEVICE.

Waking up Backing Storage on Cache Miss or Cache Clean

The RST DRIVER may D0 the BACKING STORAGE while in CONNECTED STANDBY power state, regardless of power source, for time of occurred events:

- Cache miss during I/O READ request,
- Cache cleaning due to cache volume full.

20.8.4.6 Aggressive D3 and Semi-Aggressive D3 for Backing Storage

AGGRESSIVE D3 and SEMI-AGGRESSIVE D3 of BACKING STORAGE will only be available when SRT is configured in MAXIMIZED MODE.

Registry Support for Aggressive D3 and Semi-aggressive D3 Inactivity Timeouts

The RST DRIVER reads configuration of I/O inactivity timeouts for AGGRESSIVE D3 and SEMI-AGGRESSIVE D3 from Windows Registry.

Default Inactivity Timeout for AGGRESSIVE D3

The RST DRIVER waits not less than 2 seconds for completion of an I/O request to BACKING STORAGE before attempting to D3 the BACKING STORAGE.

Default Inactivity Timeout for SEMI-AGGRESSIVE D3

The RST DRIVER waits not less than 30 seconds for completion of an I/O request to BACKING STORAGE before attempting to D3 the BACKING STORAGE.

20.8.4.7 Entering Resiliency Phase on AC Power

The RST DRIVER will SEMI-AGGRESSIVELY D3 the BACKINGS TORAGE when entering the RESILIENCY PHASE of CONNECTED STANDBY when the system is on AC power.

20.8.4.8 Entering Resilience Phase on Battery Power

The RST DRIVER will AGRESIVELLY D3 the BACKING STORAGE when entering the RESILIENCY PHASE of CONNECTED STANDBY when system is on battery power.

20.8.5 SATA Link Power Management Support

The RST DRIVER implements SATA Link Power Management as specified in the Power State Transition section of the SATA AHCI Specification.

20.8.5.1 SIPM Support for MODERN STANDBY Power Model

The RST DRIVER disables SIPM and sets ASP to SLUMBER upon entering into MODERN STANDBY, and re-enables SIPM and sets ASP to PARTIAL upon existing from MODERN STANDBY

SIPM Disabled when in MODERN STANDBY

The RST DRIVER disables SIPM on entering to MODERN STANDBY.

SIPM Re-enabled Upon Exiting MODERN STANDBY

The SIPM is re-enabled upon exiting from MODERN STANDBY.

ASP Configuration when in MODERN STANDBY

The RST DRIVER keeps ASP set to SLUMBER when the platform is in MODERN STANDBY.

ASP Configuration when not in MODERN STANDBY

The RST DRIVER keeps ASP set to PARTIAL when the platform is not in MODERN STANDBY.

§ §

21 Legacy RAID Option ROM and Utilities

There is a unified RAID Option ROM (.ffs or .bin) and RAID Utilities (RCfgSata.exe and RCmpSata.exe) beginning with the Intel® Rapid Storage Technology 9.6 and later release versions. This unified OROM package is supported on the platforms identified below and all later platform releases when enabled for RAID. With this RAID Option ROM, these platforms now support all RAID functionality based on the hardware capability.

General Requirements

Hardware	Only enabled on specific SKUs of the Intel® 8 / 9 Series chipset platforms: • H87, H97, HM87, HM97 • Q87, QM87, • Z87, Z97 • C226 • LPT-LP Premium • Broadwell U/Y Premium With one of the following Intel® Processors installed: - Intel® Core™ Branded (i3, i5, and i7 processor families) - Intel® Xeon® Processor family
Operating System	All supported Operating Systems for this release 131

§ §

22 HDD Password Support with RAID Volumes

Intel® Rapid Storage Technology supports password protected HDDs to be RAID array member disks and pass-thru disks. The product will rely on the BIOS implementing for most of the ATA Security support. There is a whitepaper available called "Implementing Intel® Matrix Storage Manager Compatible Support for ATA Security in BIOS" available on CDI that describes the necessary BIOS design for compatibility with the Intel Rapid Storage Technology. Rapid Storage Technology product will handle the RAID and hot-plug related behavior with regards to password protected disks.

22.1 HDD Password Use Cases

If at least one unlocked member disk and one locked member disk (with relevant data for the RAID array) are connected, then the RST Driver will reset the "offline disk member" status prior to a boot-time enumeration in the case that a member disk is unlocked by entering the correct password. The Intel® RST UI will display the locked disk as locked, and any member disks as offline.

Accelerated volumes containing a locked member disk will return to a normal online state upon transitioning from S4 to S1 and entering the correct password to unlock the volume/disk.

Scenario	Action	Result	Comments
RAID1 Volume Disk 1 – Locked Disk 2 - Unlocked Volume – Locked (Both disks have relevant data)	Remove Disk 1 (locked disk)	Volume becomes unlocked and Degraded. User can rebuild volume unto a new unlocked disk.	The user had authority to access Disk 2 which has the same data as Disk 1, by removing the locked drive the user can access Disk 2.
RAID1 Volume Disk 1 – Locked Disk 2 – Unlocked Volume – Degraded Disk 1 has old data and caused the volume to go Degraded.	None	N/A	The user has access to Disk 2 because the data on Disk 1 is old and irrelevant.
RAID5 Volume Disk 1 – Locked Disk 2 – Unlocked Disk 3 – Unlocked Volume – Locked (All disks have relevant data to Volume)	Remove Disk 1 (locked disk)	Volume becomes unlocked and Degraded. User can rebuild volume onto a new unlocked disk.	The user had authority to access Disk 2 and Disk 3 which has all the data needed to rebuild the volume, by removing the locked drive the user can access Disk 2 and Disk 3 as a Degraded Volume.

Intel® RRT Volume Master Disk – Locked Recovery Disk - Locked (external port docking station) Volume – Locked (Both disks have relevant data)	User connects laptop to docking station and unlocks Recovery disk and Master Disk and boots. Then user takes the laptop from the docking station and leaves the external drive connected to power	The recovery drive can be connected to a new laptop and the information can be used to rebuild an Intel® RRT volume if the power was maintained, because the drive is still in an unlocked state.	Similar situation to a user leaving a laptop unlocked and unattended.

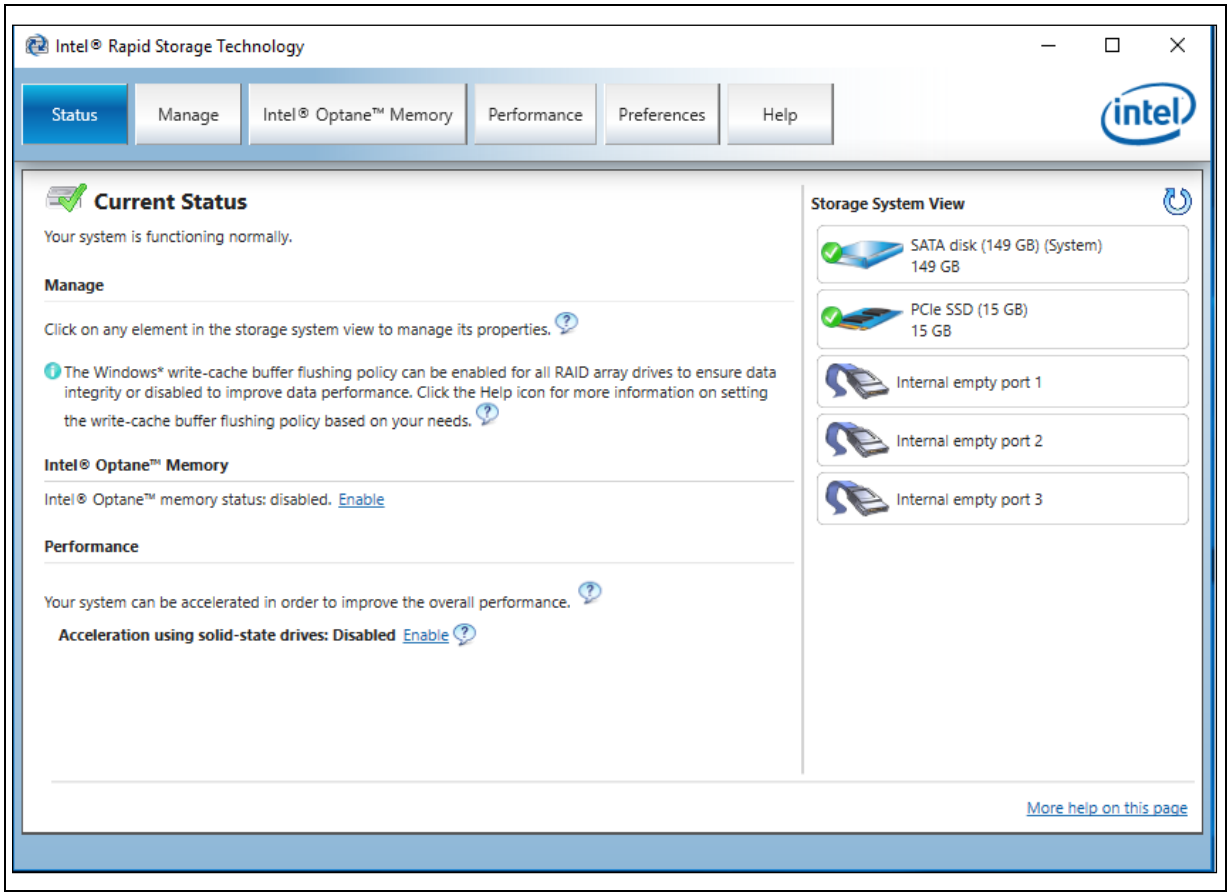
22.2 Unlocking Password Protected Disks

The ability to unlock password protected drives from within the Intel® RST GUI during OS runtime has been removed beginning with ***RST production release version 11.6.0.1030***. Users will only be able to unlock password protected drives from within the system BIOS. Consult the system manufacturer for instructions on BIOS use for this feature.

§ §

23 Intel® Rapid Storage Technology UI

Note: The Intel® Rapid Storage Technology UI is not required to be installed when the storage subsystem is operating in AHCI mode. The main benefit of the Intel® RST UI is in its management and monitoring of the Intel® RST RAID storage subsystem and the management and configuration of cache accelerated volumes.



23.1 Introduction

The Intel® Rapid Storage Technology UI is a Windows*-based application that provides users monitoring and management capabilities for the Intel® RST storage subsystem. It offers a wide range of monitoring and management activities for the Intel® RST RAID subsystem (***In AHCI mode there are no management or monitoring capabilities offered by the UI application***).

23.1.1 Getting Started

The Intel® Rapid Storage Technology software package provides high-performance SATA AHCI and SATA RAID capabilities for supported operating systems.

The Intel® Rapid Storage Technology (RST) UI requires the Microsoft .NET 4.5 framework beginning with Intel® RST 13.0 release. For prior releases, the RST UI connects and interoperates with the Microsoft .NET 3.0, 3.5, and 4.0 framework.

Refer to the System Requirements and the online user's manual to set up your system's configuration and feature support level. You can also review the Readme file installed with this software or visit Intel's online support to learn more about the full system requirements and RAID BIOS configuration.

RAID enabled systems
Redundant Array of Independent Disks (RAID) refers to multiple independent disks combined to form one logical drive. The main objective of this technology is to improve storage system performance, data protection, and increase fault-tolerance.

This technology provides support for the following features:

- **Intel® Rapid Recover Technology**
  This technology provides full data redundancy by copying data from a designated source drive (i.e., master disk) to a designated destination drive (i.e., recovery disk). Data updates of recovery volumes can be continuous or on request.
- **Intel® Rapid Storage Technology RAID**
  This technology provides the ability to create RAID 0, RAID 1, RAID 5, and RAID 10 volumes on desktop and mobile platforms. Data is distributed across two or more disks to provide data redundancy (exception of RAID0) or to enhance data storage performance.
- **Intel® Matrix RAID Technology**
  This technology allows two independent RAID volumes to be created on a single array. The first volume occupies part of the array, leaving space for the second volume. The array may consist of two to six SATA disks depending on the volume types.
- **Hot plug**
  Also referred to as hot swap, this feature allows SATA disks to be removed or inserted while the computer is turned on and the operating system is running. As an example, hot plugging may be used to replace a failed external disk.
- **Intel® Smart Response Technology caching**
  This feature allows you to use a non-system solid state disk and configure it as a non-volatile intelligent cache device in order to accelerate a disk or volume that is part of the storage system. This configuration helps improve the overall system performance.
- **Volume migration**
  This feature provides support for converting system data into a high-performance or protection RAID configuration.
- **Volume size increase**
  This feature allows you to increase the data storage capacity of a volume by using 100% of the available array space or by adding one or more SATA disk to an existing volume.
- **Password-protected disks**
  This feature provides high-level security and protection for the data on your disks with a password, denying access from any unauthorized user.
  Additional features and technology supported by the driver although not directly accessible via the Intel® RST UI:
- **TRIM**
  This feature provides support for all solid state disks (SSDs) in your storage system that meet the ATA-8 protocol requirements and are not part of an array. This feature optimizes write operations, helps devices reduce wear, and maintains unused storage area on devices as large as possible.
  Beginning with the Intel® 7 Series chipset the driver supports TRIM on SSDs in a RAID 0 configuration.
- **ODD power optimization**
  This feature allows an unused optical disk drive (ODD) to be automatically powered off when media such as a compact disk, a DVD, or Blu-ray disk are not present in the drive

and the tray** is closed. The ODD is powered back on by the operating system or user interaction with the device, including when the eject button is pressed. ODD power optimization is particularly valuable for mobile computers as battery life is negatively affected when the ODD is powered on and in an idle state. This feature is only supported on the following system configurations: Intel® 6 Series Chipset or later, compatible motherboards, and compatible ODDs. For more information about compatibility requirements, refer to the SATA specifications available at www.sata-io.org

**For slot-loadable drives, the tray condition does not apply.

- **Native command queuing**
  A feature that allows SATA disks to accept more than one command at a time. When used in conjunction with one or more disks that support NCQ, storage performance is increased on random workloads by allowing the disk to internally optimize the order of commands.
- **Hybrid Hinting**
  This feature supports the use of Solid-State Hybrid Drives (SSHD). SSHD's are hard disks that contain Flash memory for use as a cache to store frequently accessed data. The driver provides hints to the SSHD to notify the drive which data would be best to store in the data cache.

  Disks of more than two terabytes
  This feature provides support for hard disks and solid state disks with a capacity greater than 2 TB that are reported as pass-through devices (available) or used in a RAID configuration. In addition, booting from a system disk greater than 2 TB is allowed as long as the version of the option ROM in your system supports this feature.

Note: : If a source disk is greater than 2TB and using the MBR partitioning scheme, the application will not allow data preservation in order to create a volume. Instead, a new volume will be created with no partition on it. Also, if the operating system is Windows* XP, capacity expansion operations will not be allowed for volume sizes equal or greater than 2TB.

AHCI-enabled systems

Advanced Host Controller Interface (AHCI) is an interface specification that automatically allows the storage driver to enable advanced SATA features, such as Native Command Queuing and Native Hot Plug, on the SATA disks connected to your computer. The following features are supported on AHCI-enabled systems:

- **Native command queuing**
- **Hot plug**
- **Disks of more than two terabytes (if that size is supported by the RST UEFI pre-OS driver or legacy OptionROM)**
- **Password-protected disks**
- **ODD power optimization (Microsoft Windows Vista* and higher)**
- **Dynamic Storage Acceleration**
- **Hybrid Hinting**

23.1.2 Understanding the Application

The Intel® Rapid Storage Technology application allows you to optimize and maintain a healthy storage system by creating volumes, customizing performance settings and managing storage system elements. This section provides you with a general overview of a storage system configuration and an individual review of all the areas contained in this application.

23.1.2.1 Storage System Configuration

The storage system combines hardware capabilities with RAID technology to provide flexible data storage units on your computer. Each data storage unit, or RAID configuration, consists of three elements that include physical SATA disks, one or two volumes, and one array. When at least one volume is present on the system, these elements are represented in the storage system view of the Status and Manage areas.

In this section, we describe each of these RAID configuration elements and explain how they relate to each other.

- **Array**
  An array is a collection of two or more SATA disks in a RAID configuration and is the highest element in the hierarchy of a storage system. Once a volume is created, the disks you used to create that volume form an array. Refer to the Creating Additional Volumes topic for details on how you can create two volumes across the same disks. An array can include one or two RAID volumes if the hardware allows it.
- **Volume**
  A volume is the storage area on two or more disks whose type dictates the configuration of the data stored. If you created a volume for data protection, then your storage system may include a RAID 1 volume spanning two SATA disks, which mirrors data on each disk.
- **Disks**
  A disk (i.e., hard disk or hard disk drive) physically stores data and allows read/write data access. If a disk is used to create a volume, it becomes an array disk because it has been grouped with other disks to form an array.

The storage system can also include ATAPI devices, which cannot be used to create a volume. They are a mass storage device with a parallel interface, such as CD-ROM, DVD/Blu-ray disc, or tape drive.

23.1.2.2 Navigation

The application is organized into five main areas depicted by the top navigation buttons: Status, Create, Manage, Accelerate, and Preferences. Depending on your computer's configuration and available hardware, Create and Accelerate may not be available.



Status

The 'Status' area provides a general state of health of your storage system. If a status other than normal is reported, the Manage sub-section will be available to provide you with basic information and actions links necessary to return the status to normal.



Create

The 'Create' area allows you to create different types of volumes to protect data, enhance disk performance, optimize disk capacity, or create a custom volume to combine benefits.



Note

The ‘Create’ area is only available if your computer supports RAID technology, and if the volume requirements are met. Refer to the Volume Requirements topic for an exhaustive list of storage system conditions to create a volume.



Manage

The 'Manage' area combines the logical and physical view of your storage system. The area displays detailed information about each element that is part of the storage system, such as volumes and disks; the storage system view shows how the selected element relates to others. Each element has its own 'Manage' area which is accessible by clicking any element displayed in the storage system view under 'Status' or 'Manage'.
The 'Manage' area also provides the actions available for the selected element, such as renaming a volume or changing the volume type.



Accelerate

The ‘Accelerate’ area allows you to manage the cache memory configuration using a non-system solid state disk as a cache device. If the cache is reported in an abnormal state, detailed information

and troubleshooting actions will display. The Acceleration View is specific to the ‘Accelerate’ area and only displays in this location.



Preferences

The 'Preferences' area allows you to customize system settings by enabling the display of the notification area icon, and by selecting the type of notifications that you want the application to display.



Storage System View

The storage system view has two functions:

- It is a simplified representation of your storage system and displays graphic elements, such as arrays, volumes, devices, and ports. Each element provides general attribute information, such as status, name and size. Hovering over each element provides additional attribute details.
- You can also use the graphical view to access 'Manage' by clicking the storage system element you want to work with. For example, if an array is present, clicking the volume opens Manage Volume and clicking one of the array disks will open Manage Disk for the selected disk.



Acceleration View

The Acceleration View is a graphical representation of the acceleration configuration, and only displays the devices (disks and volumes) included in this particular configuration. You can use this view to access the ‘Manage’ page specific to each represented device by clicking the storage system element for which you want more detailed information.

23.1.3 Notification Area

The notification area (also called the system tray) is located on your desktop. The taskbar contains the notification area icon for Intel® Rapid Storage Technology. The icon provides storage system status and notifications such as volume and disk events based on a change of state.
The notification area icon will automatically display in the notification area once Intel Rapid Storage Technology is installed. Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Settings changes are applied on a per user basis, and do not affect other users’ settings.

Opening the application from the notification area

1. Right-click the icon.
2. Click 'Open Application'.

The notification area icon can be in the following states:

Icon	Description
	The storage system is reported in a **normal** state and your data is protected from a disk failure.
	The storage system is reported in a **warning** state and data may be at risk. We recommend that you open the application to review and resolve the reported issues.
	The storage system is reported in an **error** state and data may be lost. We recommend that you open the application to review and resolve the reported issues as soon as possible.
	The storage system is reported in a **busy** state while an operation is in progress. Once the operation is complete, all actions will be available again, allowing you to manage the storage system as long as it is reported in a normal state. You can follow the progress of the operation by hovering over the icon.
	This icon is displayed while you are attempting to open the application, but the Intel® Rapid Storage Technology service has not started running yet. The service is expected to start automatically with a delay when you launch Windows. This icon appears if you attempt to launch the application before the delay period ends. If the application fails to open, try starting the service manually using Microsoft Windows* Services.

Selecting system notifications

1. Right-click the icon.
2. Select the types of notifications you want to receive. The notification area menu allows you to select or deselect one option at a time. Repeat this procedure until you are finished with your selection. The same operation can also be completed using the application, from the 'Preferences' area.

Note: To hide the notification area icon, deselect 'Show the notification area icon' under 'System Preferences'.

Reviewing notifications

- Hover over the icon at any time to view the storage system status or the progression of an operation.
- Small pop-up windows will display for a short time to notify you of specific events, such as a missing disk or the completion of an operation.
- Open the application to view more details about storage system events in the 'Status' or 'Manage' areas.

23.2 Storage System Status

Anytime Intel® Rapid Storage Technology is launched, the application opens to the 'Status' area. This is where the general state of health of your storage system is reported, both in the storage system view and in details. Depending on the status, volume creation and management options may be available in order to enhance or repair your storage system.

23.2.1 Understanding the Status

To get the full benefits of what Intel® Rapid Storage Technology has to offer, it is critical to maintain a healthy storage system. The application helps you track and reports any disk or volume related problems that could put the safekeeping of your data at risk.

The storage system can be in the following states:



Normal

Reports that the system is functioning as expected, SATA disks are present and connected to the computer. If an array is present, volume data is fully accessible.



The Create subsection is only available if the storage system meets the minimum requirements to create a volume. Depending on the available hardware, you may be given the option to create a volume to protect data, optimize the disk performance, or create a custom volume.



The Manage subsection is only available if the storage system reports atypical conditions in a normal state. Typically, details or a recommended action are provided to help you rectify any storage system conditions. For example, if a recovery volume was reported as read-only, we would inform you that disk files must be hidden prior to requesting updates.



The Accelerate subsection is only available if a solid state disk can be used as a cache device and an eligible disk or volume can be accelerated. This area typically provides the option to enable acceleration and reports the cache and accelerated device health state, as well as the current acceleration mode.



Warning

Reports that storage system data may be at risk due to a problem detected on one or more SATA disks.

The Manage subsection displays any SATA disk or volume states reported by the storage system that may require your attention in order to keep data fully protected and accessible. Details or a recommended action are provided to help you fix any storage system problems. For example, if the master disk in a recovery volume is reported as failed, we would recommend that you rebuild the volume to another disk.

Note

In this state, we recommend that you backup any accessible data before taking action

In this state, the Accelerate subsection typically reports that the cache volume is failing possibly because the solid state disk is reported at risk of failing (smart event). Details and a recommended action are provided to help you fix the problem reported on the solid state disk.

Error

Reports that storage system data may be lost due to a problem detected on one or more SATA disks.

The Manage subsection displays any SATA disk or volume states reported by the storage system that require your immediate attention in order to keep data fully protected and accessible. Details or a recommended action are provided to help you fix any storage system problems. For example, if the data on a RAID 1 volume appears inaccessible due to a failed array disk, we would recommend that you rebuild the volume to another disk.

Note

In this state, we recommend that you backup any accessible data before taking action

In this state, the Accelerate subsection typically reports that the cache volume has failed possibly because the solid state disk has also failed and there is data loss. Details and a recommended action are provided to help you fix the problem reported on the solid state disk.

23.2.2 Storage System View

The storage system view provides a visual representation of your storage system and displays arrays, volumes, devices, and ports. Volumes and SATA disks graphics reflect their current states, which allows you to quickly identify the element that is causing the storage system to be in a state other than normal.

Note: Hovering over a designated element in the storage system view provides a snapshot of its properties. Clicking allows you to access and manage its properties.

Overview of SATA disks states

State	Description	Recommendation
	An internal hard disk is reported normal.	None
	An external hard disk is reported normal.	None
	An internal solid state disk is reported as normal.	None
	An external solid state disk is reported as normal.	None
	An internal solid-state hybrid disk reported as normal.	None
	An internal disk is reported missing.	Ensure that the disk is securely connected to the SATA port and that the SATA cable is functioning properly. Refer to the Troubleshooting section for more information.
	An internal disk is reported at risk or Incompatible.	Back up your data and replace the disks as soon as possible. Refer to the Troubleshooting section for more information.
	An external hard disk is reported at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.

	An internal solid state disk is reported as being at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.
	An external solid state disk is reported at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.
	An internal disk is reported offline.	Unlock all array disks to unlock the volume. Refer to the Troubleshooting section for more information.
	An internal recovery disk is reported offline.	• The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or, • Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.
	An external disk is reported offline.	Unlock all array disks to unlock the volume. Refer to the Troubleshooting section for more information.
	An external recovery disk is reported offline.	• The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or, • Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.
	An internal disk is reported normal and locked.	Unlock the disk to access more options.
	An external disk is reported normal and locked.	Unlock the disk to access more options.
	An internal hard disk is reported failed.	Refer to the Troubleshooting section for more information.
	An external hard disk is reported failed.	Refer to the Troubleshooting section for more information.
	An internal solid state disk is reported as failed.	Refer to the Troubleshooting section for more information.
	An external solid state disk is reported as failed.	Refer to the Troubleshooting section for more information.

Volume states

Volume type	Normal	Degraded	Failed
		Refer to Troubleshooting Degraded Volumes and Caching Issues for more information.	Refer to Troubleshooting Failed Volumes and Caching Issues for more information.
RAID 0		Not applicable	

Single-disk (cache)			
Single-disk (data)		Not applicable	
RAID 1/Recovery			
RAID 5			
RAID 10			

Other storage system elements

Element	Description	Recommendation
	A port that has no devices connected to it.	None
	An ATAPI device is present, such as CD-ROM, DVD/Blu-ray disc, or tape drive.	None

23.3 Creating a Volume

You can combine SATA disks to create a volume in order to enhance your storage system. Based on the available hardware and your computer's configuration, you may be able to create a volume by selecting an enhancement goal, such as 'Protect data' under 'Status', or by selecting a volume type under 'Create'. We recommend you get familiar with the minimum requirements in this section before starting the volume creation process.

Warning

Performing this action will permanently delete any existing data on the disks used to create a volume, unless you choose to keep the data when selecting array disks. Backup all valuable data before starting this process.

23.3.1 Volume Requirements

Creating a volume is only available as an option if the following requirements are met:

- You are logged on as an administrator.
- The computer is RAID ready (refer to the user's manual available on Intel's online support web site, for assistance on setting up a RAID ready system).
- Two or more SATA disks, including the operating system disk are connected, in a normal state, and unlocked (only applies to password-protected disks).
- Each of the SATA disks that are to be part of the RAID volume may not have any SMART events.

Enabling more disks

When configuring a volume, the application will only list the SATA disks that meet the requirements listed below. For example, a locked disk connected to your computer will not be listed as an option until it is unlocked.

Based on the first disk selected, some disks may become grayed out if one or more requirements are not met. Selecting a different disk generally helps re-enable disks that were previously grayed out.

- If the first selection is a system disk, any additional SATA disks selected must be of equal or greater size to ensure that all the system files are migrated to the new volume.
- If the first selection is a non-system disk, and a system disk is then selected, the latter must be of equal or smaller size to ensure that all the system files are migrated to the new volume.
- A system volume cannot be greater than 2 TB. If your first selection is a system disk, the total size of the other disks shall not allow the volume size to exceed 2 TB.
  Exception: If you are creating a volume using disks that have no existing data, and your operating system is a 64-bit Edition, the application will allow a volume to be greater than 2TB.

- The SATA disks used to create a volume must have the same type of connection, internal or external. An internal disk shall not be paired with an external disk to create a volume. Some systems will support mixed connection types.

Enabling more volume types
Depending on the input/output (I/O) controller hub that your computer is using and the hardware connected to the system, some volume types may not be enabled in the selection list. Refer to the Readme file located in the Program Files directory for this application or to the Device Manager to determine which controller is installed on your computer. Review the controller support table below to determine which volume types you can create.

Note: Intel® 5 Series Chipset applies to both desktop and mobile platforms as well as all later chipsets.

Volume Type	Number of Disks	Controller Support
Recovery volume	2	ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series. **Note** No other volumes can be present on the system. The master disk must include 100% of the available disk space and must be less than 1.3125 TB
RAID 0	2	ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 0	3 or 4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series
RAID 0	5 or 6	ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 1	2	ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 5	3 or 4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 5	5 or 6	ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 10	4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.

23.3.2 Creation Process

Now that you have reviewed the volume requirements, this section will guide you through the three easy steps necessary to create a volume.

23.3.2.1 Selecting a Volume Type

Before you can create a volume, you need to decide how you want to enhance your storage system based on your needs. Depending on the available hardware, you may have the option to combine volume types by creating more than one volume on a single array. Refer to 'Creating Additional

Volumes' for more information on this type of configuration. Below is an overview of the five volume types that you can create.

Creating a two-disk volume from 'Status'

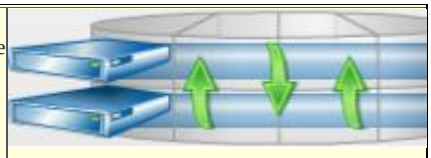
This option displays if only two disks are available, one has data such as system files, the second one doesn't, and the latter has a size that is equal or greater than the other. Based on this simple configuration, you can create a volume to protect data or optimize disk performance by clicking one of the two options listed in the Create subsection. When choosing this option, the application automatically configures the volume using the only two disks available and assigns default settings. Refer to the applicable procedure described in Completing the Process for details.

Creating a Custom Volume

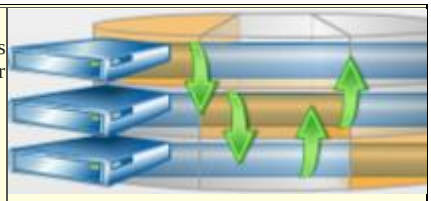
1. Click 'Create' or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click 'Next'.

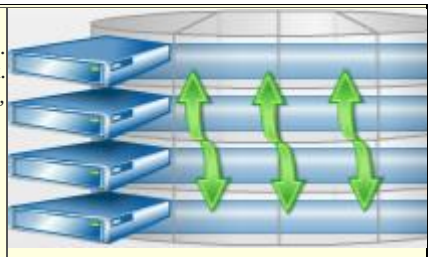
Recovery volume: **Flexible data protection** Combines two SATA disks and utilizes RAID 1 functionality to copy data from a designated master disk to a designated recovery disk. Data updates on the volume can be continuous or on request. In 'on request' mode, you can request data updates that copy changes from the master disk to the recovery disk since the last update. No other volumes can be present on the system. The master disk must include 100% of the available disk space and must be less than 1.3125 TB.	
Disks required	2
Advantage	Full data redundancy; more control over how data is copied between master and recovery disks; fast volume updates in on request mode; master and recovery disk files can be viewed in Windows Explorer*.
Disadvantage	Storage capacity is only as large as the smallest disk.
Application	Critical data protection for mobile systems; fast restoration of the master disk to a previous or default state. Available in specific mobile configurations.

RAID 1: **Real-time data protection** Combines two SATA disks where each stores an exact copy of the data to appear as a mirror of each other.	
Disks required	2
Advantage	Full data redundancy and excellent fault-tolerance; increased read transfer rate.
Disadvantage	Storage capacity is only as large as the smallest disk; slight decrease in write transfer rate.
Application	Typically used in workstations and servers to store critical data. Available in specific mobile configurations.

RAID 0: **Optimized disk performance** Combines two to six SATA disks and breaks down data into units that are spread across the array disks.	
Disks required	2 to 6
Advantage	Increased data access and storage performance; no loss in data capacity

Disadvantage	No data redundancy (if one disk fails, all data on the volume is lost).
Application	Typically used in desktops and workstations to store high performance, temporary data and software. Various RAID 0 volume configurations available in specific mobile configurations.

RAID 5: **Efficient data hosting and protection** Combines three to six SATA disks where data and parity are striped across the array disks in a rotating sequence. Parity is a mathematical method for recreating lost data to a single disk.	
Disks required	3 to 6
Advantage	Data redundancy; improved storage performance and capacity; high fault-tolerance and read performance.
Disadvantage	Time-consuming to rebuild and decreased performance during the process.
Application	Good choice for large amounts of critical data, such as file and application servers; Internet and Intranet servers. Available in mobile configurations that include the Intel® 5 Series Chipset which supports up to six SATA ports.

RAID 10 : **Balanced performance and data protection** Combines four SATA disks to create a combination of RAID types 1+0. The data is striped across a two-disk array forming a RAID 0 component. Each disk in the RAID 0 array is mirrored by a disk in the RAID 1 array, storing an exact copy of all the data.	
Disks required	4
Advantage	Combines the read performance of RAID 0 with the fault-tolerance of RAID 1, resulting in increased data access and full data redundancy, and increased storage capacity.
Disadvantage	4 disks are required, resulting in increased cost.
Application	High performance applications and high load database servers requiring data protection, such as video editing. Available in mobile configurations that include the Intel® 5 Series Chipset which supports up to six SATA ports.

23.3.2.2 Configuring the Volume

Once the volume type is selected, you are ready to configure your volume.

Recovery volume

1. Type a new volume name if you want to change the default name.
2. Select the master disk.
3. Select the recovery disk.
4. Select a different update mode, if desired.

5. Click ‘Next’. This button will not be active until all the required selections have been made.

Advanced configuration settings:
- Enable or disable the volume write-back cache.
- Select the check box to initialize the volume. You can choose to perform this action at a later time.

RAID Volume
1. Type a new volume name if you want to change the default name.
2. Select the required number of disks.
3. Select the disk from which you want to keep data, if desired. You can only keep data from one disk. If you want to keep data from more than one disk, you must back up all valuable data prior to creating a volume.
4. Click ‘Next’. This button will not be active until all the required selections have been made.
   Advanced configuration settings:
   - Select the array allocation by using the slider.
   - Select a data strip size.
   - Enable or disable the volume write-back cache.
   - Select the check box to initialize the volume. You can choose to perform this action at a later time.

Note: Currently, the application does not allow the creation of greater than 2TB volumes where the source disk is greater than 2TB and data on that disk is preserved (e.g. system volume). Target disks can be greater than 2TB but such volumes cannot. This limitation results from the lack of GPT partition scheme support. Note that volumes greater than 2TB that include member disks greater than 2TB are supported as long as array disks are unpartitioned or that no data is preserved at volume creation.

23.3.2.3 Completing the Process

If you are creating a custom volume, and have configured the volume with the disk selection and other settings, you are ready to review the projected configuration and complete the volume creation process.
If you are creating a two-disk volume for data protection or disk optimization from 'Status', you can follow the procedure provided below.

Creating a two-disk volume from ‘Status’
1. Under ‘Status’, in the Create sub-section, select the type of volume you want to create.
2. In the ‘Confirm Volume Creation’ dialog, review the volume configuration. Note that the volume name is the only setting that can be changed.
3. Click ‘Create Volume’ to confirm. The process starts immediately.
4. Once completed, a dialog displays to notify you that the volume was successfully created. Click ‘OK’ to close the dialog.
5. The page refreshes and displays the new volume in the storage system view as well as the data migration progress.

Creating a custom volume

⚠ **Warning**

You can only keep existing data from one of the disks you select to create a volume. We recommend that you backup all valuable data before proceeding.
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.
1. Under 'Confirm', review the selected configuration.
2. Click ‘Create Volume’ if you want to create the volume using the selected configuration. Otherwise, click ‘Back’ and make any necessary changes. The process starts immediately.
3. Click 'OK' to confirm.
4. Once completed, a dialog displays to notify you that the volume was successfully created. Click ‘OK’ to close the dialog.

5. The ‘Status’ area displays the new volume in the storage system view as well as the data migration progress.

If the size of the new volume is larger than the size of the source drive, the following steps apply:

6. Once the migration status reports 100% complete, restart your computer for the operating system to recognize the new volume size.
7. Create a new partition or extend the existing partition to utilize the new volume space using Windows Disk Management*. If your system is running Microsoft XP*, you may only have the option to create a new partition.

Note: To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management.

23.3.3 Creating Additional Volumes

Creating multiple volumes on a single array

You can add a volume to an existing RAID array by creating another volume that uses the available space on the array. This feature allows you to combine different volume types and their respective benefits. For example, a configuration with RAID 0 and RAID 1 on two SATA disks provides better data protection than a single RAID 0 and higher performance than a single RAID 1.

The first RAID volume occupies part of the array, leaving space for the other volume to be created. After creating the first volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to add a second volume to that array.

Note: The configuration is only available if the array allocation for the first volume created is less than 100%, and space is available on that array. The application currently supports an array to include a maximum of two RAID volumes.

1. Click ‘Create’ or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click ‘Next’.
4. Select 'Yes' to add the volume to an existing array.
5. Make any necessary changes in the Advanced section.
6. Click ‘Next’.
7. Review the selected configuration. Click 'Back' or an option in the left pane if you want to make changes.
8. Click 'Finish' to start the creation process.

Supported RAID volume combinations on a single array:

Combine	With
2-disk RAID 0	2-disk RAID 0
	2-disk RAID 1
2-disk RAID 1	2-disk RAID 0
	2-disk RAID 1
3-disk RAID 0	3-disk RAID 0
	3-disk RAID 5
3-disk RAID 5	3-disk RAID 0

Combine	With
	3-disk RAID 5
4-disk RAID 0	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10
4-disk RAID 5	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10
4-disk RAID 10	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10
5-disk RAID 0	5-disk RAID 0
	5-disk RAID 5
6-disk RAID 0	6-disk RAID 0
	6-disk RAID 5

Visit our Online Support for additional information on RAID type combinations for each I/O controller hub.

Creating additional volumes on a new array

You can choose to create two or more volumes on two different arrays, as long as the volume requirements are met.

1. Click ‘Create’ or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click ‘Next’.
4. Select 'No' in order to add a volume to a new array.
5. Select the required number of disks.
6. Select the disk from which you want to keep data, if desired. You can only keep data from one disk. If you want to keep data from more than one disk, you must back up all valuable data prior to creating a volume.
7. Make any necessary changes in the Advanced section.
8. Review the selected configuration. Click 'Back' or an option in the left pane if you want to make changes.
9. Click ‘Next’.
10. Click 'Finish' to start the creation process.

Note: Systems with an RST OROM older than 9.5, will not recognize 2 volumes on a single array if the RST Windows Driver version is 9.5 and newer.

23.4 Managing the Storage System

The 'Manage' area combines the logical and physical view of your storage system. The area displays detailed information about each element that is part of the storage system, such as volumes and disks; the storage system view shows how the selected element relates to others. Each element has its own 'Manage' area which is accessible by clicking any element displayed in the storage system view under 'Status' or 'Manage'.

The 'Manage' area also provides the actions available for the selected element, such as renaming a volume or changing the volume type.

23.4.1 Managing Arrays

You must be logged on as an administrator to perform the actions listed in this section.
You can manage arrays by clicking a selected array in the storage system view under 'Status' or 'Manage'. This allows you to review the properties and access all actions associated with that array, such as adding a disk or increasing a volume size.

23.4.1.1 Array Properties

An array is a logical grouping of physical SATA disks. The array properties listed below display to the left of the storage system view under Manage Array and report values specific to the element selected in the view.

Parameter	Value
Name	Reports the name of the array. The array name is automatically assigned and cannot be changed.
Size	Reports the total capacity of the array in megabytes (MB).
Available space	Reports the unallocated space on the array that can be used.
Disk data cache	Reports whether the data cache is enabled for all array disks.

23.4.1.2 Adding a Disk to an Array

You can add one or more SATA disks to an existing array to increase the system storage capacity. This feature can be useful if you want to change to a volume type that requires additional disks. This option is only available if:

- A RAID 0 and/or a RAID 5 volume is present,
- One or more SATA disks are connected to the computer and available,
- The available disk matches the internal or external connection type of the existing array disks.

You cannot add an external disk to an array that includes internal disks, and vice versa.
Refer to Connecting a Disk under Managing Disks for more information on installing SATA disks on your computer.

Warning
Any existing data on the available disk used to increase the array size will be permanently deleted. Backup all the data you want to preserve prior to executing this action.
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.
This action can also be performed from Manage Volume. Refer to the Adding a Disk to a Volume section for more information.

1. Under 'Status' or 'Manage', in the storage system view, click the array to which you want to add a disk. The element properties are now displayed on the left.
2. Click 'Add disk'.
3. Select the disk you want to use to increase the array capacity.
4. Click 'Add Disk'. Caution: Once the data migration starts, the operation cannot be canceled.

5. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes for which a disk was added, or add another partition.

Note: To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

23.4.1.3 Adding a Volume

You can add a volume to an existing RAID array by combining different volume types and their respective benefits. For example, a configuration with RAID 0 and RAID 1 on two SATA disks provides better data protection than a single RAID 0 and higher performance than a single RAID 1. The first RAID volume occupies part of the array, leaving space for the other volume to be created. After creating the first volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to add a second volume to that array.

Note: This configuration is only available if the array allocation for the first volume is less than 100%, and space is available on that array. The application currently supports an array to include a maximum of two RAID volumes on a single array.
You can also complete this action using the 'Create' area.

1. Under 'Status' or 'Manage', in the storage system view, click the array to which you want to add a volume. The array properties are now displayed on the left.
2. Click 'Create additional volume'.
3. In the 'Create Additional Volume' dialog, type a new name if you want to change the default name.
4. Select the volume type, and then click 'OK'. Only the volume types available for the current configuration will display. Refer to the table below for more information.
5. The page refreshes and the array now displays the additional volume.

Supported RAID volume combinations on a single array:

Combine	With
2-disk RAID 0	2-disk RAID 0
	2-disk RAID 1
2-disk RAID 1	2-disk RAID 0
	2-disk RAID 1
3-disk RAID 0	3-disk RAID 0
	3-disk RAID 5
3-disk RAID 5	3-disk RAID 0
	3-disk RAID 5
4-disk RAID 0	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10
4-disk RAID 5	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10
4-disk RAID 10	4-disk RAID 0
	4-disk RAID 5
	4-disk RAID 10

Combine	With
5-disk RAID 0	5-disk RAID 0
	5-disk RAID 5
6-disk RAID 0	6-disk RAID 0
	6-disk RAID 5

Visit our Online Support for additional information on RAID type combinations for each I/O controller hub.

23.4.1.4 Increasing Volume Size

You can increase the size of a RAID volume by using the remaining available space on the array. A minimum of 32 MB must be available for this action to be available. Hovering over the array name in the storage system view displays the amount of available space in MB.
After creating a volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to increase the volume size by the amount of available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size.
This option is only available if:
A RAID 0, RAID 1, RAID 5 and/or RAID 10 volume is present,

- The array allocation for the volume is less than 100% and space is available on the existing array.

Warning
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.
Increasing the volume size from Manage Array

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The array properties are now displayed on the left.
2. Click 'Increase size' next to the volume name. If more than one volume is present on a single array, you will need to increase the size of each volume one at a time.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Increasing the volume size from Manage Volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the left.
2. Click 'Increase size'.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Note: To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

23.4.1.5 Enabling Disk Data Cache

Enabling the disk data cache for all SATA disks on the array allows you to enable cache memory physically present on the disks and use it to speed up data access. This action is only available from Manage Array because the data cache must be in the same state across all disks that are part of a single array.

Under Manage Array, the disk data cache is reported as enabled or disabled for all SATA disks in the array. Under Manage Disk, the disk data cache is reported as enabled or disabled for a specific disk that is part of that array. The option to change this setting is only available from Manage Array.

Warning

Enabling the disk data cache increases the cache size and the amount of cached data that could be lost in the event of a power failure. The risk can be decreased if your computer is connected to an uninterruptable power supply (UPS).

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The element properties are now displayed on the left.
2. In the Advanced section, click 'Enable' or 'Disable' depending on the option available.
3. Click 'Yes' to confirm.
4. The page refreshes and now displays the new setting.

23.4.2 Managing Volumes

You must be logged on as an administrator to perform the actions listed in this section.

You can manage existing volumes by clicking a volume in the storage system view under 'Status' or 'Manage'. This allows you to review the volume properties and access all actions associated with that volume, such as renaming, changing type, and deleting.

23.4.2.1 Volume Properties

A volume is an area of storage on one or more SATA disks used within a RAID array. A volume is formatted by using a file system and has a drive letter assigned to it. The volume properties listed below display to the left of the storage system view under 'Manage' and report values specific to the element selected in the view.

RAID volume status table

Status	Description
Normal	Indicates that volume data is fully accessible.
Locked	Indicates that at least one array disk is locked with a password. The volume is visible because at least one other array disk is unlocked. Refer to Unlocking Password-Protected Disks for instructions on unlocking disks.
Degraded	Indicates that one array disk is missing or has failed. A RAID 0 volume cannot be in this state because of the striping configuration.
Failed	• RAID 0 volume: indicates that one or more array disks are missing or have failed. • RAID 1 volume: indicates that both array disks are missing or have failed. • RAID 5 or 10 volume: indicates that two or more array disks are missing or have failed.
Incompatible	Indicates that the volume was moved to another system that does not support the volume type and configuration.

Status	Description
Inaccessible	Indicates that data on the accelerated volume cannot be accessed because it is missing, or that the accelerated volume data is not synchronized with the data on the cache volume.
Unknown	Indicates that an unknown error was detected.

Recovery volume status table

Status	Description
Normal	Indicates that volume data is fully accessible.
Locked	Indicates that at least one array disk is locked with a password. The volume is visible because at least one other array disk is unlocked. Refer to Unlocking Password-Protected Disks for instructions on unlocking disks.
Degraded	• The recovery disk has failed, or • The master disk is missing or has failed and the volume is running off the recovery disk.
Failed	Indicates that both array disks have failed.
Incompatible	Indicates that the volume was moved to another system that does not support the volume type and configuration.
Unknown	Indicates that an unknown error was detected.
Power-saving mode	Indicates that the computer is running on battery power. If the volume is in continuous update mode, data updates are paused and will resume as soon as the computer is reconnected to the power supply.
Data update needed	Indicates that the recovery disk does not have a redundant copy of the data on the master disk, and you should request an update.
Running off recovery disk	Indicates that the recovery disk is the designated source drive in the volume.

Master disk read-only	Indicates that the recovery disk is the designated source drive in the volume, and that the master disk files are accessed. In this state, data recoveries from the recovery disk are not available.
Recovery disk read-only	Indicates that the recovery disk files are accessed. In this state, data updates are not available.

Busy volume states table

Status	Description
Initializing	Indicates that data on a volume is being synchronized. This step is required prior to verifying or verifying and repairing data on a volume.
Verifying	Indicates that the volume is being scanned to detect data inconsistencies.
Verifying and repairing	Indicates that the volume is being scanned to detect data inconsistencies, and errors are being repaired. This state does not apply to a RAID 0 volume because errors cannot be repaired.
Migrating data	Indicates that data is being reorganized on the volume. This state displays when a system volume is created, the volume size is increased, or the type is changed to different RAID configuration.
Rebuilding	Indicates that data redundancy is being restored across all disks associated with the volume. A RAID 0 volume cannot be in this state because of the striping configuration.
Recovering data	Indicates that data on the master disk is being overridden by all the data on the recovery disk. This state only applies to recovery volumes.
Updating data	Indicates that the latest master disk changes are being copied to the recovery disk. This state only applies to recovery volumes.

General parameters table

Parameter	Value
Details	Provides detailed information if a volume is in a state other than normal.
Type	Reports the volume type.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device.

Size	
	Enhanced: Indicates that the disk or volume is accelerated for optimized data protection.
	Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance.
	None: Indicates that no disk or volume is accelerated.
	Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.
	Reports the total capacity of the volume in gigabytes (GB) in the storage system view and in megabytes (MB) in the volume properties under Manage Volume.
Data stripe size	Reports the size of each logical contiguous data block used in the volume for RAID 0, 5, and 10 volumes. The strip size is indicated in kilobytes (KB).
Write-back cache	Reports whether the write-back cache feature is enabled for the volume.
System volume	Reports whether the volume contains system files that are required to start and run the operating system.
Initialized	Reports whether the volume is initialized.
Verification errors found	Reports the number of inconsistencies found during the last volume data verification.
Block with media errors	Reports the number of blocks with media errors found during the last volume data verification.
Physical sector size	Reports the size of each sector that is physically located on the disk.
Logical sector size	Reports the size of data collection blocks.
Details	Provides detailed information if a volume is in a state other than normal.

23.4.2.2 Renaming a Volume

You can change the name assigned to a volume present in your storage system at any time. The name change will take effect immediately.

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to rename. The volume properties are now displayed on the left.
2. Click 'Rename'.
3. Type a new volume name, and then click 'OK'.

Note: Volume names are limited to 16 English alphanumeric and special characters including spaces, but cannot include a backslash "\".

23.4.2.3 Rebuilding a Volume

When a volume is reported as degraded because of a failed or missing disk, the disk must be replaced or reconnected and the volume be rebuilt in order to maintain fault-tolerance. The option to rebuild is only available when a compatible disk is connected, available and normal. If a spare disk is available, the rebuild process will start automatically when a disk fails or is missing. For RAID 0 volumes, the rebuild process will start automatically only when one of its members is reported as at risk.

Warning

Completing this action will permanently delete existing data on the new disk and make any other volume on the array inaccessible. We recommend you backup valuable before continuing.

Rebuilding from 'Status' (manually)

1. Verify that the volume is reported as degraded in the Manage subsection. If you have more than one volume listed in this section, you will need to fix the issues reported one at a time.
2. Click 'Rebuild to another disk' next to the volume you want to rebuild.
3. In the Rebuild Volume dialog, select the disk that will replace the failed disk. Only compatible disks in a normal state will be displayed. Refer to Volume Requirements for more information.
4. Click 'OK' to confirm.
5. The volume starts rebuilding and the page refreshes displaying the progress of the operation. You can use other applications during this time and you will be notified when the process has successfully completed.

Rebuilding from 'Manage' (manually)

1. Under 'Status' or 'Manage', in the storage system view, click the volume you want to rebuild. The element properties are now displayed on the left.
2. Click 'Rebuild to another disk', and then follow the procedure described above.

23.4.2.4 Recovering Data

Recovering data to the master disk allows you to maintain full data redundancy on the recovery volume and keep the volume data healthy. This action is only available if a recovery volume is present and running off the recovery disk.

You may have to recover data if:

- Data on the recovery and master disk is not synchronized and full data redundancy is at risk.
- Data on the master disk is invalid or inaccessible.

Warning

Completing the action will override existing data on the master disk and update it with the data on the recovery disk. Backup all valuable data before continuing.

1. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'.
2. Click 'Yes' to confirm.
3. The recovery operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.

Note

If master disk is removed while the data recovery is in progress and is then reconnected, the operation will resume automatically from where it stopped as long as the volume is in on request update mode. If the volume is in continuous update mode, you will need to restart the operation by following the procedure described above,

23.4.2.5 Resetting Volume to Normal

This action is only available when a volume is reported as failed, but both array disks are present and normal, and allows you to access and try recovering healthy volume data.
In most cases, this situation will occur after one or more array disks was reported as failed or at risk, and then reset to normal.
Completing this action resets the volume state by ignoring previous events and does not repair data. Any data loss or corruption that may have occurred as a result of prior hardware failure or change of state remains. We recommend that you back up accessible data and replace failed hardware as soon as possible to prevent further data loss.

1. Under 'Status', in the Manage subsection, click 'Reset volume to normal'. You can also perform this action from Manage Volume, which is accessible by clicking the RAID 0 volume in the storage system view.
2. Click 'Yes' to confirm.
3. The page refreshes and the volume displays as normal. If the operation failed to return the volume to a healthy state, visit Intel's online support web site for more options.

23.4.2.6 Changing Volume Type

You can choose to change the type of an existing volume based on your storage system needs. The following configurations are possible:

Change type from	To
2-disk recovery volume **Note** Only available if the recovery volume is in continuous update mode	2-disk RAID 1
2-disk RAID 1	2-disk recovery volume **Note** No other volumes can be present on the system. The RAID 1 volume must be less than 1.3125 TB and include 100% of the available space on the array
2-disk RAID 1	2-disk RAID 0
	3, 4, 5 or 6-disk RAID 0
	3, 4, 5 or 6-disk RAID 5
2-disk RAID 0	3, 4, 5 or 6-disk RAID 5
3-disk RAID 0	4, 5 or 6-disk RAID 5
4-disk RAID 0	5 or 6-disk RAID 5
4-disk RAID 10	4, 5 or 6-disk RAID 5

Note
Before starting, refer to the system and volume requirements to determine which RAID types are supported by your computer and make sure the required number of SATA disks are connected. The Intel® Chipset provides support for the creation of all RAID volume types and for up to six SATA ports on a mobile platform. Changing volume type does not require re-installation of the operating system

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to modify. The volume properties are now displayed on the left.
2. Click 'Change type'.
3. In the 'Change Volume Type' dialog, type a new name if you want to change the default name.
4. Select the new volume type, and then click 'OK'. Caution: Once the data migration starts, the operation cannot be canceled.
5. Once the migration has completed, the 'Manage' page refreshes and reports the new volume type.

Warning
All applications and existing volume data remain intact, but any existing data on the disks added to enable this operation will be permanently deleted. Backup data before adding these disks.
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

23.4.2.7 Increasing Volume Size

You can increase the size of a RAID volume by using the remaining available space on the array. A minimum of 32 MB must be available for this action to be available. Hovering over the array name in the storage system view displays the amount of available space in MB.
After creating a volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to increase the volume size by the amount of available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size.
This option is only available if:

- A RAID 0, RAID 1, RAID 5 and/or RAID 10 volume is present,
- The array allocation for the volume is less than 100% and space is available on the existing array.

Warning
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

Increasing the volume size from Manage Array

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The array properties are now displayed on the left.
2. Click 'Increase size' next to the volume name. If more than one volume is present on a single array, you will need to increase the size of each volume one at a time.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Increasing the volume size from Manage Volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the left.

2. Click 'Increase size'.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

23.4.2.8 Adding a Disk to a Volume

You can add one or more SATA disks to an existing array to increase the system storage capacity. This feature can be useful if you want to change to a volume type that requires additional disks.

This option is only available if:

- A RAID 0 and/or a RAID 5 volume is present,
- One or more SATA disks are connected to the computer and available,
- The available disk matches the internal or external connection type of the existing array disks. You cannot add an external disk to an array that includes internal disks, and vice versa. In specific advanced system configurations, this condition may not apply.

Refer to Connecting a Disk under Managing Disks for more information on installing SATA disks on your computer.

Warning
Any existing data on the available disk used to increase the array size will be permanently deleted. Backup all the data you want to preserve before completing this action.
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state

This action can also be performed from Manage Array. Refer to the Adding a Disk to an Array section for more information.

1. Under 'Status' or 'Manage', in the storage system view, click the volume to which you want to add a disk. The element properties are now displayed on the left.
2. Click 'Add disk'.
3. Select the disk you want to use to increase the array capacity.
4. Click 'Add Disk'. Caution: Once the data migration starts, the operation cannot be canceled.
5. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes for which a disk was added, or add another partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

23.4.2.9 Changing Update Mode

A recovery volume gives you the flexibility to choose between updating data on the recovery disk continuously or on request.

In continuous update mode, the latest master disk changes are copied to the recovery disk automatically, as long as both disks are connected to the computer. In on request mode, the latest master disk changes are copied to the recovery disk only when you request a data update.

The current update mode is reported in the volume properties under Manage Volume. By default, the recovery volume is created in continuous update mode.

Note
This action is only available if a recovery volume is present and in normal state. If the recovery volume is read-only because the master or recovery disk files are accessed, you will need to hide the files before the update mode can be changed.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Change mode’, and then click 'Yes' to confirm.
3. The page refreshes and the volume properties report the new update mode.

23.4.2.10 Updating Data

You can manually copy the latest master disk changes to the recovery disk at any given time; this action allows you to synchronize data on the recovery volume, improving data protection and lowering the risk of losing valuable data in the event of a disk failure. When you request an update, only changes since the last update are copied.

Note
This action is only available if a recovery volume is present, and in ‘on request’ update mode.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Update data’.
3. The update process can be instantaneous or may take a while depending on the amount of data being copied. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.

Note
You can follow the progress of the update by hovering over the notification area icon or under ‘Status’ or Manage Volume.

23.4.2.11 Accessing Master or Recovery Disk Files

This action is only available if a recovery volume is present, in a normal state, and in on request update mode.

You can view the recovery or master disk files using Windows Explorer* depending on the designated source drive of the recovery volume. This feature can be useful when a data recovery from or to the master disk is necessary.

Accessing recovery disk files
This action is only available if the master disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click 'Access recovery disk files'.
3. Windows Explorer opens and displays the files located on the recovery disk.

Accessing master disk files
This action is only available if the recovery disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click 'Access master disk files'.
3. Windows Explorer opens and displays the files located on the master disk.

Note
When files have been accessed, the disk is displayed as missing from the array, and becomes available. Also, the volume is set to read-only and data updates are not available in this state. Hiding disk files will make the volume writable and allow data updates.

You can also access master or recovery disk files from Manage Disk.

23.4.2.12 Hiding Master or Recovery Disk Files

This action is only available if a recovery volume is present and disk files have been accessed.

When you are done viewing master or recovery disk files, you can hide the display of the files from Windows Explorer*. Once the disk files are hidden, the disk becomes writable, and data updates on the volume are available.

Hiding recovery disk files
This action is only available if the master disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Hide recovery disk files’.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Hiding master disk files
This action is only available if the recovery disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click 'Hide master disk files'.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Note
You can also hide master or recovery disk files from Manage Disk

23.4.2.13 Deleting a Volume

Use caution: you cannot recover data once a volume is deleted.

When a volume is deleted, you create available space that can be used to create new volumes. Note that you cannot delete a system volume using this application because the operating system needs the system files to run correctly. Also, if the volume is a recovery volume and the master or recovery disk files are accessed, you will need to hide these files before the volume can be deleted.

Warning
When a volume is deleted, all existing data on all disks that are a part of the selected volume is permanently lost. It is recommended to complete a backup of all valuable data before continuing.

1. Under 'Status' or 'Manage', in the storage system view, click the volume you want to delete. The volume properties are now displayed on the left.
2. Click 'Delete volume'.
3. Review the warning message, and click 'Yes' to delete the volume.
4. The 'Status' page refreshes and displays the resulting available space in the storage system view. You can now use it to create a new volume.

23.4.2.14 Setting the Data Strip Size

You can assign a data strip size to a volume while creating a new volume or while changing the type of an existing volume. You cannot change the strip size of an existing volume without changing its type.

The strip size refers to each logical contiguous data block used in a RAID 0, RAID 5, or RAID 10 volume. This setting is not available for RAID 1 or recovery volumes, due to their redundant configuration. The default value is the recommended strip size based on the system configuration and the volume type selected; changing the pre-selection is best suited for advanced users.

The following table describes the usage scenarios for the typical strip sizes.

Usage scenarios for supported strip sizes*

Strip Size	Description	RAID Types
4 KB	Best for Web Servers (fast read transfer rate with slow write transfer rate).	RAID 0, 10

8 KB	Best for databases (fast read transfer rate with faster write transfer rate than with 4KB strips).	RAID 0, 10
16 KB	Good for sequential transfers.	RAID 0, 5, 10
32 KB	Best for sequential transfers.	RAID 0, 5, 10
64 KB	Best general purpose strip size.	RAID 0, 5, 10
128 KB	Best for audio and video editing.	RAID 0, 5

*Disclaimer: The data provided in this table may vary based on the brand, type, size, and speed of the disks used.

Setting the strip size when creating a volume

1. Under 'Status', click 'Create' or 'Create a custom volume'.
2. Select the volume type, and then click 'Next'.
3. Make the required disk selection, and then select a new data strip size from the drop-down list in the Advanced section.
4. Complete the volume creation process as described in the Creation Process topic.

Setting the strip size when changing volume type

1. Under 'Status' or 'Manage', in the storage system view, click the RAID volume that you want to modify. The volume properties are now displayed on the left.
2. Click 'Change type'.
3. Make the necessary volume type and disk selections, and then select a new data strip size.
4. Click 'OK' to change the type of the existing volume.
5. The 'Manage' page refreshes and reports the new volume configuration.

Available Strip Size Configurations

	RAID 0	RAID 5	RAID 10
Default			
SATA disks	128 KB	64 KB	64 KB
Solid state disks	16 KB	128 KB	16 KB
Options	4 KB, 8 KB, 16 KB, 32 KB, 64 KB, 128 KB.	16 KB, 32 KB, 64 KB, 128 KB.	4 KB, 8 KB, 16 KB, 32 KB, 64 KB.

23.4.2.15 Enabling Volume Write-back Cache

You can improve the read/write performance of a RAID or recovery volume by enabling the write-back cache on one or all volumes on an array. When this feature is enabled, data may be temporarily stored in the cache memory before being written to the physical disks. Multiple I/O requests may be grouped together to improve performance. By default, the write-back cache is disabled.

Warning

While this feature highly improves the volume and array performance, it also increases the amount of cached data that could be lost in the event of a power failure. This risk can be lowered if your computer is connected to an uninterruptable power supply (UPS)

Enabling the volume write-back cache

1. Under 'Status' or 'Manage', in the storage system view, click the volume for which you want to enable the write-back cache. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Enable', and then click 'Yes' to confirm.
3. The page refreshes and the write-back cache is now enabled.

Note
If your computer is running on battery and a recovery volume is present, the option to enable the write-back cache is not available because the recovery disk is offline and data updates are not available. If this feature was enabled prior to running the battery, write-back cache activity would be temporarily disabled until you reconnect your computer to the power supply.

Disabling the volume write-back cache

1. Under 'Status' or 'Manage', in the storage system view, click the volume for which you want to disable the write-back cache. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Disable', and then click 'Yes' to confirm.
3. The page refreshes and the write-back cache is now disabled.

23.4.2.16 Initializing a Volume

Initializing a volume is the process of synchronizing all redundant data on a volume prior to verifying or verifying and repairing that data. If you attempt to start a verification process for a volume that has not been initialized, you will be prompted to do so.

Initializing a volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to initialize. The volume properties are now displayed on the left.
2. Click 'Initialize'.
3. Click 'OK' to start the initialization process. Caution: Once the data migration starts, the operation cannot be canceled.

Initializing a volume when verifying data

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the left.
2. Click 'Verify'.
3. When prompted to initialize the volume before verifying data, click 'OK' to start the initialization process. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once complete, click 'Verify' to start the verification process.

Note
While inititialization is in progress, you can view the status in the notifications area by hovering over the Intel(R) Rapid Storage Technology icon, or in the application under Staus or Manage Volume.

Warning

The initializating process could take a while depending on the number and size of the disks. You can continue using array disks and other applications during this time. Closing the application, or powering off and restarting your computer will not disrupt the progress of this operation.

23.4.2.17 Verifying and Repairing Data

You can verify data on an existing volume by identifying and repairing inconsistencies. Running this operation on a regular basis helps you keep valuable data and the overall storage system healthy.

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the left.
2. Click 'Verify'.
3. Select the check box if you want errors found to be repaired automatically during the verification process.
4. Click 'OK' to start the verification process.

Note
Data on a volume cannot be verified and repaired unless the volume has been initialized first. If you attempt to start a verification process for a volume that is not initialized, you will be prompted to first initialize the volume. Based on its configuration, a RAID 0 volume cannot be repaired because of the lack of redundancy.

23.4.2.18 Swapping Disks

You can change the order of designation for array disks in a recovery volume by setting the master disk as the destination drive and the recovery disk as the source drive. This action is best suited for advanced users.

Note
This action is only available if a recovery volume is present, normal, and in continuous update mode.

Swapping disks can be useful if:

- You selected the wrong disk as the master disk when you created the recovery volume,
- You think one of the disks is failing,
- You replaced the recovery disk with a faster, newer disk, and want to run off that device once it has been updated.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Swap master and recovery disks'.
3. Click 'Yes' to confirm.
4. Hover over each disk in the storage system view to review their new usage.

23.4.3 Managing Disks

You must be logged on as an administrator to perform the actions listed in this section.
You can manage disks by clicking a selected disk in the storage system view under 'Status' or 'Manage'. This allows you to review the properties and access all actions associated with that disk, such as unlocking a password-protected disk or marking a disk as spare. Depending on their usage or status, some actions may not be available.

23.4.3.1 Disk Properties

The disk properties listed below display to the left of the storage system view under 'Manage' and report values such as usage and status that are specific to the disk selected in the view. Based on the detailed information provided, you can make changes to the way each disk is configured, or take action on one or more disk to keep your overall storage system healthy.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.
Usage	Array disk: a disk that has been grouped with other disks to form an array containing RAID volumes.
	Master disk: the disk that is the designated source drive in a recovery volume.
	Recovery disk: the disk that is the designated destination drive in a recovery volume.
	Spare: the disk has been designated as the destination drive for automatic volume rebuilds in the event of a failed, missing or at risk array disk. For RAID 0 volumes, automatic rebuilds will only occur when one of its array disks is reported as at risk.
	Available: the disk is physically connected to the computer, healthy, and available to be used in an array or as a spare disk. **Warning** Assigning an available disk to an array or marking it as a spare will permanently delete any existing data on that disk.
	Unknown: the disk is available but contains metadata that cannot be displayed in the operating system. Even though the disk is reported as normal, you will need to clear and reset the disk to make the disk available.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device.
	Enhanced: Indicates that the disk or volume is accelerated for optimized data protection.
	Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance.
	None: Indicates that no disk or volume is accelerated.
	Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.

Status	Normal: the disk is present, functioning as expected, and unlocked. Locked: the disk is password-protected. Note: if a volume includes at least one locked disk, the volume will display as locked. At risk: an impending error condition was detected on the disk and it is now at risk of failure. Missing: the disk is not present or physically connected to the computer. Failed: the disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold. Offline: indicates that an array disk is locked, that the recovery volume is in on request update mode, or that your computer is running on battery and data updates to the recovery volume are not available.
Size	Reports the total capacity of the disk in megabytes (MB) in the disk properties and in gigabytes (GB) in the storage system view.
Serial number	Reports the manufacturer's serial number for the disk.
Model	Reports the model number of the disk.
Firmware	Reports the version of the firmware found in the disk.
System disk	Reports whether the disk contains system files that are required to start and run the operating system.
Password protected	Reports whether the disk is protected with a password.
Disk data cache	Reports whether the data cache is enabled on this disk. This feature is controlled at the array level.
Native command queuing	Reports whether the disk supports this feature.
SATA transfer rate	Reports the data transfer rate between the SATA controller and the SATA disk. The supported rates are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.
Physical sector size	Reports the size of physical sectors on the disk (bytes).
Logical sector size	Reports the size of logical sectors on the disk (bytes).

23.4.3.2 Unlocking Password-Protected Disks**

You can unlock a password-protected disk by entering the password which allows you to access data or use that disk to create a volume. The password is setup through the system BIOS. Locked disks

can be identified with the lock icon appended to them and display a 'Locked' status in the disk properties.

23.4.3.3 Marking a Disk as Spare

This action is only available for non-system disks in a normal state. Also, unless your mobile computer is equipped with the Intel® 5 Series Chipset or later, which provides support for up to six SATA ports, you will not be able to mark a disk as a spare. Most mobile computers are limited to one internal and one external disk, which are used to create the volume.

Marking a disk as a spare allows you to designate an available SATA disk as the default destination for automatic volume rebuilds in the event of a failed, missing or at risk array disk. However, for RAID 0 volumes, automatic rebuilds will only occur if one of its members is reported at risk.

1. Under 'Status' or 'Manage', in the storage system view, click the disk that you want to mark as a spare. The volume properties are now displayed on the left.
2. Click 'Mark as spare'.
3. Click 'OK'.

Note
RAID 1, 5, 10, and recovery volumes can use one or more spares.

Warning
When marking a disk as a spare, any existing data on that disk is permanently deleted. Back up all data you want to preserve before starting this action.

If you system is running a version of the RST OROM that does not support disks that are 2TB or larger, you can reset such a disk to available, but disallow the marking of it as a spare.

23.4.3.4 Resetting a Disk to Available

After a disk was marked as spare, you can choose to make that spare disk available again and use it differently. Once available, the disk can be used to create a volume or be added to an existing volume if all other requirements are met.

1. Under 'Status' or 'Manage', in the storage system view, click the disk that you want to reset to available. The volume properties are now displayed on the left.
2. Click 'Reset to available'.
3. The page refreshes and the disk usage is now reported as available.

23.4.3.5 Resetting a Disk to Normal

You can reset a SATA disk to normal when the storage system reports one of the following disk statuses:

At risk
A disk is reported at increased risk of failing in the near future that could be due to a slow degradation over time. You can choose to ignore this alert at this time by resetting the disk to normal, but it may re-appear if the disk continues to assert this condition. We recommend that you contact the manufacturer for more information to prevent potential data loss.

Failed

A SATA disk has failed to properly complete read and write operations in a timely manner, and data may be lost. We recommend that you replace the failed disk as soon as possible to return the overall storage system to normal. In this state, data may be lost, but you can try resetting the disk to normal and attempt a data recovery. If the disk operations continue to fail, the disk will return to a failed state immediately.

If the failed disk is an array disk, refer to the Troubleshooting section for guidelines on rebuilding a failed or degraded volume.

1. Under 'Status', in the Manage subsection, locate the disk reported as at risk or failed. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view.
2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.

Note
Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss.

23.4.3.6 Accessing Disk Files

This action is only available if a recovery volume is present, in a normal state, and is on request update mode.

This feature allows you to view the files on the designated destination drive in a recovery volume using Windows Explorer*. For example, you may want to review the recovery disk files prior to starting a data recovery in the event that data on the master disk is inaccessible or corrupted.

When the volume status is normal, the recovery disk is the designated destination drive and files are accessible. When the volume status is running off the recovery disk, the master disk is the designated destination drive and files are accessible. You can review the usage of each disk by hovering over the array disks in the storage system view or by clicking one of the disks to review its properties under Manage Disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery or the master disk depending on the volume status. The disk properties are now displayed on the left.
2. Click 'Access files'.
3. Windows Explorer opens and displays the files located on the disk.

Note
When files have been accessed, the disk is displayed as missing from the array, and becomes available. Also, the volume is set to read-only and data updates are not available in this state. Hiding disk files will make the volume writable and allow data updates.

Warning
Windows Explorer will not open if the disk does not have any partitions on it.

23.4.3.7 Hiding Disk Files

This action is only available if a recovery volume is present and disk files have been accessed.

When you are done viewing master or recovery disk files, you can hide the display of the files from Windows Explorer*. Once the disk files are hidden, the disk becomes writable, and data updates on the volume are available.

1. Under 'Status' or 'Manage', in the storage system view, click the disk whose files are accessed. The disk properties are now displayed on the left.
2. Click 'Hide files'.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Note
You can also hide master or recovery disk files from Manage Volume.

23.4.3.8 Connecting a Disk

Installing new hardware is one of the steps you may have to take to keep you storage system healthy or to extend the life of a computer that is running out of storage space.

Intel® Rapid Storage Technology provides hot plug support, which is a feature that allows SATA disks to be removed or inserted while the computer is turned on and the operating system is running. As an example, hot plugging may be used to replace a failed external disk.

Our application provides support for SATA 1.5 Gb/s (generation 1), SATA 3 Gb/s (generation 2), and 6 Gb/s (generation 3) data transfer rates. The rate support depends on the Intel® Chipset and SATA disks present in your system. Visit our Online Support for additional information on chipset features and benefits.

Follow these procedures to replace or connect a disk in case you need to power off your computer:

Replacing a disk

1. Power off your computer.
2. Replace the disk that reports a problem.
3. Turn your computer back on. If the replaced disk was part of an array, you will need to follow the procedure provided in the Troubleshooting section based on the volume state and type.

Note
To install an external disk, plug it into you computer and connect the power cord.

To remove and install an internal disk, you should be comfortable opening your computer case and connecting cables. Follow the manufacturer's installation guide to complete this procedure. If you are replacing the system disk, you will have to re-install the operating system after you connect the disk because the system disk contains the files required to start and run your computer.

Installing a new disk (to increase storage space)

1. Power off your computer.
2. Connect the new disk.

3. Turn your computer back on. During the system startup, the application's option ROM should automatically detect the new disk if it is installed correctly. Once you open the application, verify under 'Status', in the storage system view, that the new disk displays. You can then access management options by clicking that disk.

23.4.4 Managing Ports

A port is a connection point on your computer where you can physically connect a device, such as a SATA disk or ATAPI device. A port transfers I/O data between the device and the computer.

If a port is reported as empty in the storage system view, you can use that port to connect a new device in order to increase the storage system capacity. Currently, the maximum number of internal ports that can be used to connect devices is six.

The port properties listed below display to the left of the storage system view under 'Manage' and report values specific to the element selected in the view.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.

23.4.5 Managing ATAPI Devices

An ATAPI device is a mass storage device with a parallel interface such as a CD-ROM, DVD/Blu-ray disc, tape drive, or solid-state disk. The ATAPI properties listed below display to the left of the storage system view under 'Manage' and report values specific to the selected element.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.
Serial number	Reports the manufacturer's serial number for the device.
Model	Reports the model number of the device.
Firmware	Reports the version of the firmware found in the device.
SATA transfer rate	Reports the transfer mode between the SATA controller and the ATAPI device. The typical values for this parameter are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.

23.4.6 Managing Solid-State Hybrid Drives (SSHD)

To increase performance, some hard drive manufactures are now including non-volatile memory to be used as a data cache on hard drives. These drives are called Solid-State Hybrid Drives. Hybrid Hinting is a feature by which the RST driver will send information to a SSHD, notifying the drive which data would be best to store in the data cache. There are no extra steps to manage SSHD's.

RST enables Hybrid Hinting when the following minimum requirements are met:

- SSHD must have a minimum of 8GB of non-volatile cache.
- SSHD must have a maximum of 1TB of non-volatile cache.
- SSHD is attached to a supported SKU for Intel® Smart Response Technology for Hybrid Drive Acceleration. (See section 21.1.2 for SKU support)

23.5 Accelerating the Storage System

You can configure an internal solid state disk to be used as a non-volatile intelligent caching for a system or non-system disk or volume that's present on your system. Moving frequently accessed data over to the cache allows you to improve overall system performance, increase read/write access times, and reduce start up times without adding more system memory.

This feature also increases the power efficiency of a mobile computer by retaining stored data and reading data from the cache instead of the SATA disk itself.

Accelerate is only available if the requirements listed in this section under Cache Device Properties are met.

23.5.1 Cache Device Properties

The Performance tab-> enable acceleration page in the UI are only available if the following requirements are met:

- See Section 20 for requirements
- Operating system: all supported OS for this release (see section 2.1.5).
- CPU: Core or Premium
- BIOS: RAID-Ready system and Accelerate feature bit is set.
- An internal SATA solid state disk is present with a minimum capacity of 16 GB.
- A hard disk or volume (array members must all be hard disks) is eligible for acceleration.
- No recovery volume is present.

Limitations

- The maximum cache size is 64 GB.
- Only one disk or volume at a time can be accelerated per system.
- If two volumes are present on a single array (they share the same array of disks), neither volume can be accelerated.
- Once a volume is accelerated, a second volume cannot be added to the same array.
- Once a solid state disk is configured to be used as a cache device, the option to create a recovery volume is no longer available. Recovery volumes do not support system configurations with multiple volumes.

Solid State Disk Properties

Parameter	Value
Port	Reports the port number to which the solid state disk is attached.
Port location	Reports that the solid state disk is internal.
Status	Reports the state of health of the internal solid state disk present in the system.
	Normal: Indicates that the solid state disk is present, functioning as expected, and unlocked.
	Failed: Indicates that the solid state disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.
	At risk:
	Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that cache data is being deleted in order to disable acceleration. In some cases, these transitions will start automatically in the event that errors are detected and a risk of data loss is identified.
Type	Reports that the device is a solid state disk.
Usage	Reports that the solid state disk is configured to be used as a cache device.
Size	Reports that the solid state disk is configured to be used as a cache device.
Serial number	Reports the manufacturer's serial number for the internal solid state disk.
Model	Reports the model number of the solid state disk.
Firmware	Reports the version of the firmware found in the solid state disk.
Password protected	Reports whether the solid state disk is password-protected.
Disk data cache	Reports that the data cache is enabled on the solid state disk. When a solid state disk is configured as a cache device, this setting can only be changed at the operating system level.
Native command queuing	Reports whether the solid state disk supports this feature.
SATA transfer rate	Reports the data transfer rate between the SATA controller and the SATA solid state disk. The supported rates are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.
Physical sector size	Reports the size of physical sectors on the solid state disk (bytes).
Logical sector size	Reports the size of logical sectors on the solid state disk (bytes).
Accelerated device	Indicates the location of the disk or the name of the volume that is currently accelerated by the cache device.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device .

	Enhanced: Indicates that the disk or volume is accelerated for optimized data protection. Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance. None: Indicates that no disk or volume is accelerated.

Cache and Simple Data Volume Properties

Parameter	Value
Name	Reports the name of the volume.
Status	Normal: For simple data volumes, indicates that volume data is fully accessible. For cache volumes, indicates that cache data is fully accessible, and that caching activity with the associated disk or volume is occurring under healthy conditions. Failing: Indicates that a SMART event was detected on the solid state disk that is used as a cache device. Failed: Indicates that the cache volume has exceeded its recoverable error threshold, and that read and write operations are no longer occurring.
Type	Indicates that the single-disk RAID 0 volume is a cache volume.
Data stripe size	Indicates that the single-disk RAID 0 volume is a cache volume.
Allocated cache size	Reports the volume capacity used for cache memory.
Write-back cache	Reports whether the write-back cache feature is enabled for the volume.
Physical sector size	Reports the size of each sector that is physically located on the disk.
Logical sector size	Reports the size of data collection blocks.

Accelerated Disk or Volume Properties

Parameter	Value
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device. Enhanced: Indicates that the disk or volume is accelerated for optimized data protection. Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance. Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.

23.5.2 Enabling Acceleration

You can enable acceleration in order to improve the performance for a SATA hard disk or a RAID volume that includes only SATA hard disks. This operation caches its contents using a non-volatile memory device (a solid state disk) that is attached to an AHCI port.

Enabling acceleration allows you to:

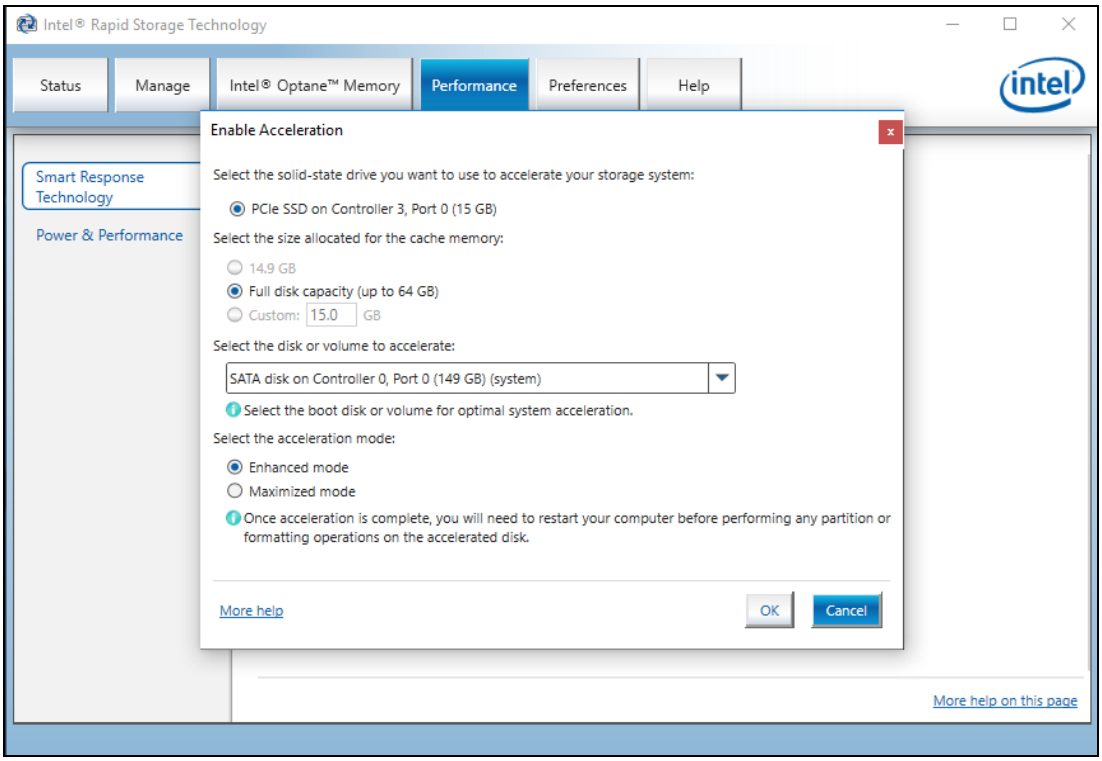
- Use a solid state disk as a cache device: The maximum cache size is 64 GB.
- Accelerate one system or non-system disk or volume present in the system by associating it with the cache volume, and subsequently caching its content.
- Configure acceleration in enhanced (optimized for data protection) or maximized mode (optimized for input/output performance). See below for more information on the acceleration modes.

- If the solid state disk used as a cache device is larger than 64 GB and has a minimum of 4 GB of additional space, a second single-disk RAID 0 volume will be automatically created which can be used for simple data storage.

Enabling Acceleration

Follow these steps to enable acceleration:

1. From the Performance Page, Click 'Intel Smart Response Technology' from the left side panel. Click Enable acceleration.



2. Select the solid state disk you want to use as a cache device.
3. Select the portion of the solid state disk you want to use to store non-volatile cache memory. Any remaining space on the solid state disk may be used for data storage using the simple data single-disk RAID 0 volume that is automatically created.
4. Select the disk or volume you want to accelerate. We highly recommend that you accelerate the system volume or system disk for maximum performance.
5. Select the acceleration mode you want to use, and then click 'OK'. By default, enhanced mode is selected.
6. The page refreshes and reports the new acceleration configuration in the Acceleration View.

Acceleration modes

Non-volatile cache memory can be enabled in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**

This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the data in the cache memory. For data safety reasons, this mode is the default acceleration setting.

- **Maximized mode: Acceleration optimized for input/output performance.**
  This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn't synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

23.5.3 Disabling Acceleration

You can disable acceleration on a disk or volume if you want to:

- Enable acceleration on a different disk or volume,
- Return the solid state disk to pass-through,
- Physically move an accelerated disk or volume to another computer.

Completing this action makes any cached data associated with the accelerated disk or volume immediately inaccessible. If the current acceleration mode is maximized, disabling acceleration may take a while to complete, depending on the cache and the solid state disk size. You can use other applications during this time.

1. Click 'Accelerate', and then click 'Disable acceleration'.
2. In the dialog, click 'Yes' to confirm.
3. The page refreshes and reports the acceleration as disabled.

In the event that you are unable to open or access Intel® Rapid Storage Technology due to an application error or operating system issue, you will need to disable acceleration using the option ROM user interface.

1. Restart your computer.
2. Press Ctrl-I to access the main menu of the option ROM user interface.
3. Select 'Acceleration Options' from the main menu.
4. Select the accelerated disk or volume.
5. If acceleration is in maximized mode, type 's' to synchronize data from the flash memory to the accelerated disk or volume. Otherwise, go to step 7.
6. Press 'Y' to confirm.
7. Type 'r' to remove acceleration.
8. Press 'Y' to confirm.

23.5.4 Changing Acceleration Mode

This action is only available if a disk or volume is currently accelerated. A disk or volume can be accelerated in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**
  This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes

inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the data in the cache memory. For data safety reasons, this mode is the default acceleration setting.

- **Maximized mode: Acceleration optimized for input/output performance.**
  This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn't synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

By default, acceleration is enabled in enhanced mode due to the lower risk of data loss, but you can change acceleration mode at any time as long as the cache volume and accelerated device are in a normal state and caching activity is occurring.

Follow these steps to change the acceleration mode:

1. Click 'Accelerate', and then click 'Change mode'.
2. Click 'Yes' to confirm the mode change to either enhanced or maximized, depending on the current acceleration mode.
3. The page refreshes and the new acceleration mode displays under the Acceleration Configuration subsection and the Acceleration View.

Warning
When a device is accelerated in Maximized mode, performance is highly improved but cached data is at higher risk of being lost in the event of a power failure or under other atypical conditions.

Acceleration in a busy state

The acceleration mode will display as busy under the following conditions (by user interaction or automatic transition):

- When changing acceleration mode from maximized to enhanced.
- When disabling acceleration while in maximized mode.

The transition time varies based on the cache and disk sizes. Disk and volume actions will not be available until the acceleration transition has completed, except for renaming and deleting volumes.

23.5.5 Accelerating a Disk or Volume

Once a solid state disk is configured to be used as a cache device, you can choose to accelerate any disk or volume in a normal state that is located on your storage system. We recommend that you accelerate the system disk or volume in order to get the full benefits of the non-volatile cache memory configuration.

Follow these steps to accelerate a disk or volume:

1. Click 'Accelerate', and then click 'Select device'.
2. In the 'Accelerate Disk or Volume' dialog, select the device you want to accelerate.

3. Select the acceleration mode you want to use, and then click ‘OK’. By default, enhanced mode is selected.
4. The page refreshes and reports the newly accelerated device in the Acceleration Configuration subsection. The Acceleration View also indicates the accelerated device with the acceleration icon appended to it.

Non-volatile cache memory can be enabled in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**
  This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the data in the cache memory. For data safety reasons, this mode is the default acceleration setting.
- **Maximized mode: Acceleration optimized for input/output performance.**
  This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn’t synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

23.5.6 Resetting a Cache Device to Available

This action is only available if a solid state disk is configured as a cache device and there is no accelerated disk or volume present (no association with the cache device). In this situation, you have two options:

- Reset the solid state disk to available and use that device for other purposes.
- Accelerate a disk or volume that is eligible and available for acceleration. Refer to Cache Device Properties for a detailed list of eligibility requirements.

Warning
In the event that a single-disk RAID 0 data volume was created along with a cache volume, resetting the solid state disk to available will delete both volumes. Data on the RAID 0 data volume will be permanently erased. Backup all valuable data before beginning this action.

1. Click ‘Accelerate’.
2. Click 'Reset to available'.
3. In the dialog, select the check box to confirm that you understand that data on the data volume will be permanently deleted.
4. Click 'Yes' to confirm.
5. The 'Accelerate' page refreshes. Under 'Status', the storage system view displays the solid state disk usage as available. The device can now be used for any purpose.

23.5.7 Disassociating the Cache Memory

This action is only available if an issue is reported on the accelerated disk or volume that is associated with the cache device and it is missing. In this state, the acceleration mode is typically reported as unavailable and caching activity is no longer occurring.

If you are unable to resolve the reported issue on the accelerated disk or volume (e.g., degraded or failed volume due to a missing array disk), the only option will be to remove the association between the cache device and the disk or volume.

Once the association between the cache and the accelerated disk or volume is removed, all cache metadata and data is deleted from the cache device. You can then reset the solid state disk to available or accelerate a different disk or volume, as long as the cache device is healthy.

Follow these steps to disassociate the cache memory and the accelerated device:

1. Click 'Accelerate'.
2. Click 'Disassociate'.
3. In the 'Disassociate' dialog, click 'Yes' to confirm.
4. The page refreshes and the Acceleration View reports the new configuration. Options to reset the solid state disk to available or to select a new device to accelerate (as long as an eligible disk or volume is available) are now available.

Note
You can also perform this action using the option ROM user interface.

23.6 Preferences

System preferences allow you to decide whether you want the notification area icon to display. In addition, you can select the types of notifications you want to receive, such as storage system warnings or errors, and be notified of any reported problems while the application is closed.

Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Settings changes are applied on a per user basis, and do not affect other users' settings.

Showing system notifications

By default, System preferences are set to show the notification area icon. If you previously chose to hide the notification area icon, follow these steps to display the icon again:

1. Under 'Preferences', select 'Show the notification area icon'.
2. Click 'Apply Changes'. Verify that the icon is now displayed in the notification area.

Hiding system notifications

Once you hide the notification area icon, the service no longer reports storage system information, warnings, or errors through the notification area. You will need to use the application to monitor the health of the storage system. Follow these steps to hide the notification area icon:

1. Under 'Preferences', deselect 'Show the notification area icon'.
2. In the 'Hide Notification Area Icon' dialog, click 'Yes' to confirm.
3. Verify that the icon is no longer displayed in the notification area.

Selecting system notifications

1. Under ‘Preferences’, make sure that ‘Show the notification area icon’ is selected.
2. Select the types of notifications you want to receive.

Note

Storage system information provides details on any changes of state other than warnings or errors, such as new disks being detected or locked.

Storage system warnings report the cause for the overall warning state of the storage system, such as a degraded RAID volume due to a missing disk.

Storage subsystem errors report the cause for the overall error state of the storage system, such as a failed volume due to a failed disk.

24 Opal Drive Support

Support for Opal drives was added beginning with Intel® Rapid Storage Technology version 13.0.

24.1 Opal ATA Commands Supported

The Intel® RST driver will support the following Opal ATA commands on systems in either AHCI or RAID mode (only on pass-thru drives):

- ATA_PASS_THRU_DIRECT IOCTL commands of ATA TRUSTED RECEIVE and ATA TRUSTED SEND.

24.2 Intel® Rapid Storage Technology UI Support

The Intel® RST UI will display information about the Opal disks attached to the SATA controller when the Opal disk is partially or fully Opal locked.

24.3 Opal SCSI Commands Supported

The Intel® RST driver supports translation of the following SCSI Commands on systems in AHCI or RAID mode (only pass-thru disks):

- SCSI_PASS_THROUGH_DIRECT IOCTL commands of SECURITY_PROTOCOL_IN and SECURITY_PROTOCOL_OUT to ATA TRUSTED RECEIVE and ATA TRUSTED SEND.

24.4 Opal UEFI Protocol Supported

The Intel® RST UEFI driver supports the following UEFI Protocol for OPAL Support for systems in RAID mode with the RST UEFI Driver installed:

- EFI_STORAGE_SECURITY_PROTOCOL

24.5 Intel® RST PCIe NVMe Support for Opal

Intel® RST PCIe NVMe Support for Opal requires security communication protocols. The following NVMe Security commands are required for Opal support:

- SECURITY SEND
- SECURITY RECEIVE

RST will allow the Security Send and Security Receive commands to be sent only when using the following protocols:

- RST UEFI

- EFI_STORAGE_SECURITY_COMMAND_PROTOCOL: EFI_STORAGE_SECURITY_COMMAND_PROTOCOL is loaded ***ONLY*** if a device supports TCG and Block IO protocol is loaded successfully

- RST OS Driver
  - SCSI SECURITY PROTOCOL: Intel® RST driver supports translating commands to NVMe Security Send and NVMe Security Receive commands on TCG supported Opal PCIe NVMe SSDs.

24.6 NVMe Device Requirements for Intel® RST Opal Support

Intel® RST will recognize Opal supported NVMe Devices that report support for the Security Send and Security Receive commands.

Intel® RST uses the register from Admin Command Set Attributes and Optional Controller Capabilities (bytes 257:256). RST requires that bit 0 of this register must be set to a value of 0x01 to indicate Opal support. If the bit is not set, RST will not recognize the device as supporting Opal.

NVM Express 1.1b

79:78	M	**Controller ID (CNTLID):** Contains the NVM subsystem unique controller identifier associated with the controller. Refer to section 7.9 for unique identifier requirements.
255:80		Reserved
Admin Command Set Attributes & Optional Controller Capabilities		
257:256	M	**Optional Admin Command Support (OACS):** This field indicates the optional Admin commands supported by the controller. Refer to section 5. Bits 15:3 are reserved. Bit 2 if set to '1' then the controller supports the Firmware Activate and Firmware Download commands. If cleared to '0' then the controller does not support the Firmware Activate and Firmware Download commands. Bit 1 if set to '1' then the controller supports the Format NVM command. If cleared to '0' then the controller does not support the Format NVM command. Bit 0 if set to '1' then the controller supports the Security Send and Security Receive commands. If cleared to '0' then the controller does not support the Security Send and Security Receive commands.

§ §

25 eDrive Support

Support for eDrives is new for Intel® Rapid Storage Technology beginning in version 13.2. eDrive provides transport of Trusted Computing Group (TCG) protocol over 1667.

Intel® RST added support for TCG communication in version 13.0 (see section 24).

25.1 What is an eDrive?

A disk that returns the following value in its IDENTIFY DEVICE DATA ATA response:
Bit 0 of Word 48 =1(Trusted Computing Group Support): Opal SSC Specification.
Bit 7 of Word 69 =1(IEEE 1667 Support): ACS-3 Specification Rev. 5

25.2 SCSI Standard Inquiry Command Response

When a disk is identified as an eDrive, the Intel® RST driver will report IEEE 1667 support in a SCSI Standard Inquiry command response.

25.3 Configurations Supported

eDrive support is similar to Opal Support and will be supported in RAID or AHCI modes on a pass-thru (non-RAID member) drive attached to the SATA controller.

Further details and specifications for eDrives can be found at the IEEE 1667 Workgroup web page: http://www.ieee1667.com

§ §

26 Using the BCFS to Differentiate Platform SKUs

Beginning with the Intel® RST 10.x Release and the Intel® 5 Series Express Chipset (codename Ibex Peak), the BIOS Control Feature Set (BCFS) has been enabled to give OEMs the opportunity to customize the Intel® RST features offered on any particular Intel® 5 Series Express Chipset model/SKU and later. OEMs no longer need special Intel® RST OROM images from Intel in order to enable/disable certain desired features for a platform SKU. OEMs can now enable/disable the desired features per platform SKU directly in their BIOS code. By clearing or setting the corresponding bits of the '***Intel RST Feature Capabilities***' register in the Intel chipset's SATA controller MMIO space, OEMs now have greater flexibility in determining what Intel® RST features will be supported per platform model/SKU.

The following sections explain the use of each of the bits in the BCFS, also known as the Software Feature Mask bits.

Note: This document does not cover details on how to setup a system BIOS. For that level of information contact your platform's BIOS vendor or your Intel field representative to put you in contact with the appropriate Intel BIOS support personnel.

26.1 Configuring the Platform's RAID Related Features

When the BIOS has set the SATA Controller's mode to RAID, the following bits of the 'Intel RST Feature Capabilities' register in the Intel chipset's SATA controller MMIO space will determine what RAID levels will be supported on the platform SKU:

Note: Clearing all RAID level related bits to '0' (that includes the Intel® RRT bit) is an unsupported configuration. The Intel® RST OROM will ignore the BIOS settings and enable all RAID levels (Intel® RRT inclusive).

26.1.1 Configuring the Standard Supported RAID Levels

There are four (4) bits that control the 4 standard RAID levels supported by Intel® Rapid Storage Technology:

Bits	Type	Reset/ Default	Description
3	RWO	1h	**RAID 5 Enable (R5E):** If set to '1', then RAID5 is enabled
2	RWO	1h	**R10 Enable (R10E):** If set to '1', then RAID10 is enabled
1	RWO	1h	**RAID 1 Enable (R1E):** If set to '1', then RAID1 is enabled
0	RWO	1h	**R 0 Enable (R0):** If set to '1', then RAID0 is enabled

26.1.1.1 Example Configuration

To configure a platform SKU that offers **only** RAID levels 0 and 10, the bits must be configured as follows:

Bit 0 == 1 (default)

Bit 1 == 0

Bit 2 == 1 (default)

Bit 3 == 0

26.1.2 Configuring Intel® RRT Related RAID Features

There are two (2) bits that control two capabilities/features related to the Intel® RRT feature:

Bits	Type	Reset/ Default	Description
8	RWO	0h	**Intel® RRT Only on ESATA (ROES):** If set to '1', then only Intel® RRT volumes can span internal and eSATA drives. If cleared to '0', then any RAID volume can span internal and eSATA drives.
4	RWO	1h	**Intel® RRT Enable (RSTE)**: If set to '1', then Intel Rapid Recovery Technology is enabled.

26.1.2.1 Example Configuration

To configure a platform SKU that offers Intel® RRT, only RAID level 5, and allows both Intel® RRT and RAID5 Volumes to span disks on both internal and external (eSATA) ports, the bits must be configured as follows (**Note**: an Intel® RRT volume cannot coexist with another RAID level volume at the same time on the platform):

Bit 0 == 0

Bit 1 == 0

Bit 2 == 0

Bit 3 == 1 (default)

Bit 4 == 1 (default)

Bit 8 == 0 (default)

26.1.3 Configuring the Behavior of the OROM UI and Banner

There are three (3) bits that control the behavior of the Intel® RST OROM UI and the Banner Splash Screen that are displayed during POST at system boot-up. Use the following bit configurations to determine whether or not the splash screen will be displayed during post and if so, how long the delay will be before the system continues the boot process:

Bits	Type	Reset/ Default	Description
11:10	RWO	0h	**OROM UI Normal Delay (OUD)**: Values of these two bits specify the delay of the OROM UI Splash Screen in a normal status. 00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs **Note:** If bit 5 == 0, then these values are disregarded **Comment:** Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.
5	RWO	1h	**Intel RST OROM UI (RSTOROMUI):** If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.

26.1.3.1 Example Configuration

To configure a platform SKU that enables the OROM Banner Splash Screen to be displayed for 6 seconds, the bits must be configured as follows:

Bit 5 == 1 (default)

Bit '10' == 0 (default)

Bit '11' = = 1

26.1.4 Configuring Intel® RST UI Capabilities

There is one capability within the Intel® RST UI that is controlled by the BCFS bits. To enable/disable the ability for the Intel® RST UI to unlock password protected disks, use the following bit configurations:

Bits	Type	Reset/ Default	Description
6	RWO	0h	**HDD Unlock (HDDLK):** If set to '1', then HDD password unlock is enabled in the OS.

The default settings for these two features are:

Bit 6 == 0

When this bit is cleared the Intel® RST UI does not display any option to use this feature.

26.1.4.1 Example Configuration

To configure a platform SKU to not allow unlocking passwords from the Intel® RST UI and to allow the UI to activate the disk/port LEDs, the bits must be configured as follows:

Bit 6 == 0 (default)

26.1.5 Configuring the Platform to Support Intel® SRT (EOL'd)

The BCFS bit is only one of the platform parameters (see section 17.1.1 for all requirements) that determines whether Intel® Smart Response Technology can be enabled on a platform, however, unless this bit is set there will be no support for this feature on the platform.

Bits	Type	Reset/ Default	Description
9	RWO	0h	**Intel® Smart Response Technology Enable Request (SEREQ):** If set to '1', then Smart Response Technology is enabled. If cleared to '0', Smart Response Technology is disabled.

Bit 9 == 1

This enables the Intel® Smart Response Technology feature on the platform SKU.

26.1.6 Configuring the Platform to Support Intel® Optane™ Memory System Acceleration

The BCFS bit is only one of the platform parameters (see Optane(TM) Requirements for all requirements) that determines whether Intel® Optane™ Memory System Acceleration can be enabled on a platform, however, unless this bit is set there will be no support for this feature on the platform.

Bits	Type	Reset/ Default	Description
13	RWO	0h	**Intel® Optane™ Memory Acceleration (Optane_Memory_Accel):** If set to '1', then the computer will be capable of enabling System Acceleration with Optane Memory. If cleared to '0', then the computer will not be capable of enabling System Acceleration with Optane Memory. **Note**: This feature is added in KBL platform and onward.

Bit 13 == 1

This enables the Intel® Optane™ Memory System Acceleration feature on the platform SKU.

26.1.7 Configuring the Platform to Support CPU Attached Storage

The BCFS bit is only one of the platform parameters that determines whether Intel® RST support for CPU Attached Storage can be enabled on a platform, however, unless this bit is set there will be no support for this feature on the platform.

Bits	Type	Reset/ Default	Description
14	RWO	0h	**CPU Attached Storage (CPUAS):** If set to '1', then the computer will be capable of Intel® RST supported PCIe NVMe drives attached to the CPU PEG ports. If cleared to '0', then the computer will not be capable of enabling Intel® RST supported PCIe CPU Attached storage. **Note**: This feature is added in CNL/CFL platform and onward.

Bit 14 == 1

This enables Intel® RST supported PCIe NVMe drives attached to the CPU PEG ports on the platform SKU.

26.2 BIOS Control Feature Set

Bits	Type	Reset	Description
15	RO	0h	Reserved
14	RO	0h	**CPU Attached Storage (CPUAS):** If set to '1', then the computer will be capable of Intel® RST supported PCIe NVMe drives attached to the CPU PEG ports. If cleared to '0', then the computer will not be capable of enabling Intel® RST supported PCIe CPU Attached storage. **Note**: This feature is added in CNL/CFL platform and onward.
13	RWO	0h	**Intel® Optane™ Memory Acceleration (Optane_Memory_Accel):** If set to '1', then the computer will be capable of enabling System Acceleration with Optane Memory. If cleared to '0', then the computer will not be capable of enabling System Acceleration with Optane Memory. **Note**: This feature is added in KBL platform and onward.
12			Reserved
11:10	RWO	0h	**OROM UI Normal Delay (OUD)**: Values of these bits specify the delay of the OROM UI Splash Screen in a normal status. 00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs If bit 5 == 0, then these values are disregarded **Comment: Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.**
9	RWO	0h	**Intel® Smart Response Technology Request Enable (SEREQ):** If set to '1', then Smart Response Technology is enabled. If cleared to '0', Smart Response Technology is disabled.

Bits	Type	Reset	Description
8	RWO	0h	**Intel® RRT Only on ESATA (ROES):** If set to '1', then only Intel® RRT volumes can span internal and eSATA drives. If cleared to '0', then any RAID volume can span internal and eSATA drives.
7	RWO	0h	**Reserved**
6	RWO	0h	**Reserved**
5	RWO	1h	**Intel® RST OROM UI (RSTOROMUI):** If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.
4	RWO	1h	**Intel® RRT Enable (RSTE)**: If set to '1', then Intel Rapid Recovery Technology is enabled
3	RWO	1h	**RAID 5 Enable (R5E):** If set to '1', then RAID5 is enabled
2	RWO	1h	**RAID 10 Enable (R10E):** If set to '1', then RAID10 is enabled
1	RWO	1h	**RAID 1 Enable (R1E):** If set to '1', then RAID1 is enabled
0	RWO	1h	**RAID 0 Enable (R0E):** If set to '1', then RAID0 is enabled

§ §

27 Testing, Certification Notes

Correcting Microsoft* Windows 7 (Win7) WHQL test failure

For Internal SATA ports with interlock switches, the RST driver will set Removable=TRUE in the IRP_MN_QUERY_CAPABILITIES handler. This causes Win7 to show the internal device in its own "container" which is used to describe devices that are external to the system. For example, a CD-ROM on an interlocked switch in a Win7 system, under 'Devices and Printers',

You can see that the CD-ROM on the interlocked SATA port shows up separately in the top-level 'Devices' view. This can result in a platform WHQL test failure.

There is a whitepaper describing use of Removable device capability bits on Win7 by bus drivers: http://www.microsoft.com/whdc/Device/DeviceExperience/ContainerIDs.mspx.

In order to correct this issue to pass the platform WHQL test, RST recommends the OEM to take the following action:

In the system BIOS, define an _EJ0 ACPI method on the interlocked port. _EJ0 will signal to the ACPI driver to set Removable for the RST driver and still mark the device as internal to the system such that it does not show in its own container. The implementation is to use a registry key for each port to tell RST whether to set Removable bit or not. If _EJ0 ACPI method is defined in the system BIOS by the manufacturers, they can tell RST not to set the Removable bit. For example:

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\Port1]

"EJ0IsDefined"=dword:1

If 1, _EJ0 will set Removable bit instead of RST. If 0, no _EJ0 defined so RST will set Removable bit. The default value is 0.

§ §

28 Glossary

Term	Definition
ATA	Advanced Technology Attachment
BIOS	Basic Input/Output System
BOM	Bill Of Materials
CD	Compact Disc
Chipset	Term used to define a collection of integrated components required to make a PC function.
Hard drives	Physical hard drives attached to a RAID controller
DEVSLP	Serial ATA Device Sleep
DOS	Disk Operating System
GB	Giga-byte
HDD	Hard Drive
I/O	Input/Output
ICH	I/O Controller Hub
ICH9	Intel® 82801IR/DO SATA RAID Controller
IDE	Integrated Drive Electronics
INF	Information file (.inf) used by Microsoft operating systems that support the Plug and Play feature. When installing a driver, this file provides the OS needed information about driver filenames, driver components, and supported hardware.
Intel® Option ROM (OROM)	Standard Plug and Play option ROM that provides a pre-operating system user interface for the Intel RAID implementation.
MB	Mega-byte
Migration	Term used to describe the movement of data from one configuration or usage model to another.
OEM	Original Equipment Manufacturer
Option ROM	A code module built into the System BIOS that provides extended support for a particular piece of hardware. For this product, the Option ROM provides boot support for RAID 0/1/5/10 volumes, and provides a user interface for configuring and managing RAID 0/1/5/10 volumes.
OS	Operating System
PATA	Parallel ATA
PCH	Platform Controller Hub is the new term for Intel chipsets
PCI	Peripheral Components Interconnect
PFW	Package for the Web

Term	Definition
PIO	Programmed Input Output
PnP	Plug and Play
Port 0..3	Term used to describe the point at which a SATA drive is physically connected to the SATA Controller. Port n is the nth of the four available ports in ICH9 systems, where n=0..3
RAID	Redundant Array of Independent Disks
RAID 0	A RAID level where data is striped across multiple physical hard drives (aka striping)
RAID 1	A RAID level where data is mirrored between hard drives to provide data redundancy (aka mirroring)
RAID 5	A RAID level where data and parity are stripped across the hard drives to provide good read/write performance and data redundancy. The parity is stripped in a rotating sequence (aka Stripping and rotating parity).
RAID 10	A RAID level where information is striped across a two disk array for system performance. Each of the drives in the array has a mirror for fault tolerance. (aka Stripping and mirroring)
RAID volume	A block of capacity allocated from a RAID Array and arranged into a RAID topology. Operating Systems typically interpret a RAID volume as a physical hard drive.
RAM	Random Access Memory. Usually refers to the system's main memory
ROM	Read Only Memory
RST	Intel® Rapid Storage Technology
RTD3	Runtime D3
SATA	Serial ATA
SCSI	Small Computer System Interface
SP#	Service Pack (number)
Strip	Grouping of data on a single physical hard drive within a RAID volume
SRT	Intel® Smart Response Technology
SSD	Solid State Drive
SSHD	Solid-State Hybrid Drive
Stripe	The sum of all strips in a horizontal axis across physical hard drives within a RAID volume
UI	User Interface

Glossary

A B C D E F H I L M N O P R S U V W

- A
  - **Accelerated disk or volume**
    A disk or RAID volume that has its non-volatile cache enabled in either maximized or enhanced mode.
  - **Access master or recovery disk files:**
    Action to view the files located on the master or recovery disk of a recovery volume using Windows Explorer*. Selecting this option sets the volume in read-only and volume updates are not allowed. This action is only available when the recovery volume is in on request update mode or running off the recovery disk.
  - **Activate port LED:**

Action to locate the port connected to a disk present on the system by activating the Light Emitting Diode (LED) light.

- **AHCI:**
  An interface specification that allows software to communicate with SATA devices such as host bus adapters, and enables advanced SATA features such as Native Command Queuing, native hot plugging, and power management. Advanced Host Controller Interface (AHCI).
- **Array:**
  An abstraction layer or collection of two or more disks used to manage RAID volumes existing on a same collection of disks. RAID arrays are not visible to the operating system.
- **At risk:**
  Status indicative that a disk or device has experienced a SMART event, and that an impending error condition was detected and the disk or device is now at risk of failure.
- **ATAPI device:**
  A mass storage device with a parallel interface such as CD-ROM, CD-RW, DVD-ROM, Blu-ray Disc, and tape drives. Advanced Technology Attachment Packet Interface (ATAPI).

· B

- **Blocks with media errors:**
  Number of inconsistencies found during the data verification of a RAID volume. This feature only applies to the verification process or the verification and repair process.

Bus Protocol Group:
A bus protocol group represents a set of bus protocols with similar performance characteristic. Bus Protocol Groups are listed here in ascending order of speed:1
1- SATA
2- PCIe

· C

- **Cache**
  A resource allocation on a storage component used for temporary data operations. Cache can be allocated in components such as RAM or non-volatile memory.
- **Cache device**
  The selected solid state disk used for cache storage in an SRT configuration.
- **Cache volume**
  The portion of the cache device (a solid state disk) that holds the non-volatile cache data in an SRT configuration. That portion is configured into a single-disk RAID-0 volume.
- **Change volume type:**
  Action to change the volume from one RAID configuration to another, and move data from one RAID volume to another. A RAID 1 volume can also be converted to a recovery volume, and vice versa.
- **Continuous update mode:**
  Update mode assigned to a recovery volume, where data on the master disk is copied to the recovery disk automatically, as long as both disks are connected to the system.

· D

- **Data stripe size:**
  Size of a grouping of data on a single physical disk within a RAID volume. Reported in kilobytes (KB).
- **Data volume**
  The portion of extra space on the cache device (a solid state disk) that can be used for data storage. That portion is configured into a single-disk RAID-0 volume.
- **Degraded:**
  Volume status indicative that one member has failed or is missing. This status only applies to recovery, RAID 1, RAID 5, and RAID 10 volumes.
- **Disassociating a cache**
  The action of removing the association between the non-volatile cache and the accelerated disk or volume.
- **Disk data cache:**
  A cache memory within a hard drive that temporarily stores frequently used data sectors for faster access. As a result, overall hard drive performance is improved.
- **Disk:**
  A hard or floppy disk. Also known as hard drive or hard disk drive.

- E
  - **E-mail notification:**
    Alert mechanism that allows the user to receive storage system information, warning, and error notifications by e-mail via SMTP. By default, this feature is disabled and requires configuration settings such as the SMTP host and e-mail addresses to be set up.
  - **Enhanced Mode**
    An acceleration mode that uses write-through, non-volatile cache to improve performance. The mode also is known as “separation safe” because all host-write requests are written to the accelerated disk or volume and possibly to the non-volatile cache.
- F
  - **Failed:**
    Volume and disk status indicative that one or more array members are missing or have failed.
  - **Firmware:**
    Permanent instructions and data programmed directly into the read-only memory (ROM) for controlling the operation of the computer. Firmware usually requires updates to fix defects or add features to the hardware.
- H
  - **Hide master or recovery disk files:**
    Action to close the display of files located on a master or recovery disk in a recovery volume after viewing them in Windows Explorer*. This option is only available when ‘Access recovery disk files’ or ’Access master disk files’ was previously selected. Once disk files are hidden, volume updates can resume.
  - **Hot plug:**
    Action to remove or insert a SATA disk when the system is powered on.
- I
  - **Increase volume size:**
    Action to expand the data storage capacity of a volume by utilizing the available array space on a RAID 0, RAID 1, RAID 5, or RAID 10 volume.
  - **Initialize:**
    Process of synchronizing all redundant data on a volume prior to creating a volume, verifying and repairing data, or changing volume type. Initialization is still required for non-redundant volumes such as RAID 0 to ensure that data is readable before starting the verification process.
  - **Intel® Rapid Recover Technology:**
    Official name for Intel’s technology that allows the user to copy data from a master disk (source) to a recovery disk (destination) either continuously or on request.
  - **Intel® Rapid Storage Technology:**
    Official name for Intel's Windows-based software to provide support for high-performance, fault-tolerant, and capacity SATA RAID arrays on select supported chipsets. Intel Rapid Storage Technology also provides support for Intel® Rapid Recover Technology, AHCI Native Command Queuing, and matrix RAID for two RAID volumes on a single array.
- L
  - **Locked:**
    Volume and disk status indicative that the data is protected with a password and cannot be accessed until disks are unlocked.
- M
  - **Mark as spare:**
    Action to designate an available and compatible SATA disk as the default destination for automatic rebuilds in the event that an array member fails or is missing.
  - **Master disk:**
    The disk that is the designated source drive in a recovery volume.
  - **Maximized Mode**
    An acceleration mode that uses write-back, non-volatile cache to improve performance better than the enhanced Mode. This mode is optimized for input/output performance and power savings.
    Memory Group
    A memory group represents a set of backend storage media types with similar performance characteristics. Memory Groups are listed here in ascending order of speed:
    1- Spindle Device (HDD)
    2- NAND Spindle Hybrid Device (SSHD)
    3- PCH SATA NAND Device (SSD)
    4- PCIe NAND Device (SSD)
  - **Migrating:**

Volume status indicative that data is being moved/transferred across selected storage devices due to a change request in the storage system configuration, such as changing volume type, creating a volume preserving existing data, increasing the volume capacity, or changing data stripe size.

- N
  - **Native Command Queuing:**
    Command protocol in SATA that allows multiple commands to be outstanding within a disk at the same time. The commands are dynamically reordered to increase disk performance.
  - **Normal:**
    Volume, disk, and device status indicative that they are in a healthy state, functioning as expected, disks are properly connected, and data is fully accessible.
- O
  - **Offline:**
    Disk status indicative that an array disk is locked (the volume status displays as locked), that the recovery volume is in on request update mode, or that your computer is running on battery and data updates to the recovery volume are not available.
  - **On request update mode:**
    Update mode assigned to a recovery volume, where data on the master disk is copied to the recovery disk when the user requests it. Only changes since the last update process are saved to the recovery disk.
  - **Option ROM:**
    Firmware that is called by the system BIOS in order to communicate and provide support for a hardware device. For this product, the option ROM provides boot support for RAID volumes as well as a user interface for configuring and managing RAID volumes. Also known as OROM.
- P
  - **PCIe Storage Device:**
    A storage device that connects to the PCIe bus. PCIe devices can contain either an AHCI or NVMe communication controller.
    Port:
    An internal or external data connection of a computer (e.g., SATA controller) to which a peripheral device (e.g., SATA disk) can be attached.
- R
  - **RAID 0:**
    A RAID type or configuration where data is striped across multiple physical disks. Data is split into manageable blocks called strips across array members. Striping does not create data redundancy but improves read/write performance.
  - **RAID 1:**
    A RAID type or configuration where data is mirrored across a second physical disk in the array. Mirroring is a key feature that ensures real-time data redundancy and increased fault tolerance. There is no striping.
  - **RAID 10:**
    A RAID type or configuration that uses four disks to create a combination of RAID type 0 and 1. The data is striped across a two-disk array forming a RAID 0 component. Each of the disks in the RAID 0 array is mirrored by a disk in the RAID 1 array.
  - **RAID 5:**
    A RAID type or configuration where data and parity are striped into manageable blocks called strips across three or more physical disks. This type is a preferred configuration as it combines efficiency, fault-tolerance, and data performance.
  - **RAID:**
    Redundant Array of Independent/Inexpensive Disks is a technology used for computer data storage schemes that divide and/or replicate data among multiple disks. RAID can be designed to provide increased data reliability or increased I/O (input/output) performance, or both. A number of standard schemes have evolved which are referred to as levels or types. Intel® Rapid Storage Technology software supports RAID 0, RAID 1, RAID 5, and RAID 10 configurations (refer to each RAID type definition for more information).
  - **Read-only:**

Recovery volume status indicative that the recovery or master disk files are accessed, allowing files on the disk to be read or copied, but not changed or saved.

- **Rebuild:**
  The process of restoring a recovery, RAID 1, RAID 5, and RAID 10 volume in the event that a volume disk has failed or is missing. If a spare disk is present and compatible, the application will automatically use it as a replacement for the failed disk. An automatic rebuild process will also occur if a RAID 1 member is removed and then reconnected, in order to re-establish the mirroring. This process does not apply to RAID 0 volumes.
- **Recover data:**
  The action of retrieving data in the event that a recovery volume has failed due to a missing or failed master disk. Data from the recovery disk is copied to a new or healthy master disk, restoring data redundancy. Selecting this option will overwrite all master disk data with data on the recovery disk.
- **Recovery disk:**
  The disk that is the designated destination drive in a recovery volume.
- **Recovery volume:**
  A two-disk redundant volume that includes a master disk (source) and a recovery disk (destination) and uses Intel® Rapid Recover Technology. This configuration provides flexibility of volume data updates and maximum data redundancy.
  Remapped:
  Remapping hardware in the chipset allows a PCIe STORAGE DEVICE to appear to the OS as if it were an extra port on the Chipset's internal AHCI controller rather than a separate controller on the PCIe bus
- **Repair:**
  The process of fixing verification errors and blocks with media errors found during the verification process. This feature is only available for volumes with a normal or at risk status. A RAID 0 cannot be repaired due to its non-redundant configuration.
- **Reset disk to normal:**
  Action to return a failed or at risk disk to a normal state. We recommend that you contact the manufacturer for more information to prevent potential data loss.
- **Reset to available:**
  Action to return a disk previously marked as a spare to an available state.
- **Reset volume to normal:**
  Action to return a failed volume where both array disks are present and normal to a normal state. This feature allows you to access and attempt a recovery of healthy volume data.

· S

- **SATA disk:**
  A disk with an interface that transmits data using a serial protocol in order to communicate with the SATA controller.
- **SATA transfer rate:**
  Rate at which the SATA controller and SATA disk communicate with each other. Transfer rates are important when large contiguous blocks of data are being used, such as video and image files. Reported in gigabits/seconds (Gb/s).
- **SATA:**
  A successor to ATA and PATA, SATA is a computer bus technology primarily designed for transfer of data between storage devices such as hard drives or optical devices, and a computer. Benefits of this technology are: usage of high-speed serial cables, air cooling to work more efficiently, faster transfers, the ability to remove devices while operating (hot plugging), enables more reliable operation with tighter data integrity checks. Also known as Serial Advanced Technology Attachment or Serial ATA.
- **Single-disk RAID 0**
  A RAID-0 volume that has one (1) disk as its array disk. This is the volume type used to create the non-volatile cache region on a solid state disk being used for caching. This volume type creates an OS-visible volume that enables access to extra space on the solid state disk that is not being used for caching.
- **Size:**
  Reports the total capacity of a physical device such as a SATA disk, a volume, or an array.
- **Storage system:**
  One or more physical disks or devices that act as a unit for data storage.
- **System disk:**
  A disk that contains system files required to start and run the operating system. By default, the Windows operating system files are in the WINDOWS folder, and the supporting files are in the WINDOWS\System32 folder.
- **System volume:**

A volume that refers to the disk(s) volume that contains the hardware-specific files that are needed to start Windows, such as Ntldr, Boot.ini, and Ntdetect.com, as well as the Windows operating system files and supporting files. The system volume can be the same volume as the boot volume.

- U
  - **Unknown:**
    Disk status indicative that its usage could not be determined, due to a possible incompatibility between this software version and the disk configuration, or a virus. Also a volume status indicative that the volume is in an unexpected state due to a configuration error. Data on the volume can no longer be accessed.
  - **Update mode:**
    Type of update assigned to a recovery volume. The update mode can be set to continuous, where master disk changes are automatically saved to the recovery disk, or set to on request, where updates of the recovery disk can be requested immediately.
- V
  - **Verify:**
    Action of scanning data to detect any types of data damage, disk read errors, and volume data inconsistencies. As an option, errors found can be corrected on redundant RAID volumes. This feature is only available for volumes with a normal or at risk status.
  - **Volume size:**
    Amount of data that can be stored on a volume; reported in bytes (B) or %.
  - **Volume type:**
    Configuration of a volume which determines how data is stored and managed to improve read/write performance, increase fault tolerance and/or storage capacity.
  - **Volume write-back cache:**
    A cache memory used to enhance the read/write performance of a RAID volume by grouping multiple I/O requests into fewer requests and by writing from the cache to the volume at defined intervals.
  - **Volume:**
    A fixed amount of space across a RAID array that is structured to emulate a single physical hard drive and appears as such to the operating system. Volumes have drive letters assigned to them and some volumes can span multiple hard disks.
- W
  - **Windows Disk Management*:**
    Microsoft Windows* system utility for managing the disks and partitions or volumes that they contain. Disk Management allows the initialization of new disks, volume creation, and formatting. Most disk-related tasks can be performed using this system utility without shutting down or restarting the computer; most configuration changes take effect immediately.
  - **Write-back cache allocation:**
    Size of the dynamic random access memory (DRAM) that is allocated for write-back caching on all volumes present on the system. The cache size is set to 16 megabytes (MB) by default.

§ §

29 Troubleshooting

This section explains how to resolve the most common problems that may occur while using the application. If you have any questions regarding installing, using or maintaining this product, you can also visit Intel's online support site which provides you with self-help resources and electronic problem submission.

29.1 Failed Volumes

RAID 0

A RAID 0 volume is reported as failed when one of its members is disconnected or has failed. In both cases, the volume and its data are no longer accessible.

Cause	Solution
Missing array disk	Follow this procedure to recover data: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer. During the system startup, the volume status will display as 'Normal' in the Intel Rapid Storage Technology option ROM user interface. 3. Once the operating system is running, open Intel Rapid Storage Technology from the Start menu or click the Intel Rapid Storage Technology notification area icon. 4. Under 'Status', verify that the volume and disks status display as 'Normal'. You can also review this information under 'Manage'.
Failed array disk	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disk with a new one that is of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 0 volume using the new disk. If the failed disk was part of the system volume, you will also need to reinstall the operating system.

RAID 5

A RAID 5 volume is reported as failed when two or more of its members have failed.	
Cause	**Solution**
Two or more array disks failed	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 5 volume using the new disks. If the failed disk was part of the system volume, you will also need to reinstall the operating system.

RAID 10 A RAID 10 volume is reported as failed when two adjacent members are disconnected or have failed, or when three or four of its members are disconnected or have failed.	
Cause	**Solution**
Two adjacent array disks missing	1. Power off your computer and reconnect the missing disks. 2. The rebuild operation will start automatically. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.
Three or four array disks missing	In most cases, the volume cannot be recovered and any data on the volume is lost. This procedure deletes the failed volume: 1. Power off your computer and reconnect the missing disks. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.

	6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. You will then need to reinstall the operating system on the new volume.
Two or more array disks failed	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. You will then need to reinstall the operating system on the new volume.

29.2 Degraded Volumes

Recovery Volume A recovery volume is reported as degraded when the recovery disk has failed or when the master disk is disconnected or has failed. Data mirroring and redundancy are lost because the system can only use the functional member.	
Cause	**Solution**
Recovery disk failed	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, but if the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the recovery disk and volume status will display as 'Normal'. 5. Once completed, the volume returns to the last update mode to which it was set before the issue was reported.

	Note: If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Master disk missing	If you can reconnect the missing master disk, follow this procedure to recover data: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer and the system will automatically boot from the recovery disk. 3. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'. 4. Click 'Yes' to confirm. 5. The recovery operation starts immediately and cannot be canceled. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 6. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. If you cannot reconnect the missing disk and a SATA disk is available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. 5. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. If you cannot reconnect the missing disk or rebuild to an available disk, you will need to power off the computer and connect a new SATA disk. Once rebuilt, the recovery volume will be limited to its original size even if the new disk is larger than the original master disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Master disk failed	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, but if the read/write data access consistently fails, the disk will likely return to a failed state immediately. 1. To reset the failed master disk and the volume to normal, follow this procedure: 2. Under 'Status', click 'Reset disk to normal'. Note that the volume is now running off the recovery disk, and that the master disk is reported as offline.

	3. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'. **Warning** Starting this action will override existing data on the master disk and update it with the data on the recovery disk. Backup all valuable data before continuing. Click 'Yes' to confirm. The recovery operation starts immediately and cannot be canceled. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: Under 'Status', click 'Rebuild to another disk'. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. **Note** If there is no available disk present, you will need to power off your computer and connect a new SATA disk. Once rebuilt, the recovery volume will be limited to its original size even if the new disk is larger than the original master disk. Once your computer is back up and running you can follow the rebuild procedure described above.

RAID 1 A RAID 1 volume is reported as degraded when one of its members is disconnected or has failed. Data mirroring and redundancy are lost because the system can only use the functional member.	
RAID 5 A RAID 5 volume is reported as degraded when one of its members is disconnected or has failed. When two or more array disks are disconnected or have failed, the volume is reported as failed.	
RAID 10 A RAID 10 volume is reported as degraded when one of its members is disconnected or has failed, or when two non-adjacent members are disconnected or have fails. When two or more array disks are disconnected or have failed, the volume is reported as failed.	
Cause	**Solution**
Missing array disk	If you can reconnect the missing disk, follow this procedure to rebuild the volume: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer and the rebuild operation will start automatically. If you cannot reconnect the missing disk and a SATA disk is available and normal, follow this procedure to rebuild the volume:

	If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. **Note:** If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Failed array disk	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, which will prompt the volume to start rebuilding automatically. But if the read/write data access consistently fails, the disk will likely return to a failed state immediately and you will need to rebuild the volume to another disk. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. **Note:** If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once your computer is back up and running you can follow the rebuild procedure described above.

29.3 Other Volume States

Locked	
Cause	**Solution**
At least one (but not all) disk included in the volume is locked with a password.	In this state, the overall storage system health is still reported as normal, but we recommend that you unlock the disks to make the volume data fully accessible. Follow this procedure to unlock a disk: 1.Under 'Status' or 'Manage', in the storage system view, click the disk you want to unlock. The disk properties are now displayed on the left. 2.Click 'Unlock'. 3.Enter the password, and then click 'Unlock'.

	4.Repeat this procedure for all locked disks included in the volume in order to unlock the volume. **Note:** If all the disks included in a volume are locked, the volume is no longer displayed
Incompatible	
Cause	**Solution**
Indicates that the volume was moved to another system that does not support the volume type and configuration.	In this situation, volume data is accessible to the operating system and can be backed up, but the volume cannot operate because your system does not support its RAID configuration. Here are your options: 1.Reconnect the volume to the computer where the volume was originally created, and continue using it. 2.Delete the volume, and then create a new volume with a RAID configuration that is supported by the current system. Follow the procedure described above to delete the volume. **Warning** When a volume is deleted, all existing data on the member disks of the selected volume is permanently erased. It's recommended that you backup all valuable data prior to beginning this action.
Unknown	
Cause	**Solution**
The volume is in an unexpected state due to a configuration error.	The application is unable to detect the exact nature of the problem. Try restarting your computer. If the error persists, back up all valuable data and delete the volume using the option ROM user interface. Refer to the user's manual accessible from the Online Support area for details on using the option ROM.
Power-Saving Mode (Recovery Volumes only)	
Cause	**Solution**
Your computer is running on battery and the volume is in continuous update mode. Data updates to the recovery disk or a data recovery operation are not occurring.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. Reconnect your computer to the power supply, and the operation that was in progress prior to running on battery will resume automatically. **Note** If a data recovery was in progress, the overall storage subsystem health is reported as degraded because the operation could not be completed.
Data Update Needed (Recovery Volumes only)	
Cause	**Solution**
The data on the recovery disk is not synchronized with the data on the master disk.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. Follow this procedure to update data on the recovery disk: 1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2.Click 'Update data'. 3.Select the check box if you don't want this confirmation message to display each time you request an update. Click 'Yes' to confirm.

	4.The update process can be instantaneous or may take a while depending on the amount of data being copied. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.
Running Off Recovery Disk (Recovery Volumes only)	
Cause	**Solution**
Your computer was booted from the recovery disk using the option ROM, and the volume is operating from that disk. With this reverse configuration, the recovery disk is the designated source drive and data updates to the master disk are not available.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. A data recovery from the recovery disk to the master disk is required to maintain full redundancy. ⚠ **Warning** Completing this action will overwrite all master disk data with the data on the recovery disk. Backup all valuable data prior to starting this action. 1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2.Click 'Recover data', then 'OK' to confirm. 3.Once complete, we recommend that you restart your computer from the master disk using the option ROM user interface to return to a normal state.
Recovery Disk Read-Only (Recovery Volumes Only)	
Cause	**Solution**
The recovery disk files have been accessed and display in Windows Explorer*.	In this state, any data written to the master disk is not copied to the recovery disk because it is read-only. Data mirroring and redundancy may be lost and we recommend that you hide the recovery files to resume data updates. Follow this procedure to hide recovery disk files from Manage Disk or from Manage Volume: 1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the recovery disk. The element properties are now displayed on the left. 2.Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume. 3.The Windows Explorer window closes. You can resume data updates by clicking 'Update data' under Manage Volume. To copy the latest changes to the recovery disk automatically, change the update mode to continuous from the same area.

Master Disk Read-Only (Recovery Volumes Only)	
Cause	**Solution**
Your computer was booted from the recovery disk using the option ROM, and the volume is operating from that disk. The master disk files have been accessed and are displayed in Windows Explorer*.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. We recommend that you hide the master disk files when finished with your review, and proceed with a data recovery to the master disk in order to maintain full redundancy. Follow this procedure to hide master disk files from Manage Disk: 1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the master disk. The element properties are now displayed on the left. 2.Click 'Hide Files' from Manage Disk or 'Hide master disk files' from Manage Volume. The Windows Explorer window closes and the volume is displayed as running off recovery disk. Refer to the 'Running off recovery disk' procedure above to recover data to the master disk.
Missing volume	
Cause	**Solution**
A driver upgrade or downgrade was performed while a data migration was in progress.	The driver cannot recognize the volume or read its data if a driver upgrade or downgrade was performed during a volume migration. Volume migrations occur after one of the following operations was initiated: • Creation of a system volume or data volume while preserving data. • Volume type change combined with disk addition to the new RAID configuration. • Volume size increase. • Disk addition to an existing array. **Troubleshooting a data volume** If the data migration involved a data volume, you will need to reverse the driver upgrade or downgrade operation and return to the original driver version. This will restore driver and volume compatibility. Once the operation has completed, restart your computer. Open the application and make sure that the volume displays again in the storage system view. Data migration operation should resume immediately. **Troubleshooting a system disk** If the data migration involved a system disk or volume, it is highly likely that you will not be able to start your system because the driver cannot read the system files. The following options may allow you to load the operating system again: Restore a known good configuration. Boot from a flash drive that supports NTFS partitioning and includes the storage driver files. Bring the corrupt disk to another system, and then replace the storage driver files from a compatible driver version. Return the disk to the original system and try booting. **Troubleshooting a system volume** If the data migration involved a system disk or volume, it is highly likely that you will not be able to start your system because the driver cannot read the system files. The following options may allow you to load the operating system again: Restore a known good configuration.

	Bring all corrupted array disks to another system, and then replace the storage driver files from a compatible driver version. Return the disks to the original system and try booting.

29.4 Disk Events

State	Cause	Solution
	An internal or external disk is protected and locked with a password.	In this state, the overall storage system health is reported as normal, but to make the data fully accessible, you will need to follow this procedure to unlock the disk: 1.Under 'Status' or 'Manage', in the storage system view, click the disk you want to unlock. The disk properties are now displayed on the left. 2.Click 'Unlock'. 3.Enter the password, and then click 'Unlock'.
At risk 	An impending error condition was detected on an internal or external disk and is now at risk of failure.	The application is detecting early warning signs of failure with a SATA disk that result from a slow degradation over time. When a disk is reported at risk, you can reset that disk to normal, but we recommend that you contact the manufacturer for more information to prevent potential data loss. Follow this procedure to reset the disk to normal: 1.Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2.Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. **Note:** Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss. If the disk reported at risk is included in a RAID volume and a compatible spare disk is available, the rebuild process will start automatically. Once complete, the disk reported at risk becomes available and you can reset it to normal to return to a healthy state.
	An unexpected error was detected on a disk that has RAID configuration data (metadata) on it.	In this state, it is likely that some or all of the disk data is inaccessible. After backing up any accessible data, you will need to clear the metadata and reset the disk to return to a normal state. **Warning**: Completing this action will permanently delete existing metadata. Back up any accessible data before continuing.

State	Cause	Solution
		1.Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2.Click 'Clear and reset disk', and then click 'Yes' to confirm. 3.Once complete, the page refreshes with the disk returning to a normal state.
Missing 	An array disk is not present or physically connected to the computer.	Ensure that the disk is securely connected to the SATA port and that the SATA cable is functioning properly. If the disk is lost or cannot be reconnected, you will need to connect a new SATA disk, and then rebuild the volume to that new disk. Refer to Degraded or Failed Volumes in this section for instructions on how to rebuild a volume.
	The recovery or master disk files have been accessed and display in Windows Explorer*.	Hide the recovery or master disk files to return the disk status to offline and resume data updates in on request mode. 1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the recovery disk. The element properties are now displayed on the left. 2.Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume. 3.The Windows Explorer window closes.
Failed 	An internal or external disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.	Back up your data and we recommend that you replace the disk as soon as possible. If the failed disk is an array disk, the volume will be reported as degraded or failed depending on its configuration. Refer to Degraded or Failed Volumes in this section for instructions on resolving the problem. In a failed state, disk data may be lost, but you can try resetting the disk to normal, and then attempt a data recovery. Follow this procedure to reset the failed disk to normal: 1.Under 'Status', in the Manage subsection, locate the disk reported as failed. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2.Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. **Note** If the failed array disk is part of a redundant volume, the volume will start rebuilding automatically as soon as the disk is reset to normal.
Offline 	An internal or external array disk is locked and data on that disk cannot be read.	We recommend that you unlock the disk to make the volume data fully accessible. If more than one array disk is locked, unlock all those disks to unlock the volume.
	The recovery volume is in on request update mode.	Change the volume update mode to continuous.

State	Cause	Solution
		1.Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2.Click 'Change mode', and then click 'Yes' to confirm. 3.The page refreshes and the volume properties report the new update mode.
	Your computer is running on battery and data updates to the recovery disk are not available as long as that disk is offline.	Reconnect your computer to the power supply in order to return the recovery disk to a normal state.

29.5 Caching Issues

Cache Volume is Missing
Regardless of what the acceleration mode currently is, cache and volume data is most likely lost. Devices that are part of the acceleration configuration display in the following states: • Solid state disk: Inaccessible • Cache volume: No longer displays • Data volume (if applicable): No longer displays • Accelerated volume (if applicable): No longer displays • Accelerated disk (if applicable): Offline • Array disks: Offline
Cause
The solid state disk was removed or the disk is present but cannot be detected.
Solution
The application provides the option to clear the metadata on the array disks or previously accelerated disk and reset these disks to a normal state. 1.Under Status, in the Manage subsection, click 'Clear and reset' next to each array disk reported as offline. You can also perform this action under 'Manage' by clicking any offline disk reported in the storage system view. 2.Click 'Yes' to confirm. The array disk now displays as an available disk in a normal state and can be used to create a new volume

Cache Volume is Failing
Cause
An impending error condition (e.g., SMART events) was detected on the solid state disk that is used as a cache device. As a result, both the disk and cache volume are at risk of failure.
Solution
Early warning signs of failure with the solid state disk are detected that result from a slow degradation over time. When a disk used as a cache device is reported at risk, you can reset that disk to normal or replace the solid state disk after resetting it to available. Regardless of which option you choose, we recommend that you contact the manufacturer for more information to prevent potential data loss. Current acceleration mode: Enhanced Follow this procedure to reset the disk to normal: 1.Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the failing disk in the storage system view. 2.Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. 3.The cache volume should also return to a normal state and caching activity should resume. Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss. Follow this procedure to replace the failing solid state disk: 1.Click 'Accelerate', and then click 'Disable acceleration'. 2.In the dialog, click 'Yes' to confirm. 3.The page refreshes and reports the acceleration as disabled. 4.Click 'Reset to available'. 5.In the dialog, select the check box to confirm that you understand that data on the cache and data volumes will be deleted. 6.Click 'Yes' to confirm. 7.The page refreshes and the storage system displays the solid state disk usage as available. 8.Power off your computer and replace the failing solid state disk with a healthy one. 9.Power on your computer. You can enable acceleration again in order to resume caching activity. **Note:** If the last acceleration mode was maximized, that application will automatically transition acceleration to enhanced mode in order to avoid data loss. While transitioning, the mode will display as busy and no acceleration actions will be available until the process is complete. **Current acceleration mode: Maximized** 1.If a compatible spare is detected, the volume rebuild operation will start automatically. Once the process is complete, the cache volume will display in a normal state and caching activity will resume. 2.If no compatible spare is detected, the acceleration mode will automatically transition to enhanced in order to avoid data loss. You can then follow the procedures described above to return the solid state disk and cache volume to normal.

Cache Volume has Failed
Cause
The solid-state disk that is used as a cache device has failed to properly complete read and write operations in a timely manner and it has exceeded its recoverable error threshold. In this state, both cache and data volumes are reported as failed and acceleration is automatically disabled.
Solution
Back up any recoverable data and replace the solid- state disk as soon as possible. In a failed state, disk data may be lost, but you can try recovering it by resetting the disk to normal.
1.In the Manage subsection, under 'Status', locate the disk reported as failed. Alternately, perform this action from Manage Disk, accessible by clicking the disk in the storage system view. 2.Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. If the disk operations continue to fail, the disk will return to a failed state immediately and should be replaced. Follow this procedure: 1.Click 'Accelerate'. 2.Click 'Reset to available'. 3.In the dialog, select the check box to confirm that you understand that data on the cache and data volumes will be deleted. 4.Click 'Yes' to confirm. The page refreshes and the storage system displays the solid state disk usage as available. 1.Power off your computer and replace the failed solid state disk with an operational one. 2.Power on your computer. To resume the caching activity, enable acceleration again. If acceleration was in maximized mode prior to being automatically disabled, the disk or volume previously associated with the cache will be reported as failed if the data cleaning was unsuccessful. If data cleaning was successful, once the mode transition is complete, the accelerated disk or volume previously associated with the cache will be reported as normal.

Accelerated disk or volume is missing
Cause
The accelerated disk or volume cannot be detected (e.g., the device was moved to another system) and is reported as inaccessible. Caching activity is no longer occurring.
Solution
If the disk or volume can be reconnected: 1.Power off your computer and reconnect the missing disk or volume. 2.Restart your computer. 3.Once the operating system is running, open the application. 4.Under 'Status', in the Accelerate subsection, verify the accelerated device is properly reported as well as the acceleration mode. Caching activity should resume immediately.

If the disk or volume cannot be reconnected, follow this procedure to disassociate the cache and the missing device: 1.Click 'Accelerate'. 2.Click 'Disassociate'. 3.Click 'Yes' to confirm. 4.The page refreshes and you can now select another disk or volume to accelerate.

Accelerated Disk or Volume is Reporting an Issue
Cause
• Accelerated disk is reported as being at risk or failed. • Accelerated volume is reported as degraded or failed.
Solution
Refer to Troubleshooting Disk Events, Failed Volumes, or Degraded Volumes for detailed procedure on fixing the issue. If you cannot fix the issue reported on the accelerated disk or volume, follow this procedure to disassociate the cache and the missing device: 1.Click 'Accelerate'. 2.Click 'Disassociate'. 3.Click 'Yes' to confirm. The page refreshes and you can now select another disk or volume to accelerate.

29.6 Software Errors

Message	Cause	Solution
An unknown error occurred while running this application. If the problem persists, restart your computer or try reinstalling the application.	This error may be related to: Missing components Corrupted application Application unable to connect to the service Application fails to start.	Restart your computer or try reinstalling the application.
Intel® Rapid Storage Technology is trying to connect to the service.	The application is launched and is attempting to connect to the service in order to run.	If the connection succeeds, the application opens and is fully functional; if the connection fails, the error message described above is displayed. Try starting the service manually using Microsoft Windows* Services, or follow the recommended solution listed above to resolve the problem.
The Intel® Rapid Storage Technology service cannot be started in safe mode.	Your computer was started in safe mode and the operating system is running with a limited set of files and drivers. Intel Rapid Storage Technology cannot start or run in safe mode.	Once you are done troubleshooting application or driver problems in safe mode, you will need to exit safe mode, restart your computer, and then let the operating system start normally.

Message	Cause	Solution
		The Intel Rapid Storage Technology service can now be started and open the application.
Multiple users cannot run the application at the same time.	One or more users are attempting to open the application while an instance of the application is already running.	Make sure only one instance of the application is running at a time.
An error occurred due to insufficient resources, and the operation could not be completed. Try again later.	The Intel® Rapid Storage Technology driver does not have sufficient resources to execute the request. Another operation may be in progress and needs to complete before being able to handle a new request.	Wait a few moments, then try performing the action again.
An unknown error occurred during the volume creation process. Try recreating the volume.	An unexpected error occurred during the operation, and the application cannot identify its origin. The volume could not be created.	Verify that your hardware is properly connected and try recreating the volume.
An error occurred while an operation was in progress. The operation could not be completed.	An unexpected error occurred during an operation, such as a data migration or a rebuild, and the application cannot identify its origin.	Restart the operation. If the error persists, try restarting your computer and then the operation.
An error occurred and the selected disk or volume could not be accelerated. Restart your computer, and then try the operation again.	The cache memory allocation was likely increased to use full solid state disk capacity (up to 64 GB) while enabling acceleration.	Follow these steps to accelerate a disk or volume: • Restart your computer to complete the process of allocating the requested cache size. • Launch the application. • Try enabling acceleration again by clicking 'Enable acceleration'.



Appendix A RST SATA Port Bitmap Implementation

A.1 Legacy OROM

RST overloaded offset 0x08 of the PNP header, with a bitmap of which ports make up the PNP header.

E.g., a RAID5 volume whose member disks located on SATA ports 0,2,3 would have a value of 0x0D (0000_1101b) and a pass-thru disk on port 4 would have a value of 0x10 (0001_0000b).

Offset	Size	Value	Description
00h	1	'$'	Signature byte 1
01h	1	'P'	Signature byte 2
02h	1	'N'	Signature byte 3
03h	1	'P'	Signature byte 4
04h	1	01h	Structure revision
05h	1	Varies	Length (in 16 byte blocks)
06h	2	Varies	Offset of next header (0000h if none)
08h	**1**	**00h**	**Reserved**
09h	1	Varies	Checksum
0Ah	4	Varies	Device identifier
0Eh	2	Varies	Offset to manufacturer string (optional)
10h	2	Varies	Offset to product name string (optional)
12h	3	Varies	Device type code
15h	1	Varies	Device indicators
16h	2	Varies	Boot Connection Vector (BCV)
18h	2	Varies	Disconnect Vector (DV)
1Ah	2	Varies	Bootstrap Entry Vector (BEV)
1Ch	2	0000h	Reserved
1Eh	2	Varies	Static resource information vector

A.2 UEFI Driver

The RST UEFI driver, in an effort to provide similar functionality as in the legacy OROM, has implemented the Port Number value in the Device Path as a bitmap representing the physical disk

*connections that the Logical Disk represents. The LSB (least significant bit) represents port 0 and increases linearly. E.g. a single pass through disk on SATA port **3** (assuming the SATA ports are enumerated 0 – X) is represented by 0000_**1**000b (or 0x08).*

A.2.1 EFI_DEVICE_PATH_PROTOCOL

For each logical disk that is exposed by the SATA RAID UEFI Driver, an EFI_DEVICE_PATH_PROTOCOL shall be created.

The Device Path Protocol for each logical disk shall be appended to the PCI SATA Controller Device Path.

The fields of the EFI_DEVICE_PATH_PROTOCOL shall be filled out differently depending on whether the device is an ODD or an HDD.

EFI_DEVICE_PATH_PROTOCOL Field	ATAPI (ODD)	HDD/Volume – Logical Device
Type	3 (Messaging Device Path)	3 (Messaging Device Path)
Sub-Type	18 (SATA)	18 (SATA)
Length	10	10
HBA Port Number	Port ID bitmap (bit #n set if device is on port #n)	Port ID bitmap (bit #n set if logical device contains device ID #n) Lowest Significant Bit (LSB) represents port 0.
Port Multiplier Port Number	0x8000 (directly connected)	0x8000 (directly connected)
Logical Unit Number	0	0 for passthrough devices, myVolRaidDevNum for RAID volumes, which is the n^{th} volume created on the array.

§ §



Appendix B Common Storage Management Interface Support (CSMI)

Intel® RST driver does not support application development that require interface access to the driver via API method. However RST does support applications developed to interface to the driver via the industry standard specification of the Common Storage Management Interface. We support a subset of the total command set. The below table has a list of the commands that are supported by Intel® RST. For more detail information on the specification you can access the http://www.t11.org/ website. The document number for the specification is 04-468v0.

Supported CSMI commands:

BASE IOCTLs:

CC_CSMI_SAS_GET_DRIVER_INFO
CC_CSMI_SAS_GET_CNTLR_CONFIG
CC_CSMI_SAS_GET_CNTLR_STATUS

RAID IOCTLs:

CC_CSMI_SAS_GET_RAID_INFO
CC_CSMI_SAS_GET_RAID_CONFIG

The following CSMI IOCTLs are supported for both SATA and PCIe AHCI devices:

CC_CSMI_SAS_GET_PHY_INFO
CC_CSMI_SAS_GET_SATA_SIGNATURE
CC_CSMI_SAS_STP_PASSTHRU

§ §

Appendix C Drive and Volume Encryption Support

C.1 ATA Security Commands and HDD Password Support

RST supports HDD passwords via the ATA Security command set, as long as the host BIOS also supports the ATA security command set. RST volumes, including SRT volumes, may be locked with HDD passwords via the ATA Security command set. In the pre-boot environment, it is required that the BIOS obtain the HDD password before the RST components are loaded.

C.2 Self-Encrypting Drives (SED)

Self-encrypting drives are supported by RST under three scenarios:

1. When SED's are treated as standard drives.
2. When SED's are enabled using HDD passwords, via the ATA Security command set, following the restrictions outlined in the previous section above.
3. When SED's are Opal or eDrive enabled, using 3rd party software, and following configuration and support outlined in Sections Opal Drive Support and eDrive Support.

C.3 Solid State Hybrid Drives (SSHD's) with Encryption

SSHD's with encryption are supported by RST only if those drives can be treated as standard SSHD's.

C.4 RAID Volume and Drive Partition Encryption

RST supports encryption of volumes (including SRT), disks, and individual partitions by software encryption tools with the following restrictions:

1. Partition tables must remain unencrypted.
2. Approximately 6MB of non-partitioned space must be reserved, and remain unencrypted, at the end of the disk. This is the same as any configuration performing a create volume from an existing disk.

It is suggested that drives be encrypted before enabling SRT. This is to ensure that the SRT cache is populated with encrypted data.



§ §

Appendix D Remapping Guidelines for RST PCIe Storage Devices

D.1 Remapping Reference Documentation

Doc #	Title	Notes
549921	*Skylake Platform Controller Hub (PCH), H and LP BIOS Specification* http://www.intel.com/cd/edesign/library/asmo-na/eng/549921.htm	See Chapter 24
546884	*Skylake H Platform Design Guide Based on Skylake H Processor* -- http://www.intel.com/cd/edesign/library/asmo-na/eng/546884.htm	See Chapters 13, 14, 15, and 34
543016	*Skylake U and Y Platform Design Guide Based on Skylake U and Y Processors* -- http://www.intel.com/cd/edesign/library/asmo-na/eng/543016.htm	See Chapters 11, 12, 13, and 36
551377	*Skylake-H Client Platform SPI Programming Guide*	Download latest version ME firmware package for latest version See Chapter 10.1
550696	*Skylake-LP Client Platform SPI Programming Guide* -	Download latest version ME firmware package for latest version See Chapter 10.1
549921	*Skylake Platform Controller Hub (PCH), H and LP BIOS Specification* -- http://www.intel.com/cd/edesign/library/asmo-na/eng/549921.htm	See Chapter 24
546717	*Skylake H Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2* http://www.intel.com/cd/edesign/library/asmo-na/eng/546717.htm	*Tables 1-5, 25-3* *Chapter 3*
546955	*Skylake Platform Controller Hub (SKL PCH-H) External Design Specification (EDS) - Volume 2 of 2* http://www.intel.com/cd/edesign/library/asmo-na/eng/546955.htm	
545659	*Skylake Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2-* http://www.intel.com/cd/edesign/library/asmo-na/eng/545659.htm	*Tables 1-3, 25-3* *Chapter 3*

545845	*Skylake Platform Controller Hub (PCH-LP) External Design Specification (EDS) - Volume 2 of 2* http://www.intel.com/cd/edesign/library/asmo-na/eng/545845.htm	

D.1.1 Remapping HW and BIOS Requirements

1) Platform must have the appropriate HW to support remapping. Those PCH SKUs that support remapping:
   Skylake: PCH-LP: Premium-U, Premium-Y
   PCH-H: QM170, HM170, H170, Q170, Z170, C236, CM236
2) Also the system UEFI BIOS must have the appropritate coding to support the remapping configuration desired **(legacy BIOS not supported)**

Reference: Doc #549921: *Skylake Platform Controller Hub (PCH), H and LP BIOS Specification*

D.1.2 Remapping Configuration Rules

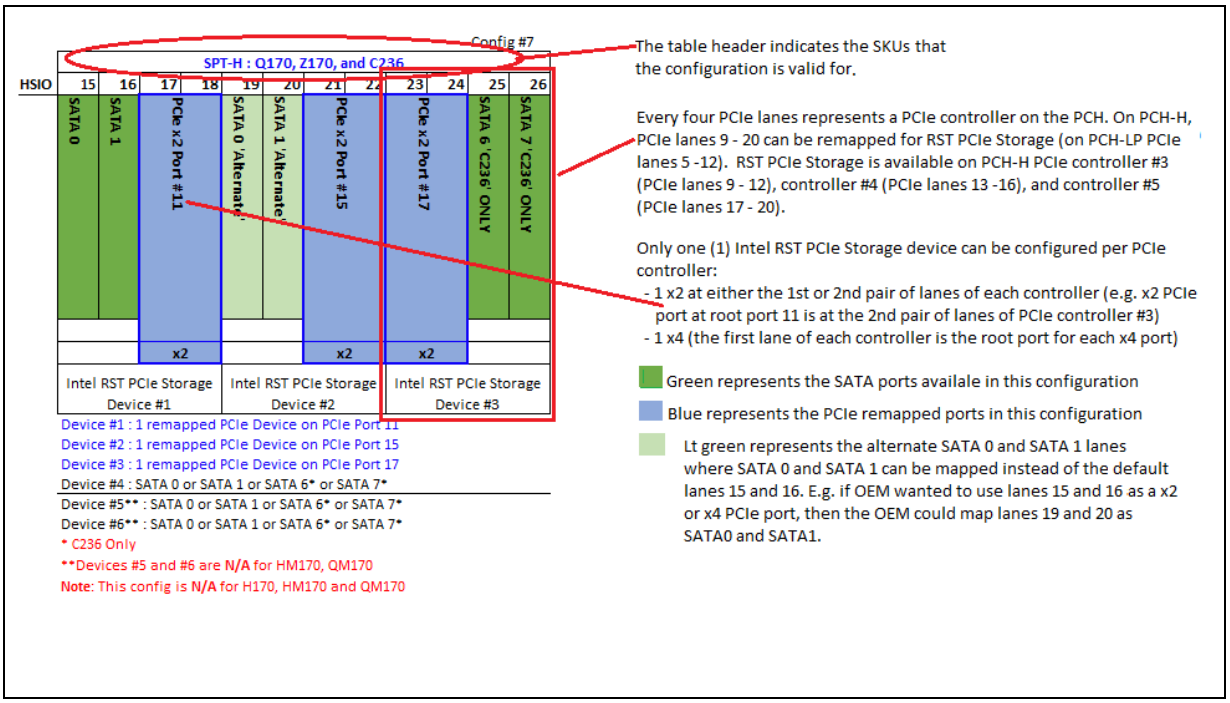


Table 33-1 PCH-H HSIO Detail (Lanes 15 -26)

	PCIe Controller #3 :| (Cycle Router #1)				PCIe Controller #4 | (Cycle Router #2)				PCIe Controller #5 | (Cycle Router #3)			
SKU	15	16	17	18	19	20	21	22	23	24	25	26
HM170	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	N / A	N / A	N / A	N / A
QM170	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	N / A	N / A	N / A	N / A
H170	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	SATA	SATA	PCIe	PCIe
Q170	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe	PCIe
Z170	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe	PCIe
C236 / CM236	PCIe / SATA	PCIe / SATA	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA	PCIe / SATA
	x4				x4				x4			
	x2		x2		x2		x2		x2		x2	
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2				Intel RST PCIe Storage Device #3			
	<----------H170**, Q170, Z170, C236---------->											
	<----------HM170, QM170---------->											
	** H170 Supports Only a Maximum of 2 Remapped PCIe Devices (Controllers) of the 3 Available											

Table 33-2 PCH-LP HSIO Detail (Lanes 9 -16)

PCH-LP HSIO Detail (HSIO Lanes 9 -16) (PCIe #5 -#12)

	PCIe Controller #2				PCIe Controller #3			
SKU	HSIO 9	10	11	12	13	14	15	16
Prem-U	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe	PCIE	PCIe / SATA	PCIe / SATA
Prem-Y	PCIe	PCIe	PCIe / SATA	PCIe / SATA	PCIe	PCIE	N/A	N/A
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			
	<----------U---------->							
	<----------Y---------->							

Remapping Rules

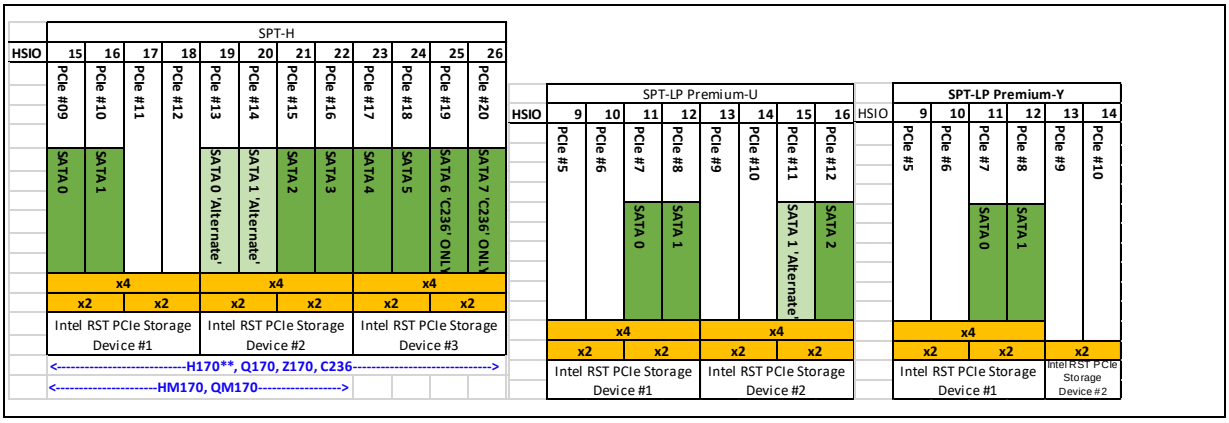
NOTE: *When remapping is enabled, the HSIO lanes that are associated with the Cycle Router that is enabled for remapping are dedicated to be used only for Intel RST PCIe Storage (PCIE NVME/AHCI or SATA devices controlled by RST). Those four lanes associated with the remapped Cycle Router cannot be used for any other purpose.*

- For a HSIO lane to be available for RST PCIe Storage remapping it must have the 'PCIe' label for that SKU in Table 33-1 and 33-2 in this section (e.g., lanes 23 and 24 cannot be remapped for RST on H170 SKU as they are labled 'SATA' only). The 'SATA' label indicates that the lane can only be mapped as SATA 3.0. The Lanes labeled 'PCIE/SATA" are flexible and can be mapped as PCIe, SATA, or remapped as RST PCIe Storage.
  -Reference: Document #546717: "Skylake H Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2" : Table 1-5
  -Reference: Document #545659: "Skylake Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2" : Table 1-3
- Only 1 Intel® RST PCIe Storage Device (x2 or x4) can be remapped per PCH PCIe controller; every 4 HSIO lanes is considered a PCIe controller (e.g. , lanes 23 - 26 are considered a controller on PCH-H; they are PCIe Controller #5 on PCH-H)
  -Reference: Document #546717: "Skylake H Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2" : Table 25-3
  -Reference: Document #545659: "Skylake Platform Controller Hub (SKL PCH) External Design Specification Volume 1 of 2" : Table 25-3
  - Controller #3 on H170 SKU can support a x4$^{Root\ 9}$, or a x2$^{Root\ 9}$, or a x2$^{Root\ 11}$

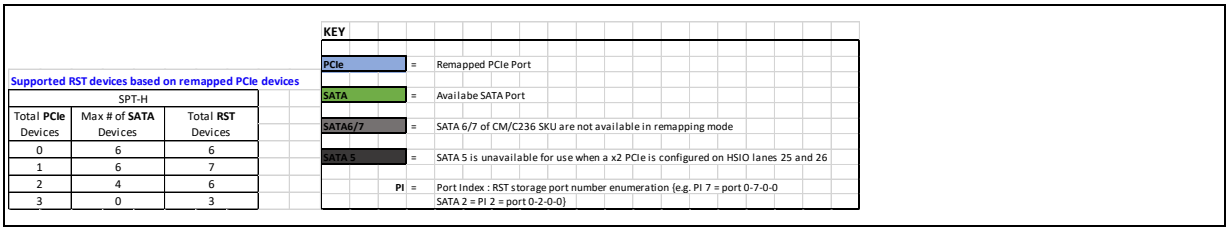
- Controller #4 on H170 SKU can support a x4[Root 13], or a x2[Root 13], or a x2[Root 15]
- Controller #5 on H170 SKU can support a single x2[Root 19] device using HSIO lanes 25 and 26 (lanes 23 and 24 cannot be remapped to the controller; they are SATA only)

Note: ***Intel RST does not support x1 PCIe lane configurations***

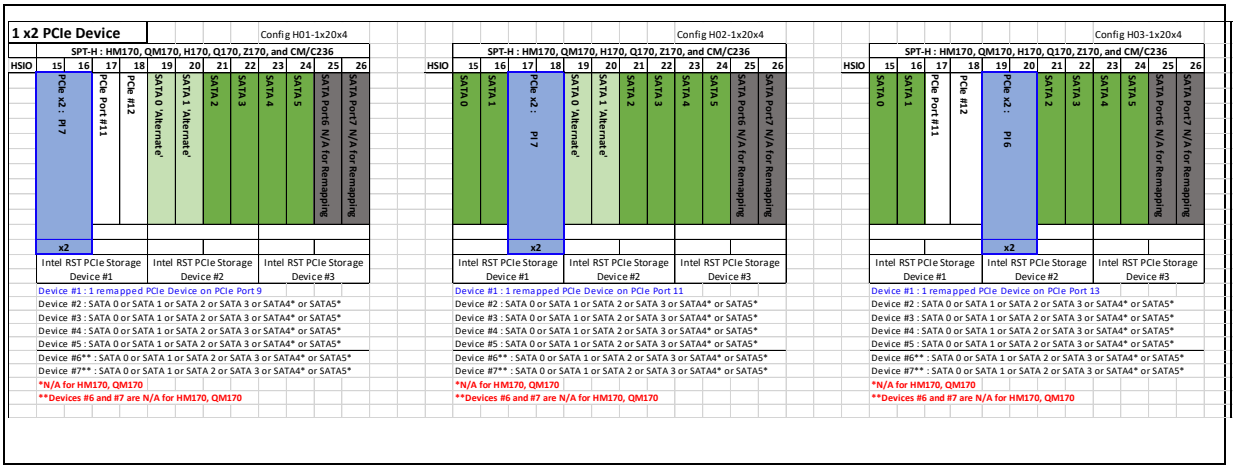
- Once a HSIO lane has been remapped for Intel RST PCIe Storage, if any SATA port is assigned to that lane, then that SATA port is no longer available for use
  - (E.g. if a x4 has been remapped on PCH-H lanes 19 and 22, then SATA 2 and 3 of the PCH are no longer available to be used)

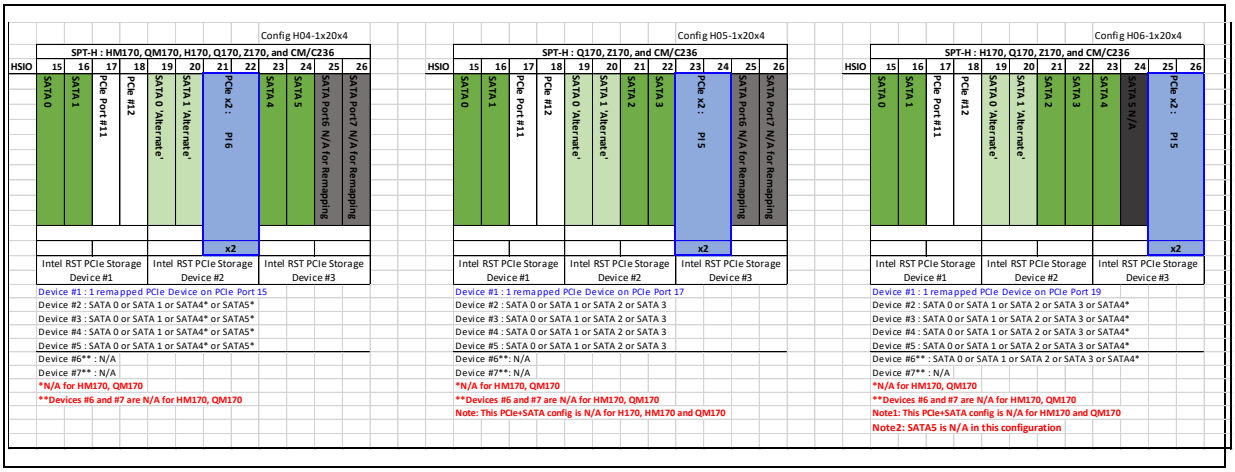


D.1.3 PCH-H Remapping Configurations

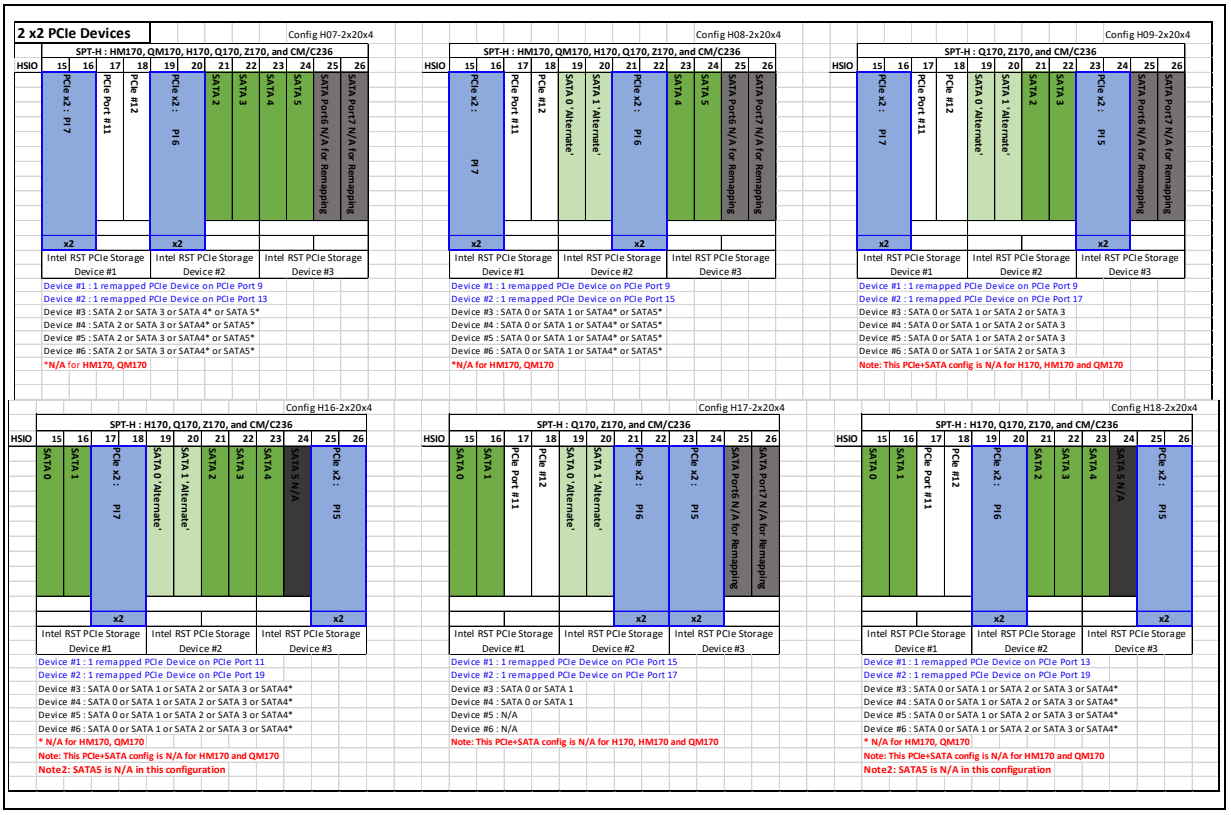


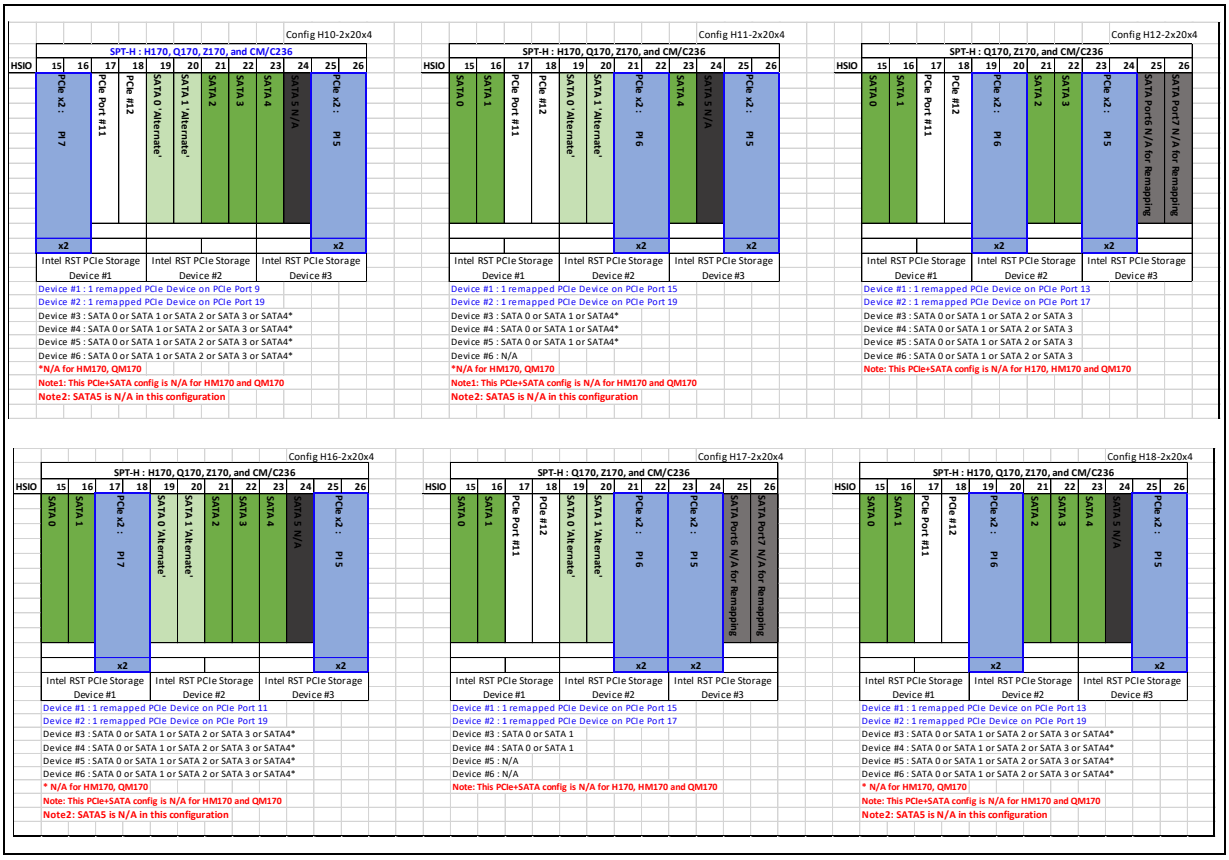
D.1.4 Configurations With 1 x2 PCIe Port Remapped



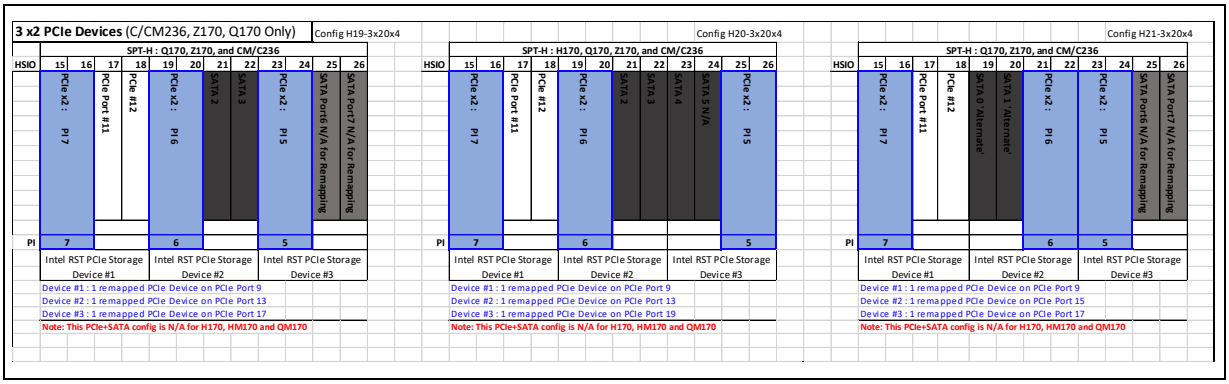


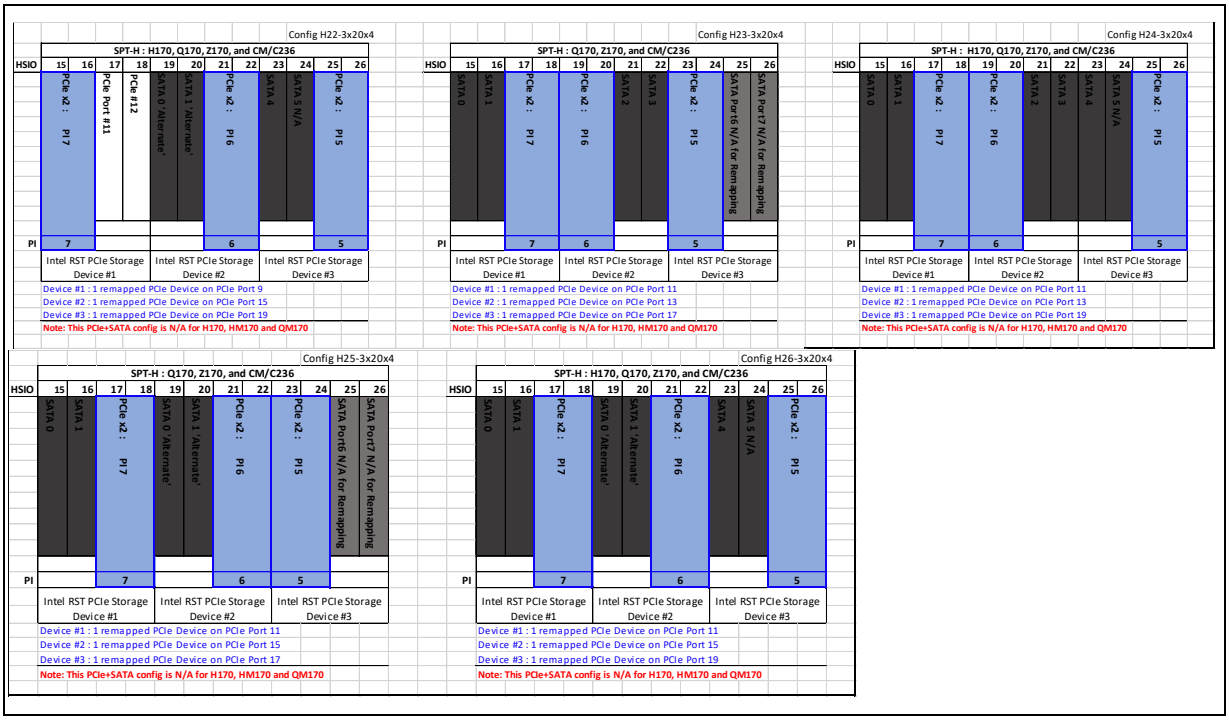
D.1.5 Configurations With 2 x2 PCIe Ports Remapped



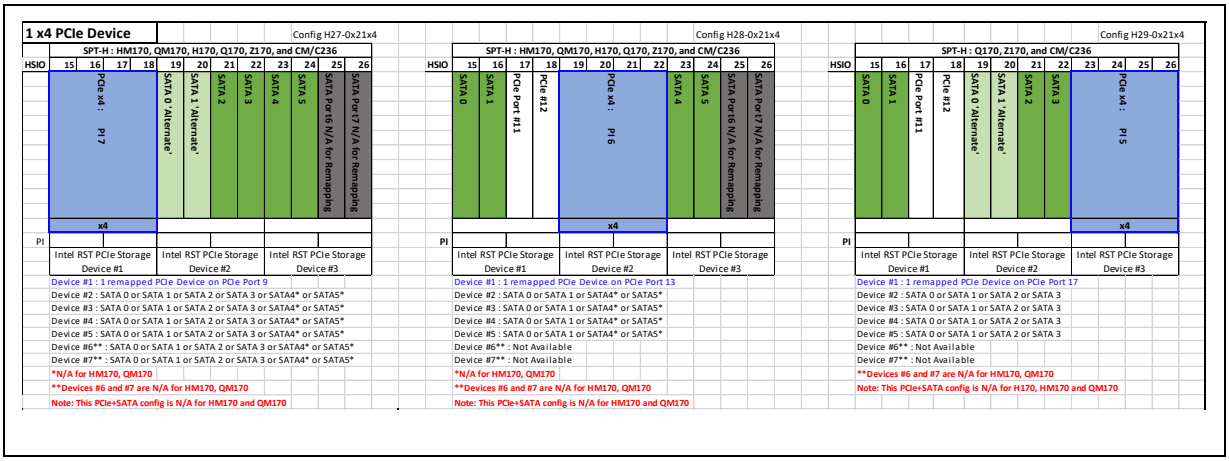


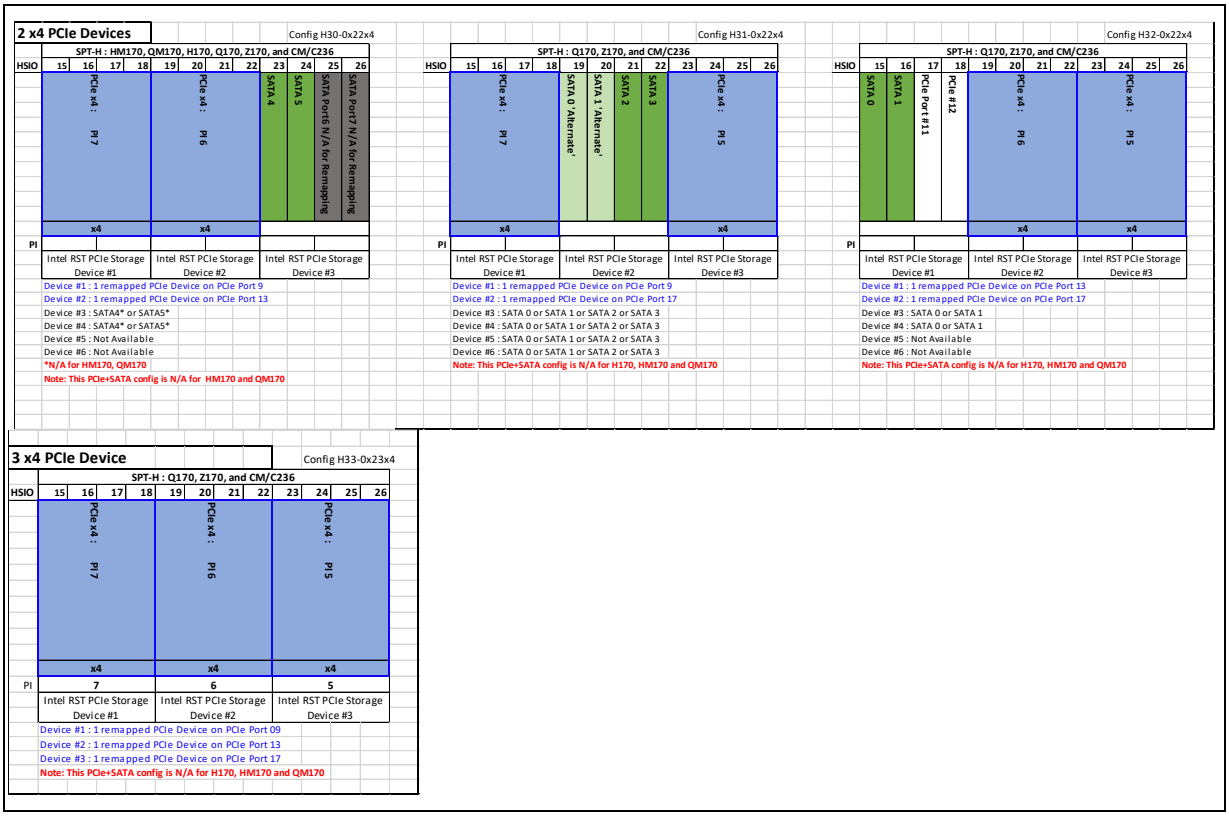
D.1.6 Configurations With 3 x2 PCIe Ports Remapped



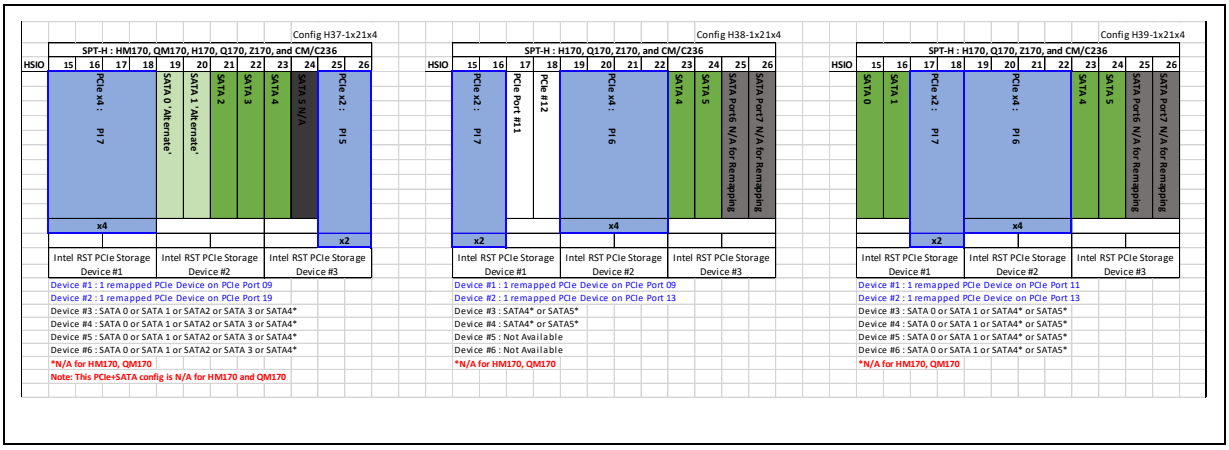


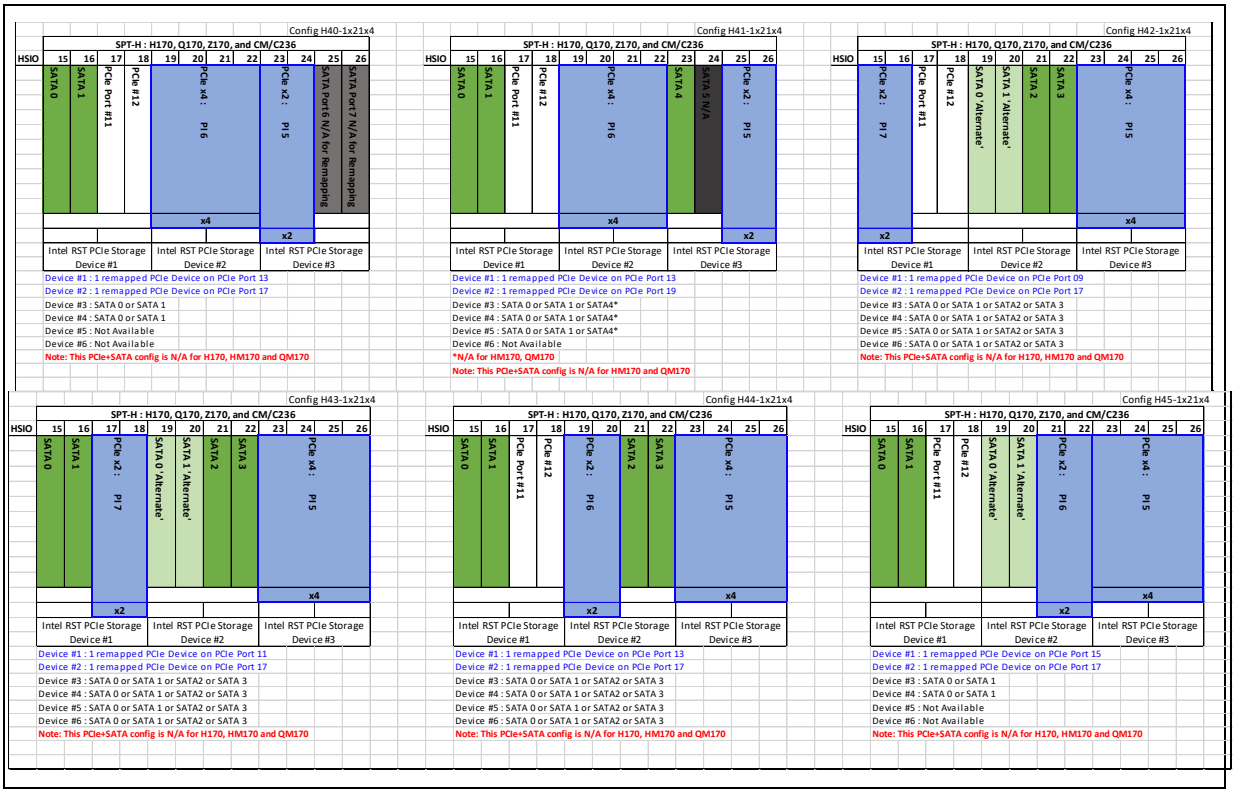
D.1.7 Configurations with 1, 2, and 3 x4 PCIe Ports Remapped



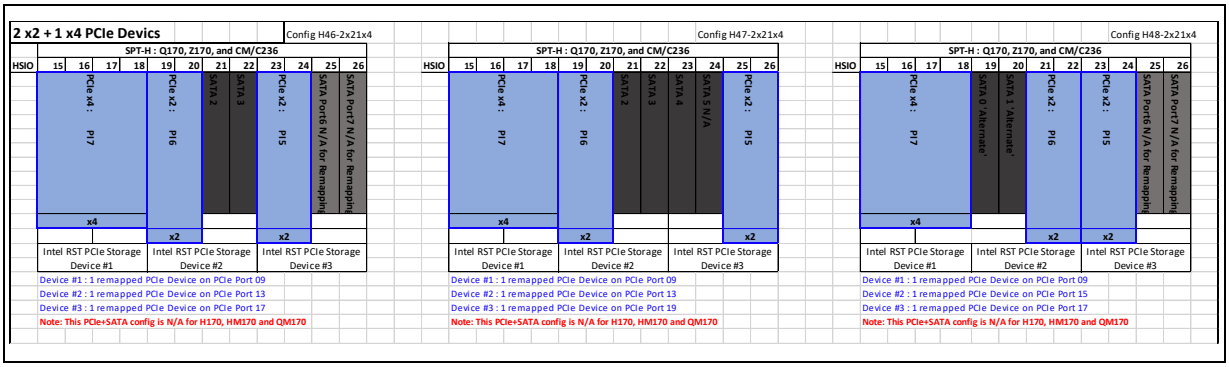


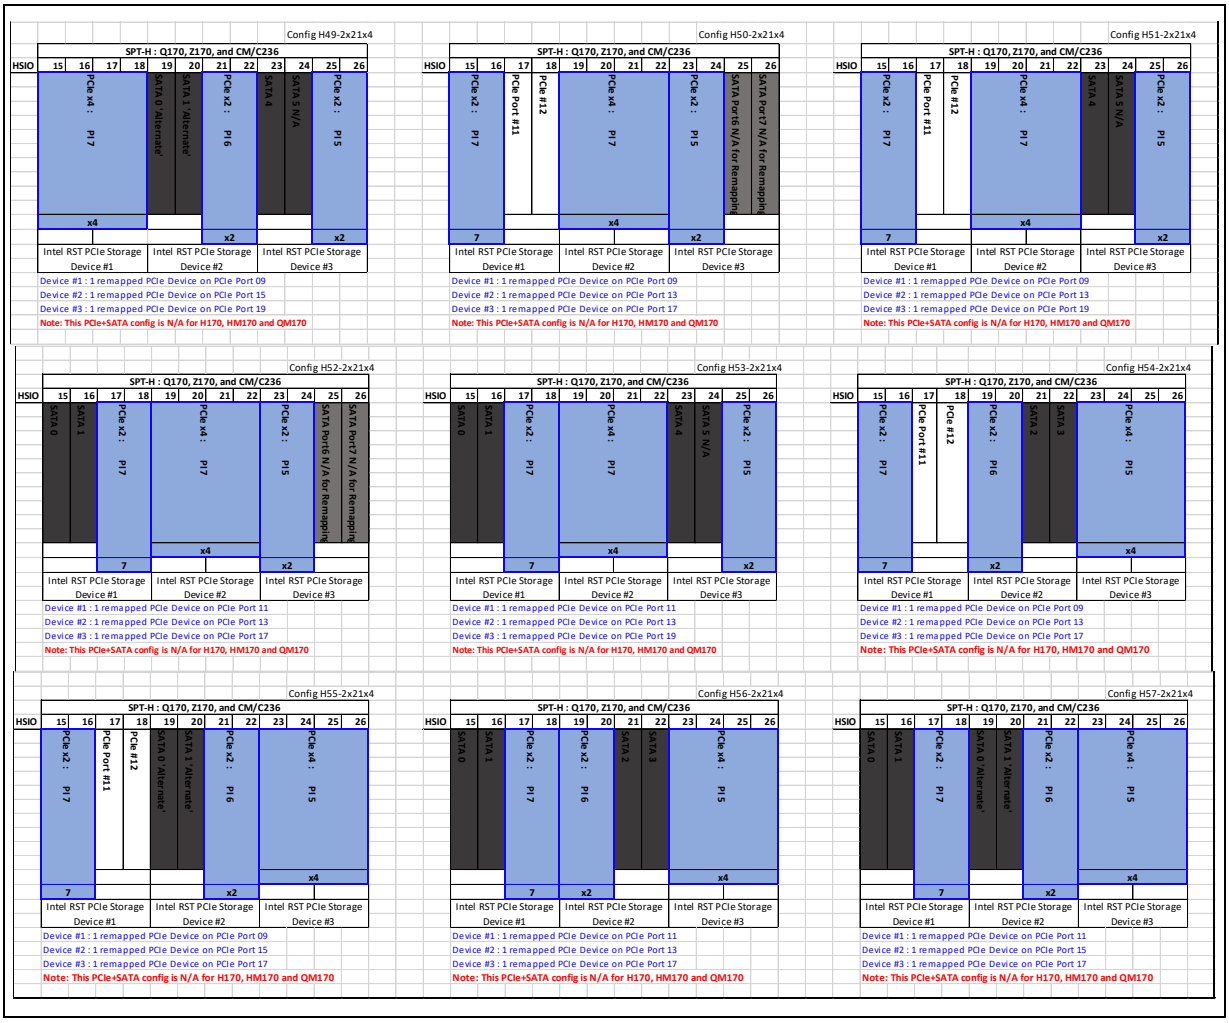
D.1.8 Configurations with (1 x2 + 1 x4) PCIe Ports Remapped



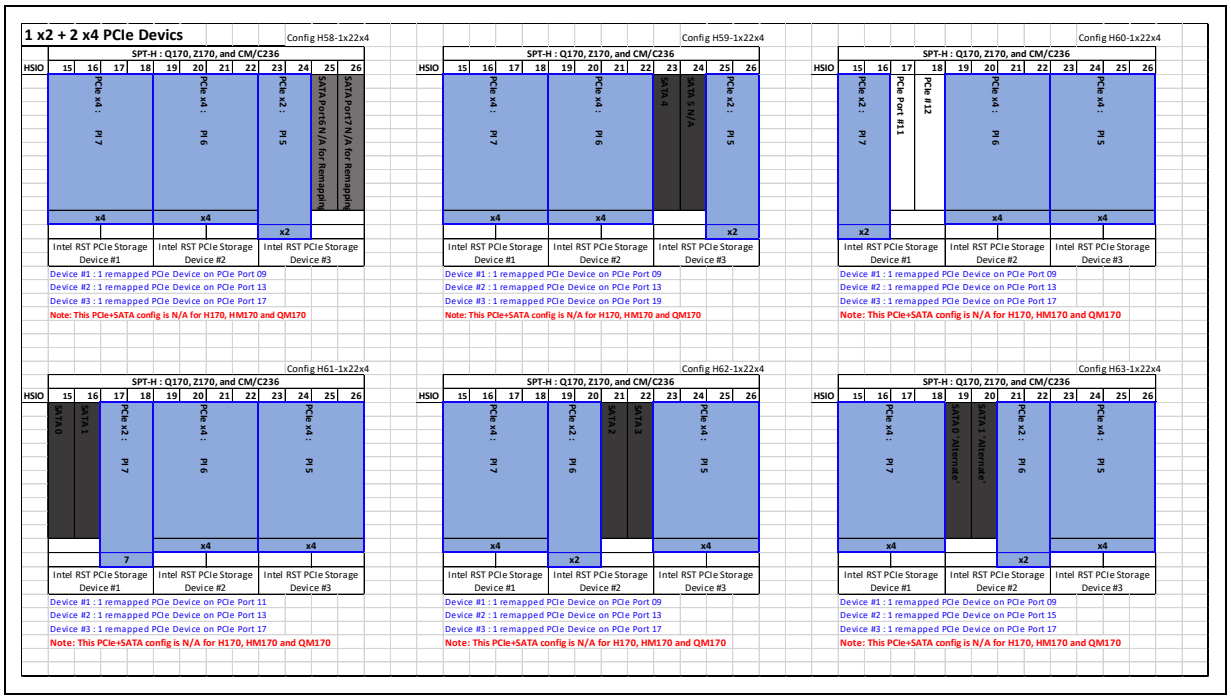


D.1.9 Configurations with (2 x2 + 1 x4) PCIe Ports Remapped




Config H49-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
PCIe x4 : PI7
SATA 0 'Alternate'
SATA 1 'Alternate'
PCIe x2 : PI6
SATA 4
SATA 5 N/A
PCIe x2 : PI5
x4
x2
x2
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 15
Device #3 : 1 remapped PCIe Device on PCIe Port 19
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H50-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
PCIe x2 : PI7
PCIe Port #11
PCIe #12
PCIe x4 : PI7
PCIe x2 : PI5
SATA Port6 N/A for Remapping
SATA Port7 N/A for Remapping
7
x4
x2
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H51-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
PCIe x2 : PI7
PCIe Port #11
PCIe #12
PCIe x4 : PI7
SATA 4
SATA 5 N/A
PCIe x2 : PI5
7
x4
x2
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 19
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H52-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
SATA 0
SATA 1
PCIe x2 : PI7
PCIe x4 : PI7
PCIe x2 : PI5
SATA Port6 N/A for Remapping
SATA Port7 N/A for Remapping
7
x4
x2
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 11
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H53-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
SATA 0
SATA 1
PCIe x2 : PI7
PCIe x4 : PI7
SATA 4
SATA 5 N/A
PCIe x2 : PI5
7
x4
x2
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 11
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 19
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H54-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
PCIe x2 : PI7
PCIe Port #11
PCIe #12
PCIe x2 : PI6
SATA 2
SATA 3
PCIe x4 : PI5
7
x2
x4
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H55-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
PCIe x2 : PI7
PCIe Port #11
PCIe #12
SATA 0 'Alternate'
SATA 1 'Alternate'
PCIe x2 : PI6
PCIe x4 : PI5
7
x2
x4
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 15
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H56-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
SATA 0
SATA 1
PCIe x2 : PI7
PCIe x2 : PI6
SATA 2
SATA 3
PCIe x4 : PI5
7
x2
x4
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 11
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170
Config H57-2x21x4
SPT-H : Q170, Z170, and CM/C236
HSIO 15 16 17 18 19 20 21 22 23 24 25 26
SATA 0
SATA 1
PCIe x2 : PI7
SATA 0 'Alternate'
SATA 1 'Alternate'
PCIe x2 : PI6
PCIe x4 : PI5
7
x2
x4
Intel RST PCIe Storage Device #1
Intel RST PCIe Storage Device #2
Intel RST PCIe Storage Device #3
Device #1 : 1 remapped PCIe Device on PCIe Port 11
Device #2 : 1 remapped PCIe Device on PCIe Port 15
Device #3 : 1 remapped PCIe Device on PCIe Port 17
Note: This PCIe+SATA config is N/A for H170, HM170 and QM170

D.1.10 Configurations with (1 x2 + 2 x4) PCIE Ports Remapped



D.1.11 PCH-LP Premium-U Remapping Configurations

Supported RST devices based on remapped PCIe devices

SPT-LP Premium-U		
Total PCIe	Max # of SATA	Total RST Devices
0	3	3
1	3	4
2	2	4

KEY

PCIe	=	Remapped PCIe Port
SATA	=	Availabe SATA Port
PI	=	Port Index : RST storage port number enumeration {e.g. PI 7 = port 0-7-0-0 SATA 2 = PI 2 = port 0-2-0-0}

D.1.12 Configurations with 1 x2 PCIe Port Remapped

1 x2 PCIe Device

Config U01-1x20x4

HSIO	9	10	11	12	13	14	15	16
SPT-LP Premium-U	PCIe x2 Port #5		SATA 0	SATA 1	PCIe #9	PCIe #10	SATA 1 'Alternate'	SATA 2
	x2							
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : SATA 0 or SATA 1 or SATA 2
Device #3 : SATA 0 or SATA 1 or SATA 2
Device #4 : SATA 0 or SATA 1 or SATA 2

Config U02-1x20x4

HSIO	9	10	11	12	13	14	15	16
SPT-LP Premium-U	PCIe #5	PCIe #6	PCIe x2 Port #7		PCIe #9	PCIe #10	SATA 1 'Alternate'	SATA 2
			x2					
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 7
Device #2 : SATA 1 or SATA 2
Device #3 : SATA 1 or SATA 2
Device #4 : Not Available

Config U03-1x20x4

HSIO	9	10	11	12	13	14	15	16
SPT-LP Premium-U	PCIe #5	PCIe #6	SATA 0	SATA 1	PCIe x2 Port #9		SATA 1 'Alternate'	SATA 2
					x2			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 9
Device #2 : SATA 0 or SATA 1 or SATA 2
Device #3 : SATA 0 or SATA 1 or SATA 2
Device #4 : SATA 0 or SATA 1 or SATA 2

Config U04-1x20x4

HSIO	9	10	11	12	13	14	15	16
SPT-LP Premium-U	PCIe #5	PCIe #6	SATA 0	SATA 1	PCIe #9	PCIe #10	PCIe x2 Port #11	
							x2	
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 11
Device #2 : SATA 0 or SATA 1
Device #3 : SATA 0 or SATA 1
Device #4 : Not Available

D.1.13 Configurations With 2 x2 PCIe Ports Remapped

2 x2 PCIe Devices

Config U05-2x20x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe x2 Port #5		SATA 0	SATA 1	PCIe x2 Port #9		SATA 1 'Alternate'	SATA 2
	x2				x2			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : 1 remapped PCIe Device on PCIe Port 9
Device #3 : SATA 0 or SATA 1 or SATA 2
Device #4 : SATA 0 or SATA 1 or SATA 2

Config U06-2x20x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe x2 Port #5		SATA 0	SATA 1	PCIe #9	PCIe #10	PCIe x2 Port #11	
	x2						x2	
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : 1 remapped PCIe Device on PCIe Port 11
Device #3 : SATA 0 or SATA 1
Device #4 : SATA 0 or SATA 1

Config U07-2x20x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe #5	PCIe #6	PCIe x2 Port #7		PCIe x2 Port #9		SATA 1 'Alternate'	SATA 2
			x2		x2			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 7
Device #2 : 1 remapped PCIe Device on PCIe Port 9
Device #3 : SATA 1 or SATA 2
Device #4 : SATA 1 or SATA 2

Config U08-2x20x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe #5	PCIe #6	PCIe x2 Port #7		PCIe #9	PCIe #10	PCIe #11	
			x2				x2	
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 7
Device #2 : 1 remapped PCIe Device on PCIe Port 11
Device #3 : Not Available

D.1.14 Configurations with 1 and 2 x4 PCIe Ports Remapped

1 x4 PCIe Device

Config U09-0x21x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe x4 Port #5				PCIe #9	PCIe #10	SATA 1 'Alternate'	SATA 2
	x4							
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : SATA 1 or SATA 2
Device #3 : SATA 1 or SATA 2
Device #4 : Not Available

Config U10-0x21x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe #5	PCIe #6	SATA 0	SATA 1	PCIe x4 Port #9			
					x4			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 9
Device #2 : SATA 0 or SATA 1
Device #3 : SATA 0 or SATA 1
Device #4 : Not Available

2 x4 PCIe Devices

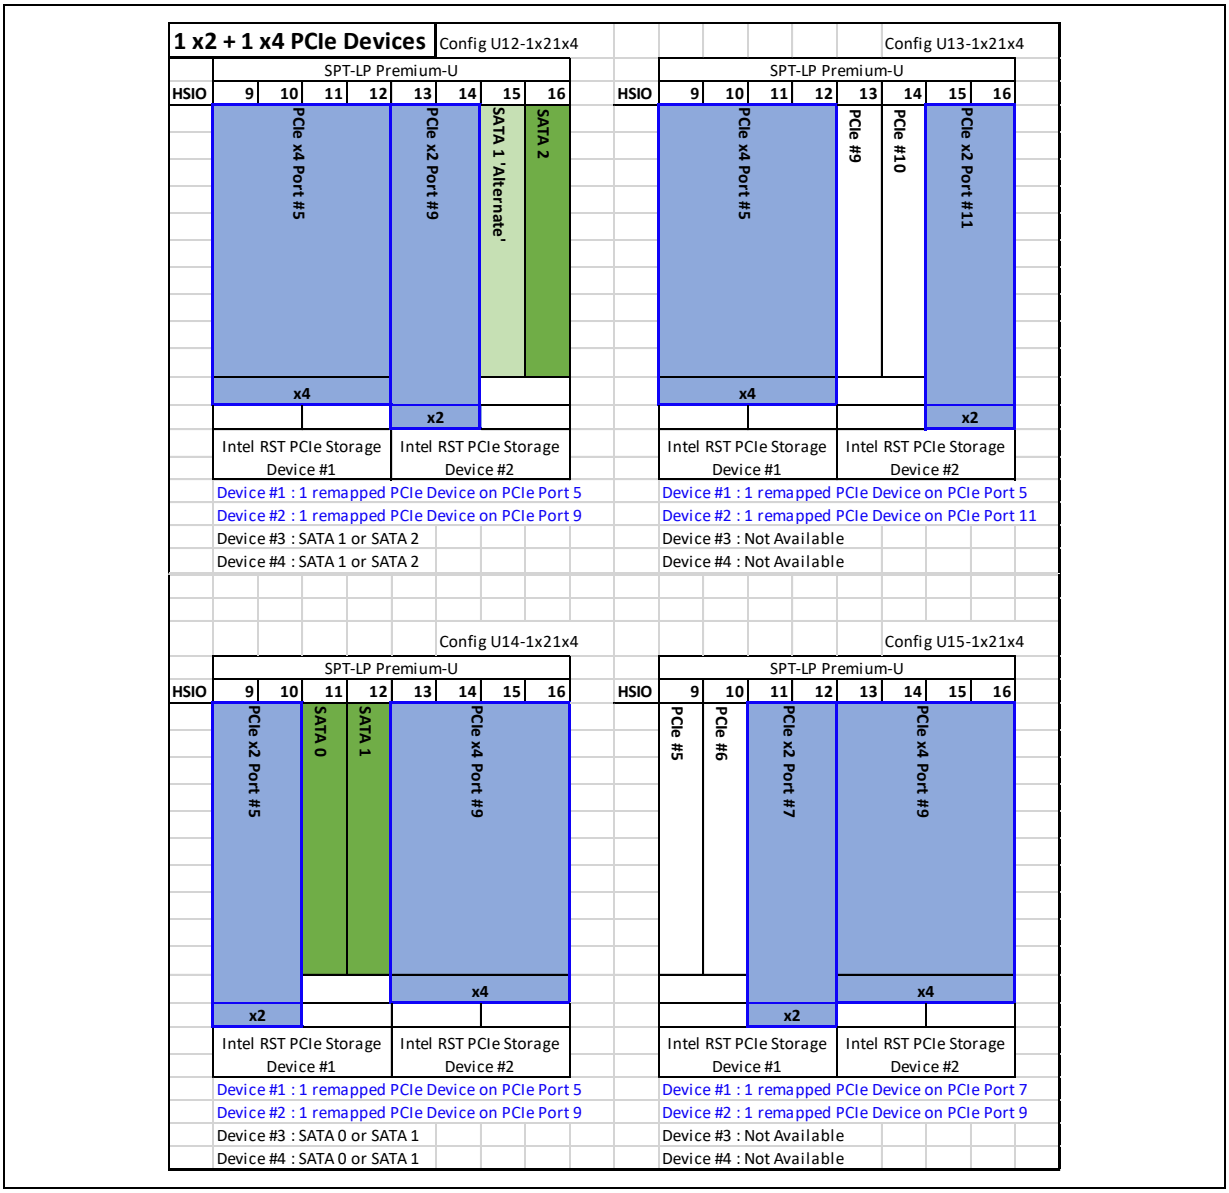
Config U11-0x22x4

SPT-LP Premium-U

HSIO	9	10	11	12	13	14	15	16
	PCIe x4 Port #5				PCIe x4 Port #9			
	x4				x4			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2			

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : 1 remapped PCIe Device on PCIe Port 9
Device #3 : Not Available
Device #4 : Not Available

D.1.15 Configurations with (1 x2 + 1 x4) PCIe Ports Remapped



D.1.16 PCH-LP Premium-Y Remapping Configurations

Supported RST devices based on remapped PCIe devices

SPT-Y		
Total PCIe	Max # of SATA	Total RST Devices
0	2	2
1	2	3
2	2	4

KEY

PCIe	=	Remapped PCIe Port
SATA	=	Availabe SATA Port
PI	=	Port Index : RST storage port number enumeration {e.g. PI 6 = port 0-6-0-0 SATA 2 = PI 2 = port 0-2-0-0}

D.1.17 Configurations With 1 x2 PCIe Port Remapped

1 x2 PCIe Device Config Y01-1x20x4

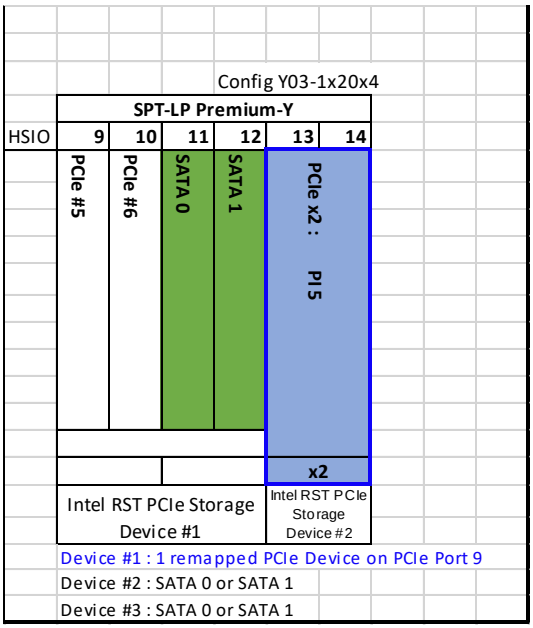
SPT-LP Premium-Y						
HSIO	9	10	11	12	13	14
	PCIe x2 : PI 6		SATA 0	SATA 1	PCIe #9	PCIe #10
	x2					
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2	

Device #1 : 1 remapped PCIe Device on PCIe Port 5
Device #2 : SATA 0 or SATA 1
Device #3 : SATA 0 or SATA 1

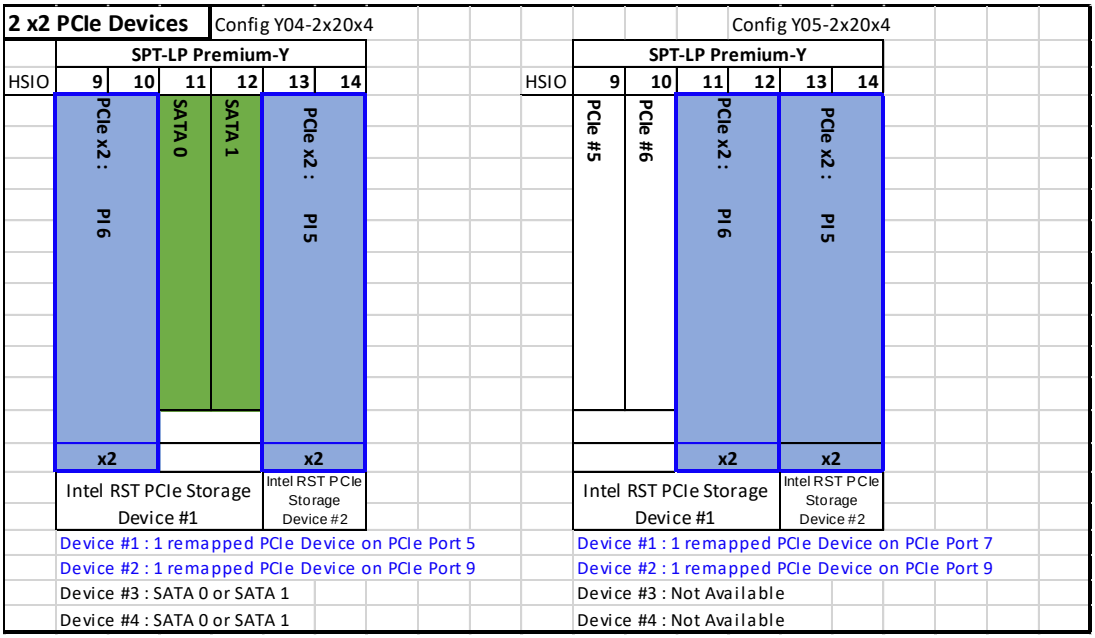
Config Y02-1x20x4

SPT-LP Premium-Y						
HSIO	9	10	11	12	13	14
	PCIe #5	PCIe #6	PCIe x2 : PI 6		PCIe #9	PCIe #10
			x2			
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2	

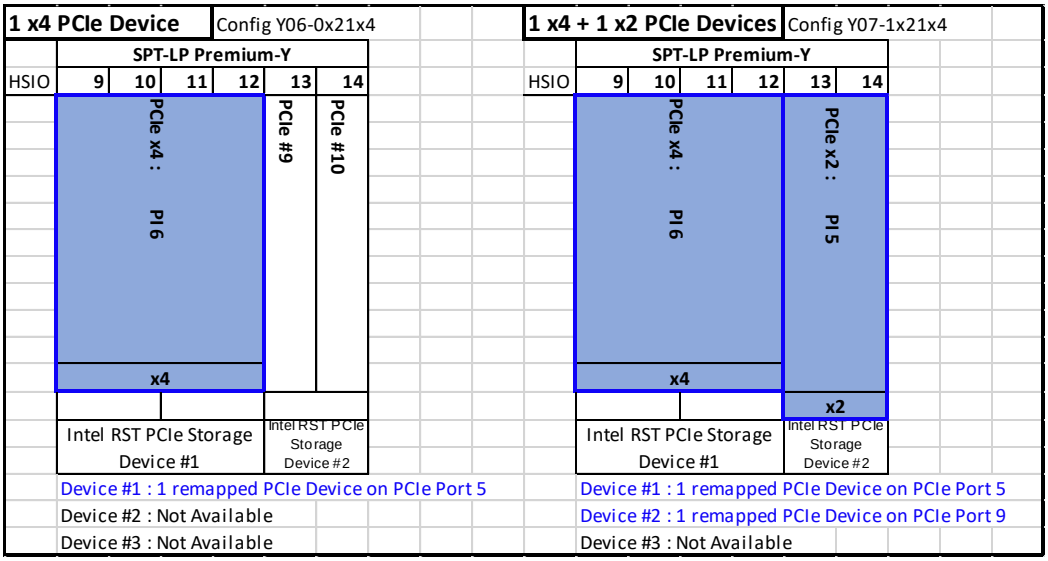
Device #1 : 1 remapped PCIe Device on PCIe Port 7
Device #2 : Not Available
Device #3 : Not Available



D.1.18 Configurations With 2 x2 PCIe Ports Remapped



D.1.19 Configurations With x4 PCIe Ports Remapped





D.2 Examples of Configurations to Meet Design Specifications

D.2.1 Example #1: SPT-H HM170 SKU With 1x2 + 1x4 + 1 SATA

Customer Design Requirement (Example of SKU Dependency):

Customer wishes to design an ***HM170*** with the following specs for Intel RST support:

- **1 x2** PCIe Storage Device
- **1 x4** PCIe Storage Device
- **1** SATA 3.0 Storage Device

Intel RST support limitations for HM170 platform

- 4 SATA 3.0 storage devices
- 2 Intel RST PCIe storage devices
- Only HSIO lanes 15 - 22 are available for remapping on the HM170 SKU
- **6 total (SATA 3.0 + RST PCIe) storage devices combined**

There are four possible configurations for 1 x4 + 1 x2 Intel® RST PCIe Storage Devices on the ***SPT-H HM170*** SKU

1 x4 + 1 x2 PCIe Devices

Config H34-1x21x4

SPT-H : HM170

HSIO	15	16	17	18	19	20	21	22	23	24	25	26
	PCIe x4 : PI7				PCIe x2 : PI6		SATA 2	SATA 3	SATA 4	SATA 5	SATA Port6 N/A for Remapping	SATA Port7 N/A for Remapping
	x4				x2							
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2				Intel RST PCIe Storage Device #3			

Device #1 : 1 remapped PCIe Device on PCIe Port 09
Device #2 : 1 remapped PCIe Device on PCIe Port 13
Device #3 : SATA 2 or SATA 3
Device #4 : Available, not required
Device #5 : Not Available
Device #6 : Not Available

Config H35-1x21x4

SPT-H : HM170

HSIO	15	16	17	18	19	20	21	22	23	24	25	26
	PCIe x4 : PI7				SATA 0 'Alternate'	SATA 1 'Alternate'	PCIe x2 : PI6		SATA 4	SATA 5	SATA Port6 N/A for Remapping	SATA Port7 N/A for Remapping
	x4						x2					
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2				Intel RST PCIe Storage Device #3			

Device #1 : 1 remapped x4 PCIe Device on PCIe Port 9
Device #2 : 1 remapped x2 PCIe Device on PCIe Port 15
Device #3 : SATA 0 or SATA 1
Device #4 : Available, not required
Device #5 : Not Available
Device #6 : Not Available

Config H38-1x21x4

SPT-H : HM170

HSIO	15	16	17	18	19	20	21	22	23	24	25	26
	PCIe x2 : PI7		PCIe Port #11	PCIe Port #12	PCIe x4 : PI6				SATA 4	SATA 5	SATA Port6 N/A for Remapping	SATA Port7 N/A for Remapping
	x2				x4							
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2				Intel RST PCIe Storage Device #3			

Device #1 : 1 remapped x2 PCIe Device on PCIe Port 09
Device #2 : 1 remapped x4 PCIe Device on PCIe Port 13
Device #3 : Not Available
Device #4 : Not Available
Device #5 : Not Available
Device #6 : Not Available

Config H39-1x21x4

SPT-H : HM170

HSIO	15	16	17	18	19	20	21	22	23	24	25	26
	SATA 0	SATA 1	PCIe x2 : PI7		PCIe x4 : PI6				SATA 4	SATA 5	SATA Port6 N/A for Remapping	SATA Port7 N/A for Remapping
			x2		x4							
	Intel RST PCIe Storage Device #1				Intel RST PCIe Storage Device #2				Intel RST PCIe Storage Device #3			

Device #1 : 1 remapped x2 PCIe Device on PCIe Port 11
Device #2 : 1 remapped x4 PCIe Device on PCIe Port 13
Device #3 : SATA 0 or SATA 1
Device #4 : Available, not required
Device #5 : Not Available
Device #6 : Not Available

1) **Config H38-1x21x4 does not meet the requirement on HM170 SKU since there are no SATA ports available in this PCIe + SATA configuration**

2) **Conclusion: The SKL SPT-H HM170 SKU with RST remapped PCIe Storage has three possible configurations (Configs #'s H34-1x21x4, H35-1x21x4, and H39-1x21x4) that will meet this customer design requirement.**

§ §