

Intel® Rapid Storage Technology enterprise (RSTe)

OEM Technical Guide

For the Intel® Rapid Storage Technology enterprise

Revision 1.1

Intel Confidential



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL® PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER, AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT. Intel products are not intended for use in medical, lifesaving, or life sustaining applications.

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

Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined." Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them.

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

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

Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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

Copyright © 2015, Intel Corporation. All rights reserved.



Contents

1 About This Document 10
1.1 Purpose and Scope of this Document 10
2 Intel® Rapid Storage Technology enterprise (RSTe) 11
2.1 Overview 11
2.1.1 Product Release Numbering Scheme 11
2.1.2 RAID Levels 12
2.1.3 Typical Usage Model for RAID Levels 14
2.1.4 Supported Platforms for This Release 14
2.1.5 Supported Chipset SKU Overview 15
2.1.6 Supported Operating Systems for This Release 16
3 Intel® RSTe Componenets 17
3.1 Intel® RSTe Pre-OS Images 17
3.1.1 Intel® RSTe RAID Legacy Option ROMs 17
3.1.2 Intel® RSTe RAID Unified Extensible Firmware Interface (UEFI) Drivers 17
3.2 Intel® RSTe OS Runtime Drivers 18
3.2.1 Drivers supporting the Intel® C220/C230 series chipset 18
3.2.2 Drivers Supporting the Intel® C610 series chipset 18
3.3 Intel® RSTe RAID Tools and Utility 19
3.3.1 Intel® RSTe Command Line Interface (CLI) Utility 19
3.3.2 Intel® RSTe Pre-OS Tools 19
3.4 RSTCLI (32/64 bit) Windows Utilities 39
4 Intel® RSTe Pre-OS Usage 49
4.1 Intel® RSTe Pre-OS Compoments 49
4.1.1 UEFI System BIOS and the Intel® RSTe UEFI/RAID Package 49
4.1.2 Legacy OROM System BIOS Requirements and the Intel® RSTe Legacy OROM Package 55
4.1.3 Choosing Boot Mode (Legacy Option ROM vs UEFI) 58
5 New in Release Version 4.3 59
5.1 General Overview 59
5.1.1 Windows 10 OS Upgrade Recommendations 59
5.1.2 RSTe Pre-OS RAID Volume Rebuild Drive Identification 59

5.1.3 SCSI Pass-Thru Support 60
5.1.4 Device Information Display in UEFI 61
5.1.5 Intel® RSTe PCH/SATA SGPIO LED Management 61
5.1.6 Common Storage Management Interface Support (CSMI) 61
5.1.7 Disk Coercion 62
5.1.8 Mechanical Ejection Support 62
5.1.9 Flush Unit Access (FUA) 63
5.1.10 Intel RSTe NVMe pass-through IOCTL support 63
5.1.11 Intel RSTe NVMe Compliance 63
5.1.12 Intel RSTe RAID Volume Support with Multiple Controllers 63
5.1.13 NVMe Admin Security Commands Support 64

6 Intel Rapid Storage Technology enterprisefor PCIe NVMe Storage Devices 65
6.1 OEM System BIOS Requirements 65
6.2 General Requirements 65
6.3 Feature Limitations 65
6.4 PCIe NVMe Device Usage Model 66
6.5 Intel® RSTe UEFI Compliance Utitlity for PCIe Storage 66

7 Creating a RAID Volume 67
7.1 Using the Intel® Rapid Storage Technology enterprise GUI 67
7.2 Using the Intel® Rapid Storage Technology enterprise Legacy Option ROM User Interface 73
7.3 Using the Intel® Rapid Storage Technology enterprise UEFI User Interface 77
7.4 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows) 83

8 Define a Hot Spare 85
8.1 Using the Windows User Interface Utility 85

9 Deleting a RAID Volume 88
9.1 Using the Intel® RSTe GUI 88
9.2 Using the Intel® RSTe Legacy Option ROM User Interface 88
9.3 Using the Intel® RSTe UEFI HII User Interface 90
9.4 Using the RAID Configuration Utilities (DOS, UEFI Shell, and Windows) 92

10 Common RAID Setup Procedures 93
10.1 Build a SATA RAID 0, 1, 5 or 10 System (OS Installation) 93
10.1.1 Using the Legacy OROM User Interface 93
10.1.2 Using the UEFI HII User Interface 94

10.2 Build a “RAID Ready” System 96
10.3 Migrate to RAID 0 or RAID 1 on an Existing “RAID Ready” System 96
10.4 Migrate an Existing Data Hard Drive to a RAID 0 or RAID 1 Volume 97
10.5 Migrating From one RAID Level to Another 98
11 Intel® RSTe NVMe LED Mangement 100
11.1 An Overview of LED Managment Support in PCIe Enclosures 100
11.2 LED Blinking Support in RSTe 4.3 100
11.2.1 LED Management Information Flow 100
11.3 LED Blinking Plug-in 101
11.3.1 The interface 101
11.3.2 ‘Name’ Property 102
11.3.3 ‘Blink’ Method 102
11.4 Plug-in Deployment 103
11.4.1 Potential Deployment Issues 104
11.5 Example Enclousre Plugin Code 104
11.6 Example of the Build Envronment 104
12 Intel® RSTe PCH/SATA SGPIO 107
12.1 Activate LED 108
12.2 FAIL LED 109
12.3 SMART LED 109
12.4 Rebuild LED 109
12.5 Migrate LED 109
13 Intel® RSTe RAID Volume Initialization 110
14 RAID Volume Verify and Repair Feature 112
14.1 Starting Verify and Repair Process 112
14.1.1 Dirty Shutdown 112
14.1.2 Manual Initiation 112
15 Pre-OS Installation of the Intel® RSTe Driver 116
15.1 Pre-OS Driver Installation Using the “Load Driver” Method 116
16 Intel® RSTe Component Versions 117
16.1 Using System Report 117
16.2 Intel® RSTe Graphical User Interface (GUI) Version 120
16.3 Using Intel® RSTe File Properties to Obtain Driver Version 121
16.4 Determining the Version of the Intel® RSTe Option ROM 122

16.4.1 Using the Intel® RSTe GUI System Report 122
16.4.2 Using the Intel® RSTe Option ROM User Interface 122
16.5 Determining the Version of the Intel® RSTe UEFI Driver 123
16.5.1 Using the Intel® RSTe GUI System Report 123
16.5.2 Using the UEFI HII 123
16.5.3 Using the UEFI Shell 124
17 Un-installation 125
17.1 Uninstalling the Intel® RSTe Software (except the RAID Driver) 125
17.2 Disabling the RAID Driver by Disabling the RAID Controller 125
18 Registry Customizations 127
18.1 E-mail Notification UI Visible Enable/Disable 127
18.1.1 Email Alerting and Notification 128
19 Intel® Rapid Storage Technology enterprise GUI 131
19.1 Introduction 131
19.1.1 Getting Started 132
19.1.2 Understanding the Application 133
19.1.3 Notification Area 138
19.2 Storage System Status 140
19.2.1 Understanding the Status 140
19.3 Creating a Volume 142
19.3.1 Volume Requirements 142
19.3.2 Creation Process 143
19.3.3 Creating Additional Volumes 147
19.4 Managing the Storage System 149
19.4.1 Managing Arrays 149
19.4.2 Managing Volumes 154
19.4.3 Managing Disks 167
19.4.4 Managing Ports 172
19.4.5 Managing ATAPI Devices 173
19.5 Preferences 174
19.6 Email Alerting and Notification 175
19.6.1 Configuration 175
19.6.2 Email Message Format 176
19.6.3 Protocol Support 177
19.6.4 Error Conditions 177
20 Using the BCFS to Differentiate Platform SKUs 178

20.1 Configuring the Platform's RAID Related Features 178
20.1.1 BCFS Bit Setting 178
20.1.2 Configuring the Behavior of the OROM UI and Banner 181
20.1.3 Configuring Intel® RSTe UI Capabilities 182
20.1.4 BIOS Control Feature Set 182
21 Testing, Certification Notes 185
21.1 Correcting Microsoft* Windows 7 (Win7) WHQL test failure 185
22 Glossary 186
23 Troubleshooting 197
23.1 Failed Volumes 197
23.2 Degraded Volumes 199
23.3 Other Volume States 202
23.4 Disk Events 203
23.5 Caching Issues 205
23.6 Software Errors 207
24 RSTe SATA Port Bitmap Implementation 209
24.1 Legacy OROM 209
24.2 UEFI Driver 210
24.3 EFI_DEVICE_PATH_PROTOCOL 210
25 Intel® C600 series chipset Support 212
25.1 Intel® C600 series chipset Supported Platforms 212
25.1.1 Intel® C600 series chipset (Romley/Romley Refresh) Platforms 212
25.1.2 Intel® C600 series chipset SKU Breakdown 212
25.1.3 Intel® C600 series chipset Upgrade ROM Options213
25.1.4 Drivers Supporting the Intel® C600 series chipset213
25.2 SCU OEM Parameters 214
25.2.1 Intel® C600 series chipset OEM Parameters 214
25.3 Structure of OEM Parameter Block 215
25.3.1 Recommended Location in SPI Flash 222
25.3.2 OEM Parameter and SPI Flash Tools 224
26 Storage System Events Detected by Monitor Service (IAStorDataMgrSvc) 231





Revision History

Revision Number	Description	Intel® RSTe Release Version	Revision Date
1.0	Initial Release of this document	4.3	June, 2015
1.1	Microsoft Windows* 2016 x64 supported as of RSTe 4.5.0.1335 and 4.5.0.2125	4.5	August, 2016



1 About This Document

1.1 Purpose and Scope of this Document

This document will assist OEMs/ODMs in evaluating, testing, configuring, and enabling RAID and AHCI functionality on platforms using the *Intel® Rapid Storage Technology enterprise* software support the platforms identified in section 2.1.4 Supported Platforms for This Release.

This document also describes installation procedures, RAID volume management such as creating, deleting, and modifying volumes, common usage models, and other RSTe features along with any special notes necessary to enable customers to develop their RAID-compatible products.



2 Intel® Rapid Storage Technology enterprise (RSTe)

Intel® Rapid Storage Technology enterprise (Intel® RSTe) provides added performance and reliability for systems equipped with serial ATA (SATA) hard drives and/or solid state disk (SSD) drives and/or NonVolitile Memory Express Solid State Drive (NVMe SSD's) to enable an optimal PC storage solution. It offers value-add features such as RAID, advanced Serial ATA* capabilities (for detailed OS support, review the Release Notes for each software release).

The RAID solution supports RAID level 0 (striping), RAID level 1 (mirroring), RAID level 5 (striping with parity) and RAID level 10 (striping and mirroring). Specific platform support is dependent upon the available SATA ports.

A configuration supporting two RAID levels can also be achieved by having two volumes in a single RAID array that use Intel® RSTe. These are called Matrix RAID Arrays. Typical for enterprise servers and workstations, Intel® RSTe RAID solution addresses the demand for high-performance or data-redundant platforms.

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. **4**.**0**.**0**.**1001**).

Number / Section	Description
AA: **Major Release Number**	This section represents the major release version of the product. Itis 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 enterpriseallows 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. Please 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.

2.1.3 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).

2.1.4 Supported Platforms for This Release

Intel® RSTe provides enhanced management capabilities and detailed status information for Serial ATA AHCI, NVMe, and RAID subsystems. Basic support for this release is on the following hardware components.



Note: Some RSTe features are limited to hardware and/or OS versions and will be documented in this guide under each feature's requirements.

2.1.4.1 Intel® C220 series chipset (Denlow/Denlow Refresh) Platforms

Intel® Server Series (Haswell) Chipset-based platform CRBs with the Intel® C220 series chipset (PCH)

- Agate Beach (Server Platform, 1 CPU)
- Flathead Creek

2.1.4.2 Intel® C610 series chipset (Grantley/Grantley Refresh) Platforms

Intel® Server Series (Haswell/Broadwell) Chipset-based platform CRBs with the Intel® C610 series chipset (PCH)

- Mayan City 2S -EP
- Yakima City 2S -EN
- Aztec City 2S -EP Workstation
- Inca City 4S -EP

2.1.4.3 Intel® C230 series chipset (Greenlow) Platforms

Intel® Xeon® processor E3-1200 v5 (Sky Lake) product family w/ Intel® C230 series chipset family

- Zumba Beach
  - Intel® C232 series chipset for Servers with 6 6Gb/s SATA 3.0 Ports
  - Intel® C236 series chipset for Servers and Workstations with 8 6Gb/s SATA 3.0 Ports

2.1.5 Supported Chipset SKU Overview

2.1.5.1 Intel® C610 series chipset SKU's

SKU Features	Grantley	Grantley Refresh



6 Gb/s SATA3 Ports	6	4	8	6
Intel® RSTe PCH/SATA Device IDs	SATA 2826	sSATA: 2827	SATA 2826	sSATA 2827
Intel RSTe NVMe Device IDs	2F9c		6F9C	
Intel® RSTe SATA RAID Support	0/1/5/10	0/1/5/10	0/1/5/10	0/1/5/10

2.1.5.2 Intel® C220/230 series chipset SKU's

SKU Features	Denlow	Greenlow
6 Gb/s SATA3 Ports	6	8
Intel® RSTe PCH/SATA Device ID	SATA 2826	SATA 2826
Intel® RSTe NVMe Device ID	N/A	A135
Intel® RSTe SATA RAID Support	0/1/5/10	0/1/5/10

2.1.6 Supported Operating Systems for This Release

- Microsoft Windows* 10#
- Microsoft Windows* 10 x64
- Microsoft Windows* 8.1#
- Microsoft Windows* 8.1 x64
- Microsoft Windows* 7
- Microsoft Windows* 7 x64-bit Edition
- Microsoft Windows* Server 2008 R2 x64#
- Microsoft Windows* Server 2012 x64-bit (All Editions) #
- Microsoft Windows* Server 2012 R2 x64-bit (All Editions)
- Microsoft Windows* Server 2016 x64

#- These OS versions pass validation testing, but are not WHQL certified.



3 Intel® RSTe Componenets

Intel® RSTe product package is comprised of several components that provide a complete platform solution. The components are:

- RSTe Pre-OS images
- RSTe OS Runtime Drivers
- RSTe Tools and Utilities

3.1 Intel® RSTe Pre-OS Images

Intel® RSTe will provide a Pre-OS package for BIOS vendors and OEM/ODMs developing their own BIOS. The Pre-OS package will contain the binary images necessary to provide RAID management within a BIOS environment to support RAID mode booting. These images are intended to be compiled into the system BIOS on the supported platforms (Intel® C600/C610/C220 and C230 series chipsets).

3.1.1 Intel® RSTe RAID Legacy Option ROMs

Intel® RSTe will provide Legacy RAID OROM images that support standard Int13 BIOS environments. These images provide a BIOS level menu-driven configuration tool for managing RSTe RAID volumes outside of an OS. There are individual images for each of the PCH controllers on the Intel® C600/C610/C220/ and C230 series chipset platforms:

- There is a single RSTe RAID Legacy Option ROM image that will support the AHCI/SATA controller for each of the mentionsed chipsets.
- For the Intel® C600 series chipset there is also an SCU RAID Legacy Option ROM image. Please reference section 25 Intel® C600 series chipset Support for additional support information.
- For the Intel® C610 series chipsets there is also a RAID Legacy Option ROM for the sSATA controller.

3.1.2 Intel® RSTe RAID Unified Extensible Firmware Interface (UEFI) Drivers

Intel® RSTe will provide UEFI drivers to support a UEFI BIOS environment. The UEFI drivers provide an HII interface to support a BIOS level menu-driven configuration tool that is accessed in the BIOS setup. The RSTe UEFI images for the Intel® C600/C610/C220/C320 series chipsets are:



- There is a single RSTe RIAD UEFI driver image to support the SATA controller for these PCHs
- For the Intel® C600 series chipset there is also an SCU RAID UEFI driver image.
- For the Intel® C610 series chipsets there is also a RAID UEFI driver image for the sSATA controller.

3.2 Intel® RSTe OS Runtime Drivers

The RSTe OS runtime drivers are based on the Windows* Storport Miniport device driver model. This provides for improved performance along with simplified functional maintainability. There are three sets of Windows* runtime drivers and they are as follows:

3.2.1 Drivers supporting the Intel® C220/C230 series chipset

There are two SATA (AHCI) Controller drivers and GUI described above also support these platforms. The first driver supports the SATA controller in AHCI mode and a second driver supports RAID mode.

Note: The system BIOS Configuration utility is used to select either AHCI or RAID modes for the SATA controller.

3.2.2 Drivers Supporting the Intel® C610 series chipset

There are two SATA (AHCI) Controller drivers and GUI described above also support this platform. The first driver supports the SATA controller in AHCI mode and a second driver supports RAID mode.

In addition, there are two additional drivers included to support the sSATA controllers. The first driver supports the sSATA controller in AHCI mode and a second driver supports RAID mode.

NOTE: The system BIOS Configuration utility is used to select either AHCI or RAID modes for the SATA and sSATA controllers.

NOTE: RAID spanning across controllers is not supported in the RSTe Windows driver.



3.3 Intel® RSTe RAID Tools and Utility

3.3.1 Intel® RSTe Command Line Interface (CLI) Utility

The Intel® RSTe package provides the OEM/ODM's factory/manufacturing environment utility that will support command line scripting. This application is designed to operate as a scripting tool and runs in either Windows* PE or a Windows* command prompt window. There are separate tools for each controller for the Intel® C600/C610/C220 and C230 series chipsets.

3.3.2 Intel® RSTe Pre-OS Tools

The Intel® RSTe package provides the OEM/ODM's factory/manufacturing environment with a set of Pre-OS tools to preconfigure and test the platform environment to help make sure the system is setup and configured as needed.

3.3.2.1 Intel® RSTe SATA/sSATA RAID Configuration

The RAID configuration tools are designed to help OEM/ODM manufacturing and testing environment. The intent of the tools allow OEM/ODM manufacturing and testing environments configure their platforms automantically using scripts.

There are two sets of Intel® RSTe SATA/sSATA RAID configuration tools. One set are DOS based utilities that runs from a DOS bootable USB memory stick and are used with the SATA and/or sSATA controllers are configured to boot using the RAID Legacy Option ROM. The other set of tools are UEFI based tools that also run from a USB memory stick from within the platform's UEFI Shell environment. These tools are used when the SATA and/or sSATA controllers are configured to boot using the RAID UEFI drivers.

3.3.2.1.1 Rcfgsata.exe/rcfgssata.exe:

Intel(R) RAID Utility for Serial ATA - v4.3.0.1018

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]

/? Displays Help Screen. Other options ignored.

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



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

COMMANDS - Only one at a time.

/C Create a volume with the specified name.

/S, /DS, /SS, & /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.

/ST Display Volume/RAID/Disk Status.

/SP Marks the selected drive(s) as spare(s). Use with /DS

/V Display version information

3.3.2.1.2 Rcfgsata.efi/rcfgssata.efi:

Intel(R) UEFI RAID Utility for Serial ATA - v4.3.0.1018

RCfgSata.efi [/?] [/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]



/? Displays Help Screen. Other options ignored.

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

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

COMMANDS - Only one at a time.

/C Create a volume with the specified name.

/S, /DS, /SS, & /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.

/ST Display Volume/RAID/Disk Status.

/SP Marks the selected drive(s) as spare(s). Use with /DS

/V Display version information

3.3.2.2 Intel® RSTe SATA/sSATA RAID Compatibility

The RAID compatibility tools are designed to support OEM/ODMs and BIOS developments/vendors. The tools help identify potential integration issues the BIOS developer/vendor may be having incorporating the RSTe PreOS images into the



system BIOS for platforms containing the Intel® C600/C610/C220 and C230 series chipsets.

There are two sets of Intel® RSTe SATA RAID compatibility tools. One set of tools are DOS based utilities that runs from a DOS bootable USB memory stick. These tools are used with the SATA and/or sSATA controllers are configured to boot using the RAID Legacy Option ROM.

The other set of tools are UEFI based tools that runs from a USB memory key from within the system's UEFI Shell environment. These tools are used with the SATA and/or sSATA controllers are configured to boot using the RAID UEFI driver images.

3.3.2.2.1 Rcmpsata.exe/rcmpssata.exe output example:

```
Intel(R) RAID Option ROM Integration Test Suite, 4.3.0.1018

Copyright(C) 2003-14 Intel Corporation.  All Rights Reserved.

RAID Option ROM Version Information

==================================

Version Table Found: D084:0000

 Version Header Revision: 2.2

 ROM Version:          4.3.0.1018

 Version Table Checksum:  Valid

 Supported RAID Levels:   RAID0, RAID1, RAID10, RAID5

 Max Disks per Array:     8

 Max Disks Supported:     8

 Max Volumes per Array:   2


 Supported Strip Sizes:

   2k 4k 8k 16k 32k 64k 128k 256k 512k 1M 2M 4M 8M 16M 32M 64M

    0  1  1  1  1  1   1   0   0  0  0  0  0  0  0  0  0
```



OROM Supported Attributes:

RAID0

RAID1

RAID10

RAID5

>=2TB DISK

BAD BLOCK MGMT

2TB VOLUMES

CHECKSUM VERIFY

OROM Supported Capabilities:

ESATA

HDD PASSWORD

DISK COERCION

OROM Supported Driver Features:

HDD UNLOCK

SGPIO

RSTe ENTERPRISE

ROHI

READ PATROL

OROM PCIR Header Location:D080:00A0

OROM is PCI Version 3.0



OROM PCI Version 3.0 Device List

Device: 0x2826

Pnp Header 1

(0x06): Next Header 0x0000

(0x08): PortBitMap 0x03

(0x0E): Mfr Str (0x00C0) Intel Corporation

(0x10): Pro Str (0x00D2) Intel Volume0

(0x16): BCV 0x1CE0

MULTI-SEGMENT OROM

1. Intel(R) SATA RAID Controller Presence Tests

1.1 Intel(R) Chipset Detected.....PASS

1.2 RAID Controller DevID at D31:F2.....PASS

1.3 RAID Controller Name.....C610 Series SATA Premium

1.4 RAID SubClass Code (0x04) at D31:F2.....PASS

1.5 Chipset Revision.....Rev 4

1.6 AIE Bit is CLEAR(0) at Dev31 Funct2.....PASS

1.7 Raid Platform ID.....8D06

1.8 RAID Controller DevID.....2826

1.9 CPU Brand String.....



2. PnP Installation Tests

2.1 $PnP Anchor String Present in 0xF000 Segment.....PASS

2.2 PnP Header Checksum Valid.....PASS

3. BBS Enabled BIOS Tests

3.1 PCI BIOS Present (INT 1A, Func 0xB101).....PASS

3.2 PCI BIOS Version.....3.0

3.3 PCI BIOS Version Supports OROM PCI Version.....PASS

4. AHCI Register Settings Tests

4.1 HBA Port Multiplier Support Status.....DISABLED

4.2 HBA Supports Command List Override.....PASS

4.3 HBA Supports Interlock Switch.....DISABLED

4.4 HBA Capabilities Register Locked.....PASS

4.5 AHCI Enable Bit Status.....ENABLED

4.6.1.1 Port Implemented Bit Set for Port 0.....DISABLED

4.6.1.2 PI Bit Matches PE Bit for Port 0.....PASS

4.6.2.1 Port Implemented Bit Set for Port 1.....DISABLED

4.6.2.2 PI Bit Matches PE Bit for Port 1.....PASS

4.6.3.1 Port Implemented Bit Set for Port 2.....DISABLED

4.6.3.2 PI Bit Matches PE Bit for Port 2.....PASS

4.6.4.1 Port Implemented Bit Set for Port 3.....ENABLED

4.6.4.2 PI Bit Matches PE Bit for Port 3.....PASS

4.6.4.3 Aggressive Link Power Management on Port 3.....DISABLED



4.6.4.4 Hot Plug Support on Port 3.....DISABLED

4.6.4.5 Interlock Switch Attached to Port 3.....N/A

4.6.4.6 ALPE and HPCP Cannot Both be Set on Port 3.....PASS

4.6.4.7 ISP Requires HPCP on Port 3.....PASS

4.6.4.8 Port 3 Command Register Locked.....PASS

4.6.4.9 External Port Bit Set for Port 3.....DISABLED

4.6.5.1 Port Implemented Bit Set for Port 4.....ENABLED

4.6.5.2 PI Bit Matches PE Bit for Port 4.....PASS

4.6.5.3 Aggressive Link Power Management on Port 4.....DISABLED

4.6.5.4 Hot Plug Support on Port 4.....DISABLED

4.6.5.5 Interlock Switch Attached to Port 4.....N/A

4.6.5.6 ALPE and HPCP Cannot Both be Set on Port 4.....PASS

4.6.5.7 ISP Requires HPCP on Port 4.....PASS

4.6.5.8 Port 4 Command Register Locked.....PASS

4.6.5.9 External Port Bit Set for Port 4.....DISABLED

4.6.6.1 Port Implemented Bit Set for Port 5.....DISABLED

4.6.6.2 PI Bit Matches PE Bit for Port 5.....PASS

4.6.7.1 Port Implemented Bit Set for Port 6.....DISABLED

4.6.7.2 PI Bit Matches PE Bit for Port 6.....PASS

4.6.8.1 Port Implemented Bit Set for Port 7.....DISABLED

4.6.8.2 PI Bit Matches PE Bit for Port 7.....PASS

4.7 CAP.SIS and PxCMD.ISP Agreement.....PASS

4.8 Software Feature Mask Bit.....ENABLED

Feature| OCF| BCFS|



RAID0| X| X|

RAID1| X| X|

RAID10| X| X|

RAID5| X| X|

RRT| -| X|

OROM UI/BANNER| X| X|

HDDUNLOCK| X| X|

LED LOCATE| X| X|

RRT on ESATA| X| X|

NVM| -| X|

DELAYBIT0| X| -|

DELAYBIT1| X| -|

RESERVED| X| -|

RESERVED| -| -|

RESERVED| -| -|

RESERVED| -| -|

5. Chipset RAID Controller Settings Tests

5.1 RAID I/O Space Enabled.....PASS

5.2 RAID Memory Space Enabled.....PASS

5.3 RAID Programming Interface Set to 0.....PASS

5.4 RAID Programming Interface Locked.....PASS

5.5 Primary Command BAR Programmed.....PASS

5.6 Primary Control BAR Programmed.....PASS

5.7 Secondary Command BAR Programmed.....PASS



5.8 Seconday Control BAR Programmed.....PASS

5.9 Legacy Bus Master BAR Programmed.....PASS

5.10 AHCI BAR Programmed.....PASS

5.11 MAP Bits Set to Non-Combined (00b).....PASS

5.12 SATA Mode Select Set to RAID (10b).....PASS

5.13.1.1 Port 0 Enable Bit Set.....DISABLED

5.13.2.1 Port 1 Enable Bit Set.....DISABLED

5.13.3.1 Port 2 Enable Bit Set.....DISABLED

5.13.4.1 Port 3 Enable Bit Set.....ENABLED

5.13.4.2 SATA Device Present on Port 3.....HDD

5.13.5.1 Port 4 Enable Bit Set.....ENABLED

5.13.5.2 SATA Device Present on Port 4.....HDD

5.13.6.1 Port 5 Enable Bit Set.....DISABLED

5.13.7.1 Port 6 Enable Bit Set.....DISABLED

5.13.8.1 Port 7 Enable Bit Set.....DISABLED

5.14 At least 2 SATA HDDs found.....PASS

5.15 Primary IO Decode Enabled.....N/A

5.16 Secondary IO Decode Enabled.....N/A

5.17 Primary Timing Register Programmed.....N/A

5.18 Secondary Timing Register Programmed.....N/A

6. Intel(R) RAID Option ROM Tests

6.1 Option ROM Found in BIOS Compatibility Segments.....PASS

6.2 Option ROM Build Matches Chipset.....PASS

6.3 Option ROM Checksum Valid.....PASS



6.4 Main Memory Size < 639kB.....PASS

6.5 Option ROM Caching.....ENABLED

6.6 Option ROM Runtime Size.....33.5kB

7. System BIOS Tests

7.1 BIOS Does Not Export HDD's on RAID Controller.....N/A

7.2 BIOS Updating INT 15h E820 Memory Map.....PASS

7.3 BIOS Reporting Correct Base Memory Size in 40:13.....PASS

8. Interrupt Settings Tests

8.1 RAID Controller IRQ.....10

8.2 Clear Spurious Interrupts on RAID Controller.....N/A

8.3 Performance Enhancement Semaphore at 40:8E.....DISABLED

10. INT13 Tests

DOS Drive 0x80:

10.1.1 INT13H FN41H Support.....ENABLED

Fixed disk access subset supported: TRUE

Enhanced disk drive subset supported: TRUE

10.1.2 INT13H FN48H Option ROM Support.....ENABLED

Disk 0x80 Does Support INT13 FN48H, Data Returned:

BufferSize: 0x4A

Flags: 0x6

NumPhysCylinders: 0x3C2

NumPhysHeads: 0xFF



NumPhysSectorsPerTrack: 0x3F

NumPhysSectors: 0x0EC0000

BytesPerSector: 0x200

EDDPointer: 0xFFFFFFFF

DevicePathSignature: 0xFABEDD

DevicePathLength: 0x2C

HostBusType: 'PCI '

InterfaceType: 'USB '

Interface path:

BusNumber: 0x0

SlotNumber: 0x0

FunctionNumber 0x0

ChannelNumber: 0x0

Device path:

PortNumber: 0x88

PortBitMap 60h: 0x0

DOS Drive 0x81:

10.2.1 INT13H FN41H Support.....ENABLED

Fixed disk access subset supported: TRUE

Enhanced disk drive subset supported: TRUE

10.2.2 INT13H FN48H Option ROM Support.....ENABLED

Disk 0x81 Does Support INT13 FN48H, Data Returned:

BufferSize: 0x1E

Flags: 0x1

NumPhysCylinders: 0x0



NumPhysHeads: 0x0

NumPhysSectorsPerTrack: 0x0

NumPhysSectors: 0x054F4B000

BytesPerSector: 0x200

EDDPointer: 0xFFFFFFFF

DevicePathSignature: 0xFABEDD

DevicePathLength: 0x2C

HostBusType: 'PCI '

InterfaceType: 'RAID '

Interface path:

BusNumber: 0xFF

SlotNumber: 0xFF

FunctionNumber 0xFF

ChannelNumber: 0xFF

Device path:

RaidArrayNumber: 0x0

PortBitMap 60h: 0xF

DOS Drive 0x82:

10.3.1 INT13H FN41H Support.....ENABLED

Fixed disk access subset supported: TRUE

Enhanced disk drive subset supported: TRUE

10.3.2 INT13H FN48H Option ROM Support.....ENABLED

Disk 0x82 Does Support INT13 FN48H, Data Returned:

BufferSize: 0x1E

Flags: 0x1



NumPhysCylinders: 0x0

NumPhysHeads: 0x0

NumPhysSectorsPerTrack: 0x0

NumPhysSectors: 0x01A8C7800

BytesPerSector: 0x200

EDDPointer: 0xFFFFFFFF

DevicePathSignature: 0xFABEDD

DevicePathLength: 0x2C

HostBusType: 'PCI '

InterfaceType: 'RAID '

Interface path:

BusNumber: 0xFF

SlotNumber: 0xFF

FunctionNumber 0xFF

ChannelNumber: 0xFF

Device path:

RaidArrayNumber: 0x0

PortBitMap 60h: 0x3

Total Test(s): 44; PASS: 44 FAIL: 0

3.3.2.2.2 Rcmpsata.efi/rcmpssata.efi output example

RCmpSata.efi

Locating Intel RSTe SATA Controller:

Intel RSTe SATA Controller Found

Vendor ID: 8086

Device ID: 2826 (Valid)



Recognized RPID found: C610 Series SATA Premium (0x8D06)

PCI device 8086:2826 configuration space:

Revision ID: 0x04

Class code: 0x010400

Cache line size: 0x00

Latency timer: 0x00

Header type: 0x00

BIST capable: 0x00

Base Address Registers:

Address	Type	Size
2088	I/O	0008
20AC	I/O	0004
2080	I/O	0008
20A8	I/O	0004
2020	I/O	0020
91A20000	Mem32	00000800

Cardbus CIS ptr: 0x00000000

Subsys vendor: 0x8086

Subsys ID: 0x7270

Exp. ROM base: 0x00000000

Caps ptr: 0x80

Reserved1: 0x000000

Reserved2: 0x00000000

Interrupt line: 0x0A



Interrupt pin: 0x01

Min GNT: 0x00

Max LAT: 0x00

Checking ACPI support protocol

Checking required protocols

SATA Driver found.

Protocols Supported:

Driver Binding Protocol: Installed

Component Name Protocol: Installed

Component Name 2 Protocol: Installed

Driver Supported EFI Version Protocol: Installed

RAID Config Protocol: Installed

Version Table

ACPI table and EFI variable version tables match.

Version Table Found:

Version Header Revision: 2.2

ROM Version: 4.3.0.1018

Version Table Checksum: Valid

Supported RAID Levels: RAID0, RAID1, RAID10, RAID5

Max Disks per Array: 8

Max Disks Supported: 8

Max Volumes per Array: 2



Supported Strip Sizes:

2k 4k 8k 16k 32k 64k 128k 256k 512k 1M 2M 4M 8M 16M 32M 64M

0 1 1 1 1 1 1 0 0 0 0 0 0 0 0 0

OROM Supported Attributes:

RAID0

RAID1

RAID10

RAID5

>=2TB DISK

BAD BLOCK MGMT

2TB VOLUMES

CHECKSUM VERIFY

OROM Supported Capabilities:

ESATA

HDD PASSWORD

DISK COERCION

OROM Supported Driver Features:

HDD UNLOCK

SGPIO

RSTe ENTERPRISE

ROHI

READ PATROL



Feature	OCF	BCFS
RAID0	X	X
RAID1	X	X
RAID10	X	X
RAID5	X	X
RRT	-	X
OROM UI/BANNER	-	X
HDDUNLOCK	X	X
LED LOCATE	X	X
RRT on ESATA	X	X
NVM	-	X
DELAYBIT0	X	-
DELAYBIT1	X	-
RESERVED	X	-
RESERVED	-	-
RESERVED	-	-
RESERVED	-	-

---

1. Intel RSTe SATA Controller Presence Tests

1.1 Controller Present: PASS

1.2 CPU Brand String: FAIL ()

1.3 Recognized RPID: PASS

2. System-Provided Protocols and Tables



2.1 ACPI Support Protocol: PASS

3. Driver-Provided Protocols and Tables

3.1 EFI Driver Protocols: PASS

3.2 Version Table (ACPI Table and Variable): PASS

Number of successes: 5, Number of failures: 1, Total tests: 6

3.3.2.3 Intel® RSTe SCU RAID Configuration

Like the SATA/sSATA controllers, the Intel® C600 series chipset SCU controller also has two sets of configuration tools. There is a DOS based utility that runs from a DOS bootable USB memory stick. This tool is used when the SCU controller is configured to boot using the RAID Legacy Option ROM.

The other tool is a UEFI based tool that runs from a USB memory key from within the system's UEFI Shell environment. This tool is used when the SCU controller is configured to boot using the RAID UEFI driver image.

Tool information similar to the information provided above.

3.3.2.4 Intel® RSTe SCU RAID Compatibility

Like the SATA/sSATA controllers, the Intel® C600 series chipset SCU controller also has two sets of RAID compliance tools. There is a DOS based utility that runs from a DOS bootable USB memory stick. This tool is used when the SCU controller is configured to boot using the RAID Legacy Option ROM.

The other tool is a UEFI based tool that runs from a USB memory key from within the system's UEFI Shell environment. This tool is used when the SCU controller is configured to boot using the RAID UEFI driver image.

Tool information similar to the information provided above

3.3.2.5 Intel® RSTe UEFI LED Test Tool

The Intel® RSTe UEFI LED test tool is designed to work with the SATA and sSATA controllers of the supported platforms. The purpose of this tool is to support the OEM/ODM manufacturing environments in testing out the SGPIO LED functionality. This tool is run from a USB memory key from within the systems UEFI environemtn. To use this tool the SATA and/or sSATA controllers must be configured for RAID



mode. There also has to be an SGPIO complient backplane contained within the platform.

Ledtoolsata.efi/Ledtoolssata.efi

LedTool.efi [/?] [/HELP] [/D:state_number] [/D0-5:state_number] [/TIMEOUT:time]
/? Same as /HELP. Displays Help Screen. Other options ignored.
/D Send the specified state to all ports. /D0-9 option will be ignored.
/D0-9 Send the specified state to selected port.
/TIMEOUT After sending the states waits the specified number of seconds and then turns off all leds and performs a resets the Messages sending logic.

NOTE: After a reset, the first sent state will turn on all other LEDs.

Valid state_number is from 0 to 9 and they are interpreted as following: 0 = NO_DISK, 1 = NO_ACTIVITY, 2 = ACTIVITY, 3 = LOCATE, 4 = FAIL, 5 = REBUILD, 6 = PFA, 7 = HOTSPACE, 8 = CRIT_ARRAY, 9 = FAIL_ARRAY

Example use: "Ledtool.efi /D2 0 /D0 2 /TIMEOUT 2"
this will set state NO_DISK on port 2 and state ACTIVITY on port 0. After 2 secs delay Messages will be reset.

3.3.2.6 Intel® RSTe UEFI Agentless Support

Intel® RSTe 4.3 provides support for Agentless reporting to provide the required information to the BIOS. There is a tool that OEM/ODMs can use to test this interface to make sure that it is implemented correctly. This feature requires the system BIOS to contain the Intel® RSTe 4.3 based UEFI driver. The tool just prints out the system information associated with agentless manement support.

AimTestTool.efi

3.3.2.7 Intel® RSTe Disk Clearing Utility

Intel® RSTe 4.3 provides support to clear the metadata and boot records from an existing to allow it wo be reused.

NOTE: Intel recommends that new disk drives always be used with build a system.

Rclrsata.efi/Rclrssata.efi

Intel(R) UEFI RAID Clear Metadata Utility for Serial ATA

--HELP--



: Selects disks to have deleted metadatas.
List should be delimited by spaces. Evaluation example:
'RCfgSata.efi 0 1 2' - it clears metadata in all selected drives

/I Displays Drives Information Screen.
/? Displays Help Screen.

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 Options:

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 0-1-0-0 0-2-0-0 *(format of the disk ID is "0-SATA_Port-0-0" where the second digit from the left represents the SATA port on the platform where the disk is located; thus 0-1-0-0 represents SATA port # 1)*

--create -l 0 -z 5 --name RAID0Volume 0-3-0-0 0-4-0-0 0-5-0-0

-C -l 1 -E 0-1-0-0 -n VolumeWithData 0-2-0-0

-C --rrt -n RRTVolume 0-1-0-0 0-2-0-0 --rrtMaster 0-1-0-0

-C --rrt -n RRTVolume 0-1-0-0 0-2-0-0 --rrtUpdate Continuous

--create --help

Information Options:

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 0-5-0-0

--information --array Array_0000

--information --help

Manage Options:

Flag	Name

Intel Confidential



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

Modify Options:

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 0-3-0-0 0-4-0-0
-m --volume ModifyVolume --level 5
--modify -v Volume -n RenameVolume
--modify --help

Accelerate Options:

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 0-3-0-0 --cache-size **X** [where 16 ≤ **X** ≤ 64]
--accelerate --setAccelConfig --disk-to-accel 0-5-0-0 --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

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.

--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.

--disassociate
Disassociates the Cache volume from acceleration

--disk-to-accel <<host>-<bus>-<target>-<lun>>
Specifies a disk if accelerating 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.

-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.

-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.

-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.

-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 enterprise(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.



4 Intel® RSTe Pre-OS Usage

This section will address the various Intel® RSTe Pre-OS components in more detail. It will address the usage of the different compnents and outline the system or platforms requirement need to properly support the usage of the Intel® RSTe family of products.

4.1 Intel® RSTe Pre-OS Compoments

This section will outline the Pre-OS or BIOS environment that comes with the Intel® RSTe family of products an the BIOS requirement for properly incorporating the Intel® RSTe Pre-OS compoents into the system or platform BIOS image.

4.1.1 UEFI System BIOS and the Intel® RSTe UEFI/RAID Package

Intel® RSTe provides UEFI drivers for OEMs and their BIOS vendors to integrate into their RAID-enabled platforms (Not required for AHCI mode platforms).

4.1.1.1 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® RSTe 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)
UEFI Shell Specification version 2.0	(http://www.uefi.org/specsandtesttools)

4.1.1.2 What Intel® RSTe Provides to OEMs/BIOS Vendors

4.1.1.2.1 Intel® RSTe UEFI Driver

These are the main components of the Intel® RSTe pre-OS EFI solution. It is provided in three different formats:

SataDriver.efi or sSataDriver.efi:

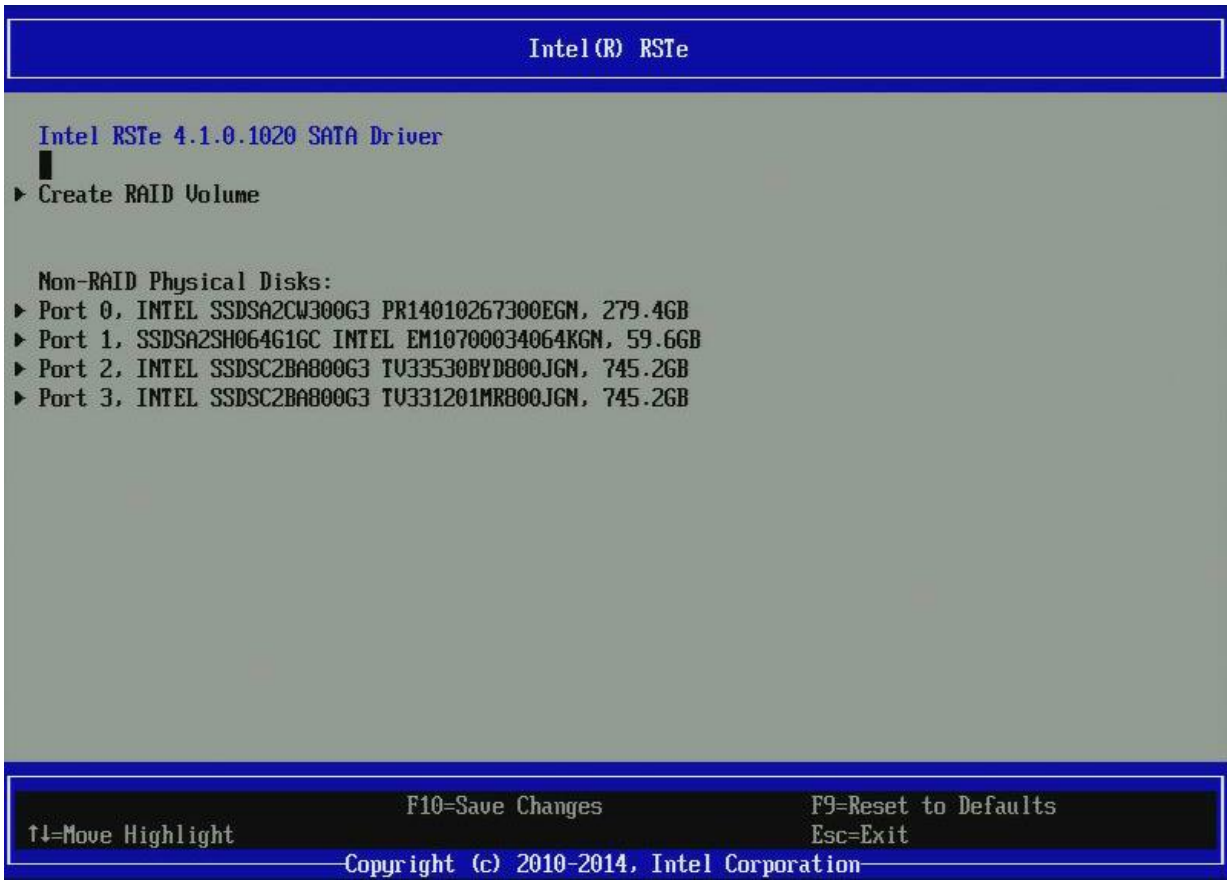


UEFI driver that requires integration into the UEFI System BIOS by the OEM's BIOS vendor. This file can be placed into the OEMs' UEFI BIOS source build where their tools can integrate it.

4.1.1.2.2 Intel® RSTe UEFI User Interface

An **HII-compliant** user interface is provided for the pre-boot configuration of the RAID system. The same functionality provided in the legacy OROM UI is available in this HII UI.

- The HII UI is integrated within the UEFI driver binary (RaidDriver.efi)
- Per the UEFI specification, we publish the HII UI as string and forms packages
- The HII 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 enterprise' will be displayed as the selection to enter the HII UI
- Some OEMs may want to hard assign where the Intel® RSTe UEFI GUI will be located within their BIOS.





- **The Intel® RSTe UEFI Driver FORMSET_GUID is:**
  FORMSET_GUID { 0xd37bcd57, 0xaba1, 0x44e6, { 0xa9, 0x2c, 0x89, 0x8b, 0x15, 0x8f, 0x2f, 0x59 } }
  {D37BCD57-ABA1-44e6-A92C-898B158F2F59}

4.1.1.3 Command line RAID Configuration Utility

RcfgSata.efi and ***RcfgsSata.efi***:

- 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/RcfgsSata.exe***) in previous releases of the Intel® RSTe product.
- Requires the implementation of the same RSTe UEFI Driver version in the system BIOS in order to operate properly.

```
fs1:\> RCfgSata.efi /st


--SATA CONTROLLER--
--VOLUME INFORMATION--
    (None found)

--DISK INFORMATION--

ID:             0x0
ID bitmap:      1
State:          Non-RAID Disk
Size:           111.7GB
Serial Number:  CV223101A9120BGN
Model Number:   INTEL SSDSC2CW120A3

ID:             0x1
ID bitmap:      2
State:          Non-RAID Disk
Size:           59.6GB
Serial Number:  EM107101EM064KGN
Model Number:   SSDSA2SH064G1GC INTEL

fs1:\> _
```

4.1.1.4 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 RSTe UEFI Driver are present. Also provides a list of the protocols published by the RSTe UEFI Driver and the capabilities/features of the RSTe UEFI Driver.

```
fs1:\> RCmpSata.efi > rcmpsata.txt

fs1:\> _
```



4.1.1.5 UEFI System BIOS Requirements for Platform Compatibility with Intel® RSTe UEFI

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

4.1.1.5.1 Required Protocols/Functions to be Provided by the UEFI System BIOS

The Intel® RSTe 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
- CopyMem
- LocateProtocol

EFI_RUNTIME_SERVICES:

- SetVariable
- GetVariable
- GetTime

Other Protocols:

- EFI_ACPI_TABLE_PROTOCOL (or EFI_ACPI_SUPPORT_PROTOCOL (EDK117))

4.1.1.5.2 Optional Protocols/Functions to be Provided by the UEFI System BIOS

If the OEM plans to use the Intel® RSTe 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



4.1.1.5.3 Protocols Provided by the Intel® RSTe UEFI Driver

The Intel® RSTe 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)

4.1.1.6 How-to-Enable the Platform with Intel® RSTe 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® RSTe UEFI driver.

4.1.1.6.1 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® RSTe UEFI UI

4.1.1.6.2 Step2: Download and Integrate the Intel® RSTe UEFI Package

1. Download the latest kit from the Intel VIP (Validation Internet Portal) website. From the kit select the ***PreOS zip*** file (e.g. PreOS-1028.zip) which will contain all of the RSTe PreOS images and tools. Go to the efi_sata and efi_ssata directories to find the UEFI driver binary file (*SataDriver.efi and sSATADriver.efi*)
2. Select and extract the binary file based on the planned integration method:
   - ***SataDriver.efi***: Use this binary if planning to integrate at the time of the BIOS image build to support the SATA Controller



- ***sSATADriver.efi***: Use this binary if planning to integrate at the time of the BIOS image build to support the sSATA Controller

3. Use the proper integration tools based on the binary file selected above

4.1.1.6.3 Step3: Verify Compliance

1. From the extracted PreOS files 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

4.1.1.7 Known Compatibility Issues with the 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.

4.1.1.7.1 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® RSTe 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® RSTe 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."

4.1.1.7.2 EXT_SCSI_PASSTHRU_PROTOCOL Test

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath - call BuildDevicePath with NULL DevicePath.

- *Test Index*: 5.7.8.2.1

4.1.1.7.3 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® RSTe UEFI driver does not support RouteConfig – EFI_UNSUPPORTED is returned. RouteConfig is not supported so that Intel® RSTe HII form values are only modified by the Intel® RSTe 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 RSTe UEFI Driver. The RSTe 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 RSTe UEFI Driver. The RSTe software does not use it as it is not necessary.

4.1.2 Legacy OROM System BIOS Requirements and the Intel® RSTe Legacy OROM Package

Intel® RSTe provides Legacy OROM images for OEMs and their BIOS vendors to integrate into their RAID-enabled platforms (Not required for AHCI mode platforms).

Due to limited shadow RAM and EBDA space, and the fact that a platform may require multiple OROMs be loaded for other functions, there might not be enough runtime space for both the RSTe SATA RAID controller OROM and the RSTe sSATA RAID controller OROM to provide int13h support simultaneously. Even so, each RSTe OROM still needs to initialize so that it can configure each controller based on platform dependencies and store data needed by the OS drivers in the Shadow RAM area even if it does not provide full int13h runtime support. This outlines what is required to be able to allow/require the user to select the boot controller.



RSTe OROM requires that the BIOS is V2.0c+ compliant.

INT 1A functions:

AX = b101h: PCI BIOS installation check
AX = b102h: PCI BIOS find PCI device
AX = b108h: PCI BIOS read configuration byte
AX = b109h: PCI BIOS read configuration word
AX = b10ah: PCI BIOS read configuration dword
AX = b10dh: PCI BIOS write configuration dword.

4.1.2.1 Intel® C610 series Chipset Legacy OROM Boot Option

4.1.2.1.1 INT 15 / AX=F300h / BX=0002h (Get RSTe OROM Boot Info)

Description:

Through this function, BIOS provides user-settable RSTe boot information to the RSTe legacy OROM drivers. These values are visible to the user through the BIOS Setup menus. The menu options should be linked to legacy OROM selections in the PCH-IO section.

Inputs:

AX = F300h (Function)

BX = 0002h (Sub Function)

EAX = 0000F300h

EBX = 4F450002h ('OE') + Sub Function

EDX = 424F4F54h ('BOOT')

Normal Outputs:

CF = clear if successful

EAX = 424F4F54h ('BOOT')

BL = legacy_orom_boot_controller_selection:



Due to limited shadow RAM and EBDA space, and the fact that a platform may require multiple OROMs be loaded for other functions, there might not be enough runtime space for both the RSTe SATA RAID controller OROM and the RSTe sSATA RAID controller OROM to provide int13h support simultaneously. Even so, each RSTe OROM still needs to initialize so that it can configure each controller based on platform dependencies and store data needed by the OS drivers in the Shadow RAM area even if it does not provide full int13h runtime support. Thus, through this setup option BIOS can avoid the runtime space conflict by allowing the user to select the boot controller according to the following values:

0000h = No runtime space restrictions. BIOS indicates that both RSTe SATA and sSATA runtime code should provide full int13h support for RSTe devices.

(NOTE: The BIOS should allow this option if it knows that there is room in shadow RAM for both OROMs' runtime code. If the BIOS can always guarantee this condition, then it does NOT need to make Legacy OROM boot controller selection visible to the user in BIOS setup.)

0001h = The sSATA controller is selected as boot controller. BIOS will load RSTe SATA OROM first, but the SATA OROM will only initialize and then leave pertinent RAID configuration information for the SATA OS RAID Driver in runtime space. The RSTe sSATA OROM will initialize, relocate to runtime space, and provide full int13h support for sSATA attached devices.

0002h = The SATA controller is selected as boot controller. BIOS will load RSTe sSATA OROM first, but the sSATA OROM will only initialize and then leave pertinent RAIDconfiguration and sSATA OEM parameter information for the sSATA OS RAID Driver in runtime space. The RSTe SATA OROM will initialize, relocate to runtime space, and provide full int13h support for SATA controller attached devices.

0003h = Neither SATA nor sSATA controller is selected as boot controller. Boot support is being provided by anotherdevice. BIOS will load both RSTe OROMs, but each will only initialize and leave pertinent RAID configuration and sSATA OEM parameter information



for the RSTe OS RAID Drivers in runtime space. Therewill NOT be int13h support for RSTe devices.

Error Outputs:

CF = set on error

AH = error code

? = 86h Function Not Supported = (boot_controller_selection = 0000h assumed)

4.1.3 Choosing Boot Mode (Legacy Option ROM vs UEFI)

Intel® RSTe Pre-OS environment does not support loading both type of Pre-OS imagates (Legacy Option ROM vs. EFI driver). The platform BIOS environment must provide an option to choose one or the other. When the Legacy Option ROM is selected as the boot mode, the BIOS should not load the Intel® RSTe UEFI driver. Likewise, when UEFI mode is selected, the BIOS should not load the Intel® RSTe Legacy Option ROM images.



5 New in Release Version 4.3

5.1 General Overview

This section will outline deatild information on the key featues included in the Intel® RSTe release package.

5.1.1 Windows 10 OS Upgrade Recommendations

Intel® RSTe 4.3.0.1219 PCH/SATA release package and Intel® RSTe 4.3.0.1681 NVMe release package both support Windows 10.

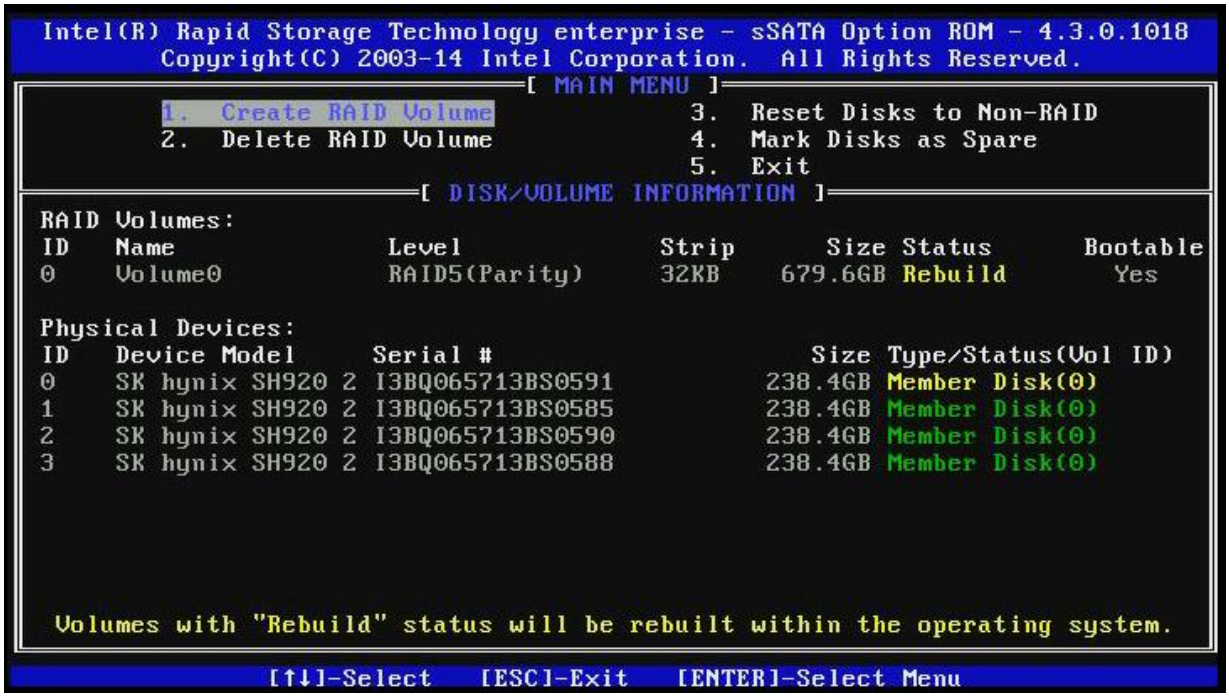
OS	Steps:
Windows 8.1 x64 **Windows 8 x64** **(Legacy / UEFI)**	• Download Windows10 x64 iso from official Microsoft Insider website: https://insider.windows.com/ • Intel recommends upgrading the system to 4.3.0.1199 prior to performing the Windows 10 upgrade process. a. For SATA/sSATA driver the minimum version required to upgrade to Windows10 is the **4.1.0.1047 Kit** b. For the SCU driver the minimum version required to upgrade to Windows 10 is the **4.2.0.1143 Kit** • Open iso and launch setup.exe to update directly from current system **Windows 8.1 x64 / Windows 8 x64** • Follow the Windows 10 setup steps
Windows 7 x32 / x64 **(Legacy / UEFI)**	1. Download Windows10 x32 / 10 x64 iso from official Microsoft Insider website: https://insider.windows.com/ • 2. Intel recommends upgrading the system to 4.3.0.1199 prior to performing the Windows 10 upgrade process. a. For SATA/sSATA driver the minimum version required to upgrade to Windows10 is the **4.1.0.1047 Kit** b. For the SCU driver the minimum version required to upgrade to Windows 10 is the **4.2.0.1143 Ki** 3. Open iso and launch setup.exe to update directly from current system **Windows 7 x86 / Windows 7 x64** 4. Follow the Windows10 setup steps

5.1.2 RSTe Pre-OS RAID Volume Rebuild Drive Identification

With the release of RSTe 4.3.0.1219 PCH/SATA production package, the RSTe Pre-OS user interfaces will begin to show drives in a RAID Rebuild state as “Online

member" is a yellow colored text. This will be seen in both the RSTe Legacy Option ROM as well as the RSTe UEFI drivers.

The following is an example of what is seen in the Legacy OROM user interface.



5.1.3 SCSI Pass-Thru Support

The Intel® RSTe 4.3.0.1219 PCH/SATA release package an the Intel® RSTe 4.3.0.1681 NVMe release package both provide support for Extended SCSI Passthru commands. Specifically:

- EFI_EXT_SCSI_PASS_THRU_PROTOCOL protocol as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D. The protocol allows SCSI request packets and ATAPI command blocks to be sent to a device on the following controllers:
  - SCU - only physical disk devices supported
  - SATA/sSATA - only ATAPI devices supported
- Enabled support for EFI_DEVICE_PATH_PROTOCOL for physical disk devices on SCU controller as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D. The protocol allows a physical device behind the handle object to be identified.



- Enabled support for EFI_SCSI_IO_PROTOCOL for physical disk devices on SCU controller as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D. The protocol allows basic operations to be performed on SCSI devices.

5.1.4 Device Information Display in UEFI

The Intel® RSTe 4.3.0.1219 PCH/SATA release package includes the support for the RSTe UEFI driver to report the physical port for devices connected to the SATA/sSATA controllers. In some of the releases of the Intel® RSTe UEFI driver, the device information provided (i.e. to the UEFI shell environment) was based off of the enumeration values created during the discovery of the devices attached. To support a manufacturing environment that relies of this information to identify the physical port the device is connected too, the UEFI driver now reports out the physical port value instead of the enumerated value. The data is displayed as a set of 3 values in the following order: X-Y-Z

X – 0: passthru disk, 1: volume

Y– PHY number: 1 (phy0), 2 (phy1), 4 (phy3), 8 (phy3), 16 (phy4), 32 (phy6), 64 (phy7)

Z – disk number on PHY (in case of an expander)

5.1.5 Intel® RSTe PCH/SATA SGPIO LED Management

5.1.6 Common Storage Management Interface Support (CSMI)

Intel® RSTe 4.3.0.1219 PCH/SATA release package does not support application development that requires interface access to the driver via Intel® RSTe API method. However Intel® RSTe 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® RSTe. 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.

5.1.6.1 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

5.1.7 Disk Coercion

The Intel® RSTe 4.3.0.1219 PCH/SATA release package and the Intel® RSTe 4.3.0.1681 NVMe release package both provide support for Disk Coercion. When a RAID volume is created, this feature will analyze the physical disks and will automatically adjust (round down) the capacity of the disk(s) to 95% of the smallest physical disk. This allows for the variances in the physical disk capacities from different vendors.

The option is available at the time that the RAID volume is created to consume the entire RAID volume if desired. When this is done, it increases the changes of issues when replacing a failed drive. Intel always recommends using the same vendor and sizes for drives in the RAID array.

5.1.8 Mechanical Ejection Support

The Intel® RSTe 4.3.0.1219 PCH/SATA release package provides support for a drive to be ejected from an enclosure managed by a mechanical ejection switch. The feature is enabled by default and is functional as long as the BIOS supports this functionality. If this feature needs to be disabled, add the following registry key:

Enable:
HKLM, System\CurrentControlSet\Services\iaStorA\Parameters\Device, EJ0IsDefined, %REG_TYPE_DWORD%, 0x00000001

Disable:
HKLM, System\CurrentControlSet\Services\iaStorA\Parameters\Device, EJ0IsDefined, %REG_TYPE_DWORD%, 0x00000000



5.1.9 Flush Unit Access (FUA)

Intel® RSTe 4.3.0.1219 PCH/SATA release package and the Intel® 4.3.0.1681 release package both honor all FUA commands received from the OS. This ensures that the cache on the attaced drive is flushed whenthe command is received. This can have a negative impact on performance. For customer who are willing to accept the risk of potential data loss can improve the overall platform performance by having Intel® RSTe ignore FUA commands received.

Enable:
HKLM, System\CurrentControlSet\Services\iaStorA\Parameters\Device, IgnoreFUA, %REG_TYPE_DWORD%, 0x00000001

Disable:
HKLM, System\CurrentControlSet\Services\iaStorA\Parameters\Device, IgnoreFUA, %REG_TYPE_DWORD%, 0x00000000

5.1.10 Intel RSTe NVMe pass-through IOCTL support

The RSTe DRIVER provides a private API that allows user-space applications to send and execute NVMe commands to remapped PCIE NVMe devices. This new API is based on the new IOCTL definition to implement NVMe pass-through channel.

- Supported:
  - Pass-through Disks
- Unsupported:
  - RAID Volumes
  - SRT Cache Devices

5.1.11 Intel RSTe NVMe Compliance

The RSTe SOLUTION adheres to the NVMe Specification.

5.1.12 Intel RSTe RAID Volume Support with Multiple Controllers

The Intel® RSTe NVMe release package provides support for the creation of RAID volumes that across multiple Intel® NVMe SSD controllers when attached to PCIe express slots managed by the CPU. When installing the Intel® RSTe NVMe package, please take care to ensure that there are no Intel® NVMe SSDs attached to PCI express slots managed by the PCH (please refer to platform documentation). When there are Intel® NVMe SSD attached to the PCH, Intel® RSTe NVMe installation process will install the Intel® RSTe NVMe driver against those devices. Those devices



will show up in the Intel® RSTe GUI. Intel® RSTe NVMe package will not impact the operation these devices in this configuration.

NOTE: Intel® RSTe NVMe does not support spanning NVMe RAID volumes across CPU and PCH PCI express slots.

5.1.13 NVMe Admin Security Commands Support

The RSTe DRIVER supports admin security commands from NVMe specification.



6 Intel Rapid Storage Technology enterprisefor PCIe NVMe Storage Devices

6.1 OEM System BIOS Requirements

System BIOS	System BIOS is not required to have any special configuration to support RSTe NVMe.

6.2 General Requirements

Hardware	Intel C610/230 series chipset platafomrs.
Operating System	32 and 64 bit versions of Windows 7; Windows Server 2008 R2; Windows 8.1; Windows Server 2012/2012R2/2016; Windows 10
Supported Devices	Intel NVMe PCIe Devices
SATA Mode	N/A
Specificcation Support	Will adhere to the NVMe Specification

6.3 Feature Limitations

Intel® RSTe NVMe release package will provide support for Intel PCIe NVMe Storage Devices with the following feature limitations:

- No support for:
  - Hot Plug
  - RAID Volume Boot Support



    - Spanning agress PCI Express slots manages by the CPU and those managed by the PCH
    - Third pary NVMe SSDs
- Validated with a maximum of 8 Intel NVMe SSDs
- If used in a RAID volume, all member device must be on a PCI Express slot managed by the CPU

6.4 PCIe NVMe Device Usage Model

The Intel® RSTe NVMe release package will provide support for Intel NVMe SSD in the following usages:

- As a bootable single Pass-Thru device with Intel® RSTe NVMe drivers
- In a Data RAID volumes only (no boot).
- As a spare disk for a RAID volume (has to be on the same bus type as the RAID member devices)

6.5 Intel® RSTe UEFI Compliance Utitlity for PCIe Storage

The Intel® RSTe NVMe release package does not support booting from a RAID volume, so RSTe NVMe does not contains a PreOS UEFI driver. It is not required nor is it needed.



7 Creating a RAID Volume

RAID volumes can be created three different ways.

1. The Intel® RSTe GUI in Windows*.
2. The Intel® RSTe Pre-OS environment
    a. Intel® RSTe Legay option ROM user interface.
    b. Intel® RSTe pre-OS UEFI HII UI.
3. The Intel® RSTe Command Line Interface tools. These tools are restricted to an OEMs manufacturing environments.
    a. Intel® RSTe Pre-OS CLI tools
        i. RCfgSata/RCfgsSata
    b. Windows (or WinPE) based RAID configuration
        i. RSTCLI 32/64 utilities.

7.1 Using the Intel® Rapid Storage Technology enterprise GUI

1. Run the Intel Rapid Storage Technology enterprise GUI from the following Start menu link within Windows:

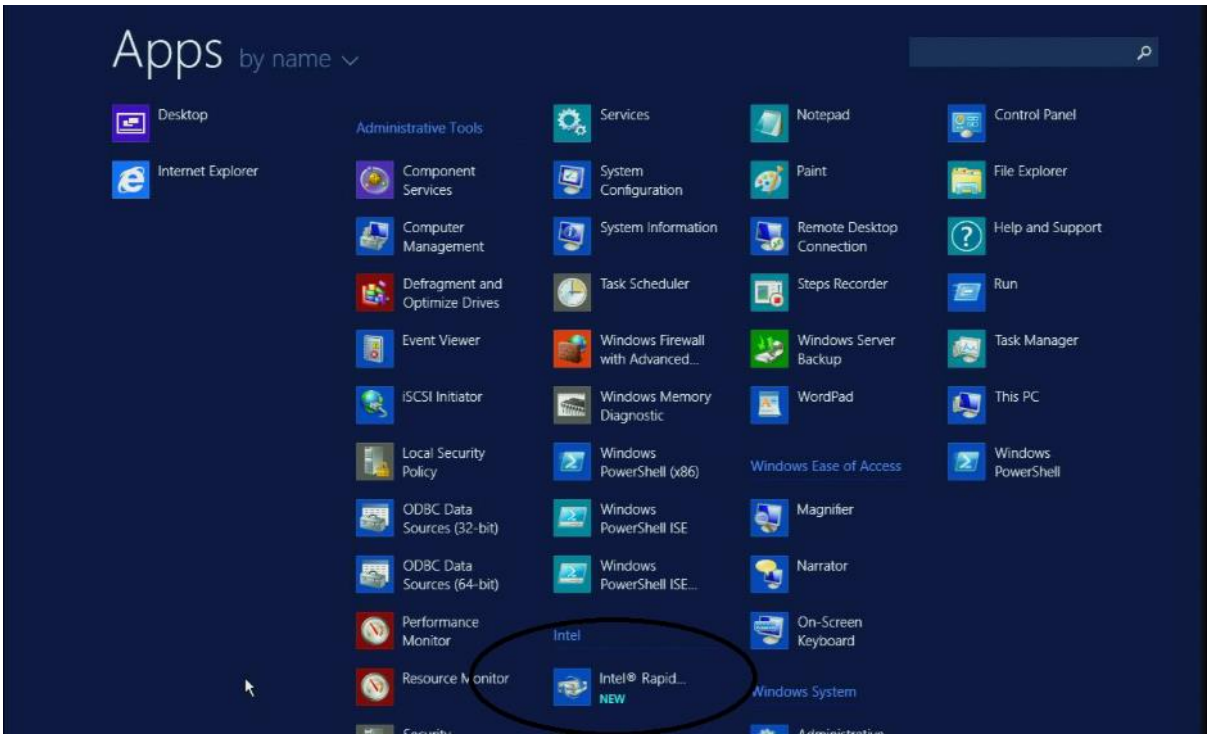
    Pre-Windows* 8

    Start→Programs→Intel Control Center (optional) ->Intel® Rapid Storage Technology enterprise→Intel Rapid Storage Technology enterprise GUI

    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 enterprise->IAStorUI.exe (right click and pin to start)



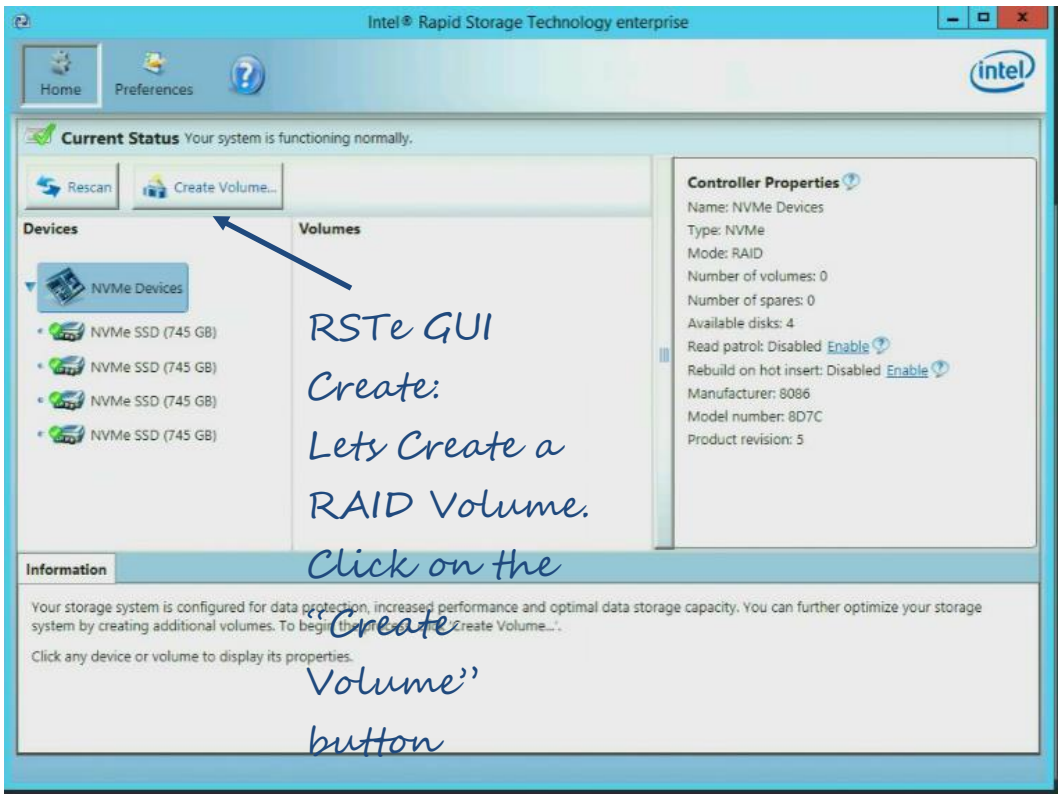


2. Based on the available hardware and your computer's configuration, you may be able to create a RAID volumes on drives by selecting a volume type under 'Create'. Based on the number of non RAID disks available to you and the size of the disks the user will only be able to see the possible volume creation options... (e.g. if you have only two disks ...you can only see options to create RAID 0, RAID1. If you have three disks, you can only see options for creating RAID 0, RAID 1, RAID5)

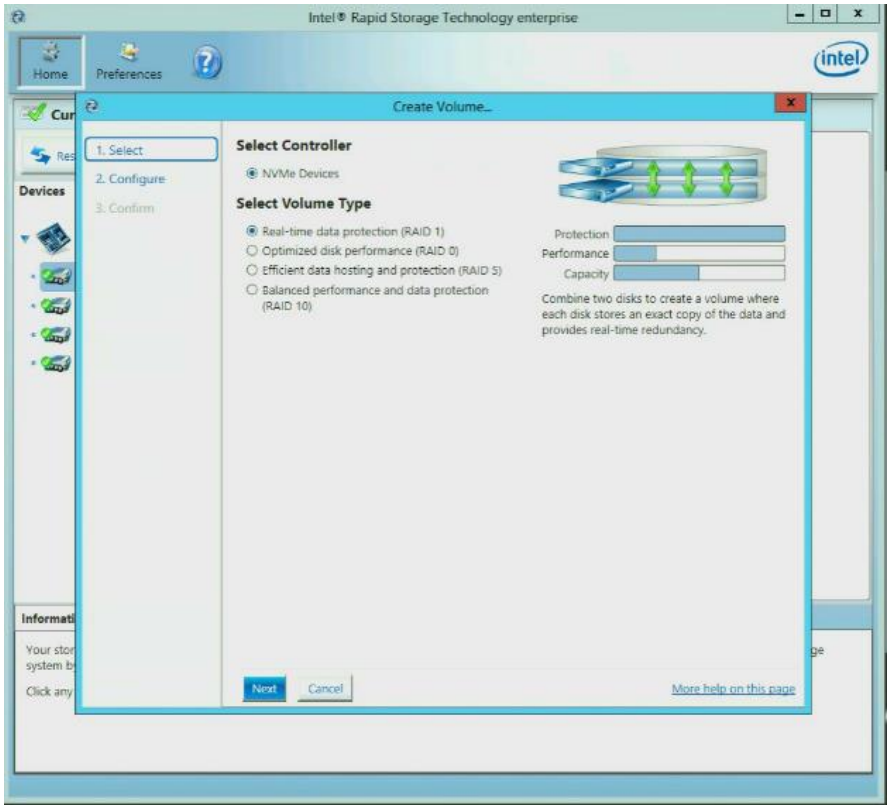
   NOTE: To create a RAID volumes the user must be in admin mode. For Intel® RSTe PCH/SATA, the system must be in RAID Ready mode with two or more disks

connected to the PCH controller(s). For Intel® NVMe the system must have two or more Intel NVMe SSDs attache to PCI express slots managed by the CPU.

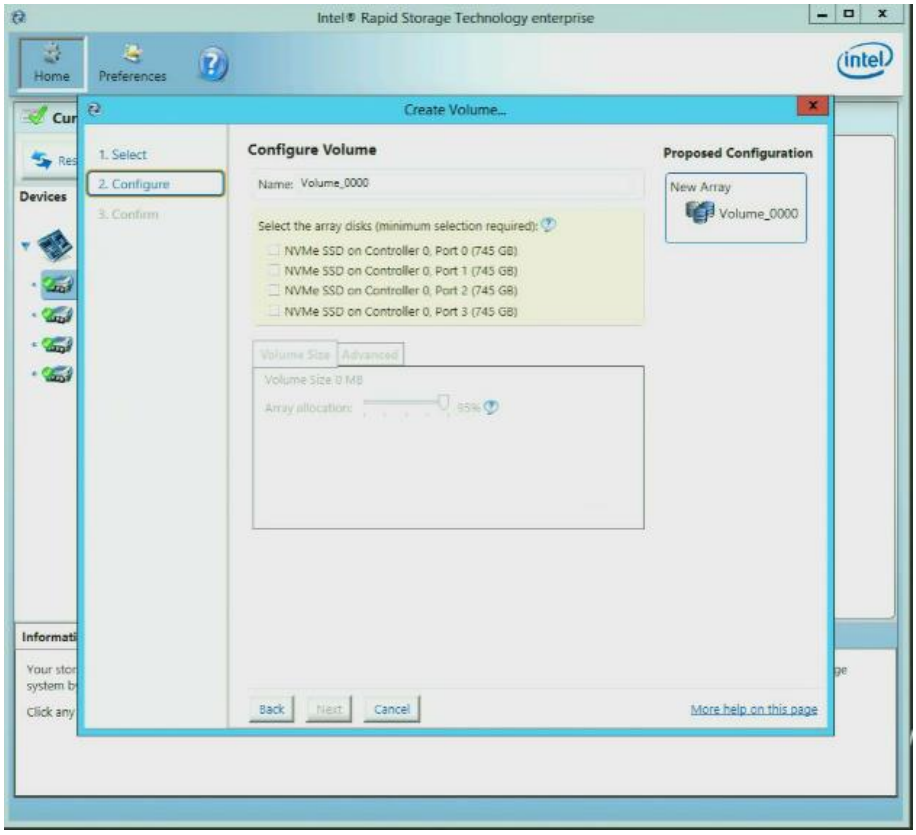
3. Instructions to create a volume by selecting volume type under ‘Create’



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

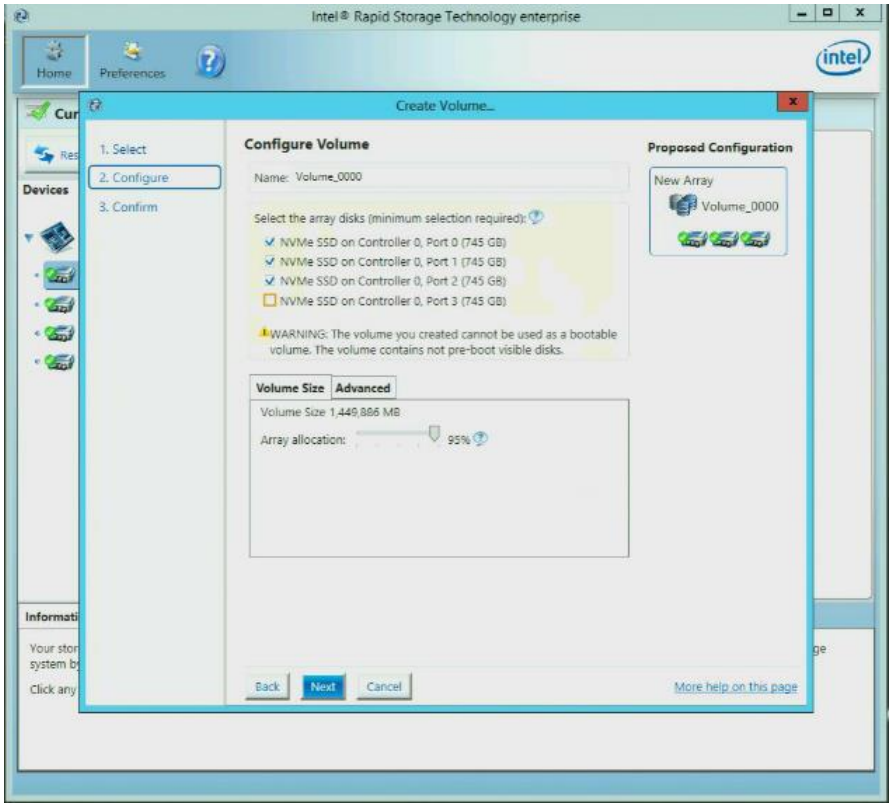


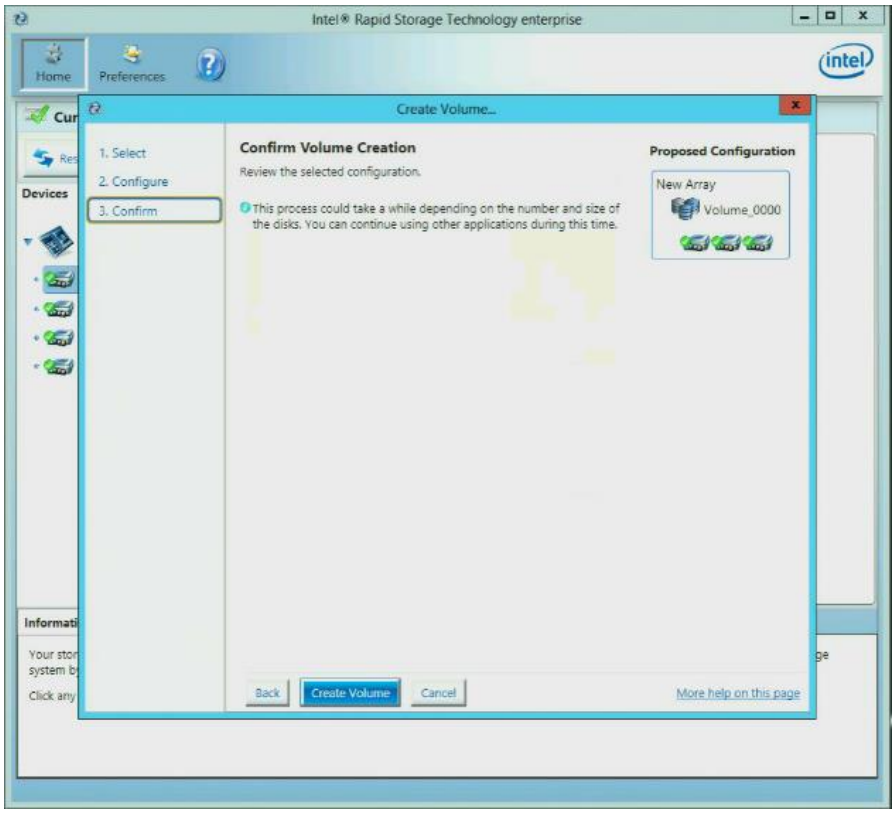
b. Now configure the volume by providing the volume name, selecting the hard disks to be part of the volume and strip size if applicable



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 GUI.

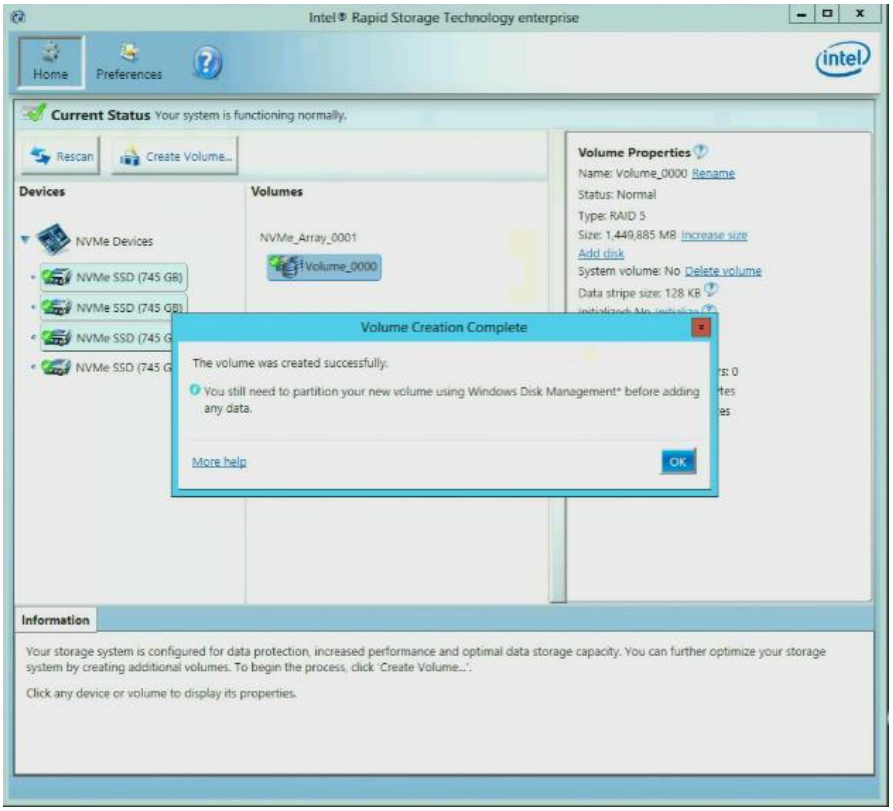
c. Once the disks are selected for volume creation, the user will presented with option, if you want preserve data on which selected disk. Click on 'Next' and select the 'Create Volume' button.





Intel Confidential

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.

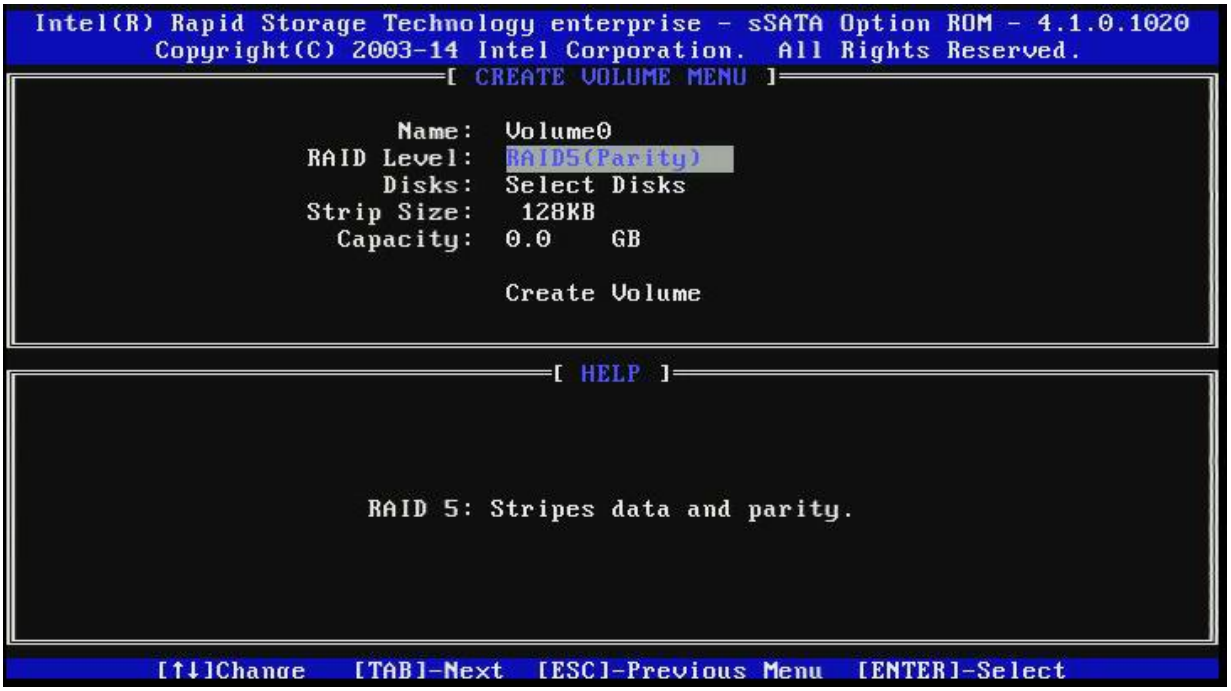
7.2 Using the Intel® Rapid Storage Technology enterprise Legacy Option ROM User Interface

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 enterprise option ROM user interface.

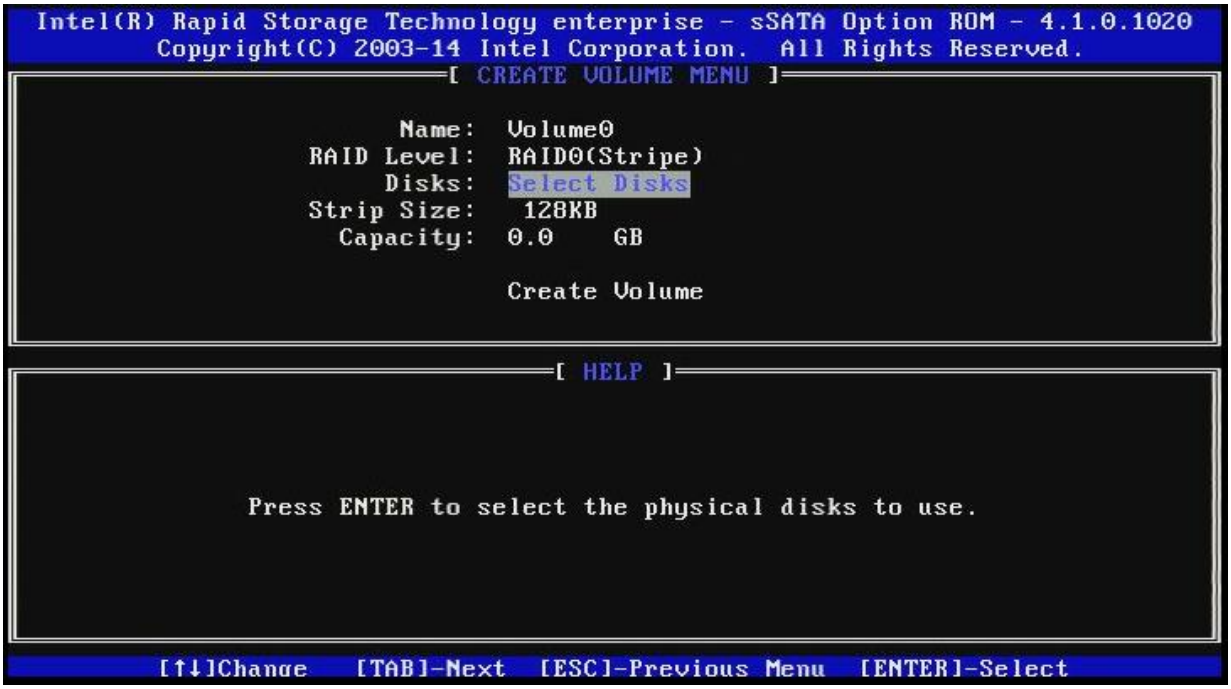
2. In the Main Menu, select option #1 'Create RAID Volume'. Enter the name you want to use for the RAID volume, then press Enter.

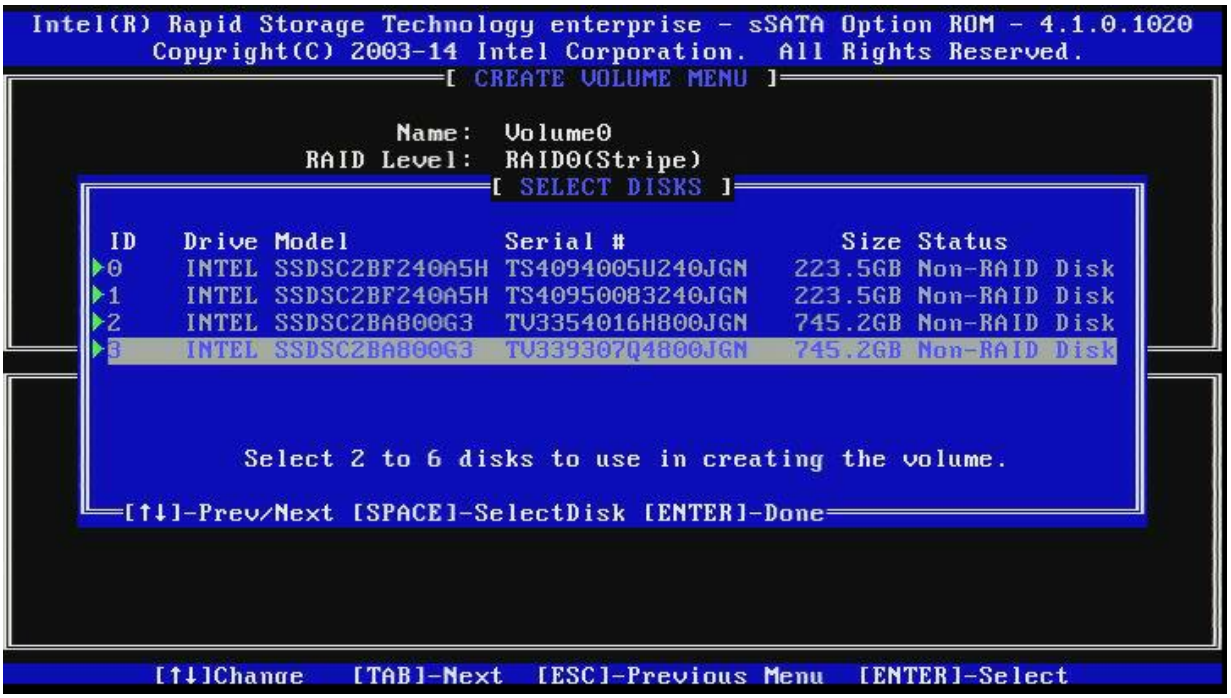


3. Select the RAID Volume Name and the RAID Level

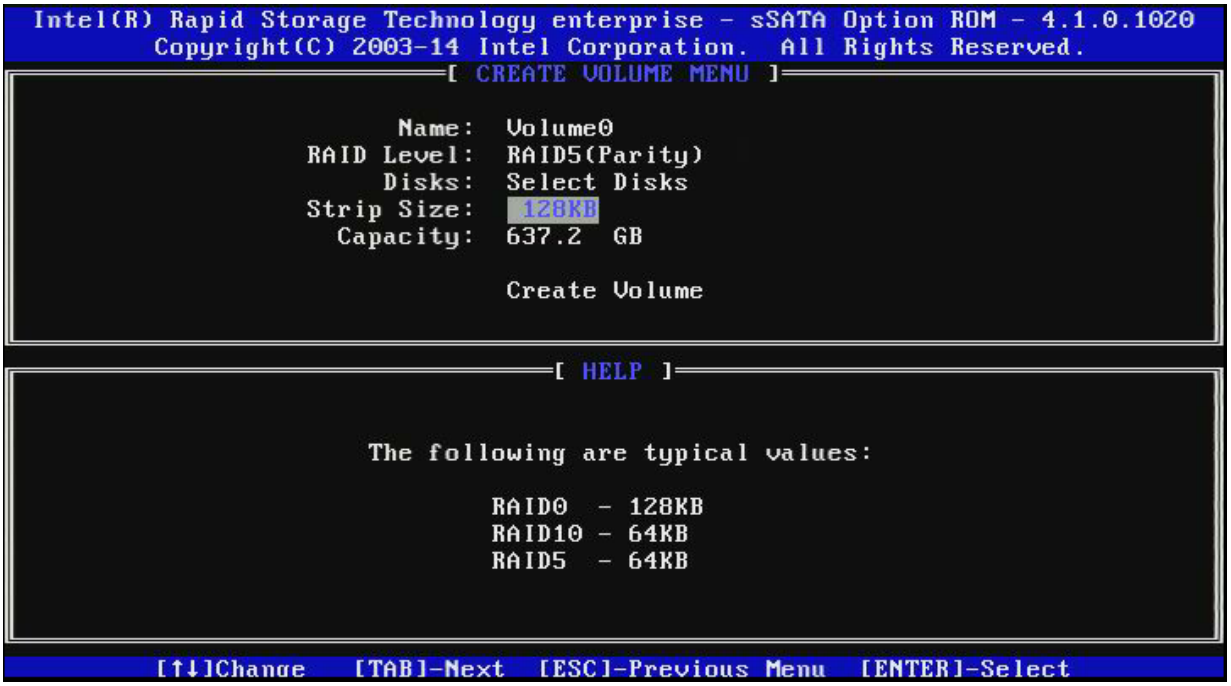


4. Press Enter to select the disks to be used by the array that the volume will be created on. Press Enter when done.

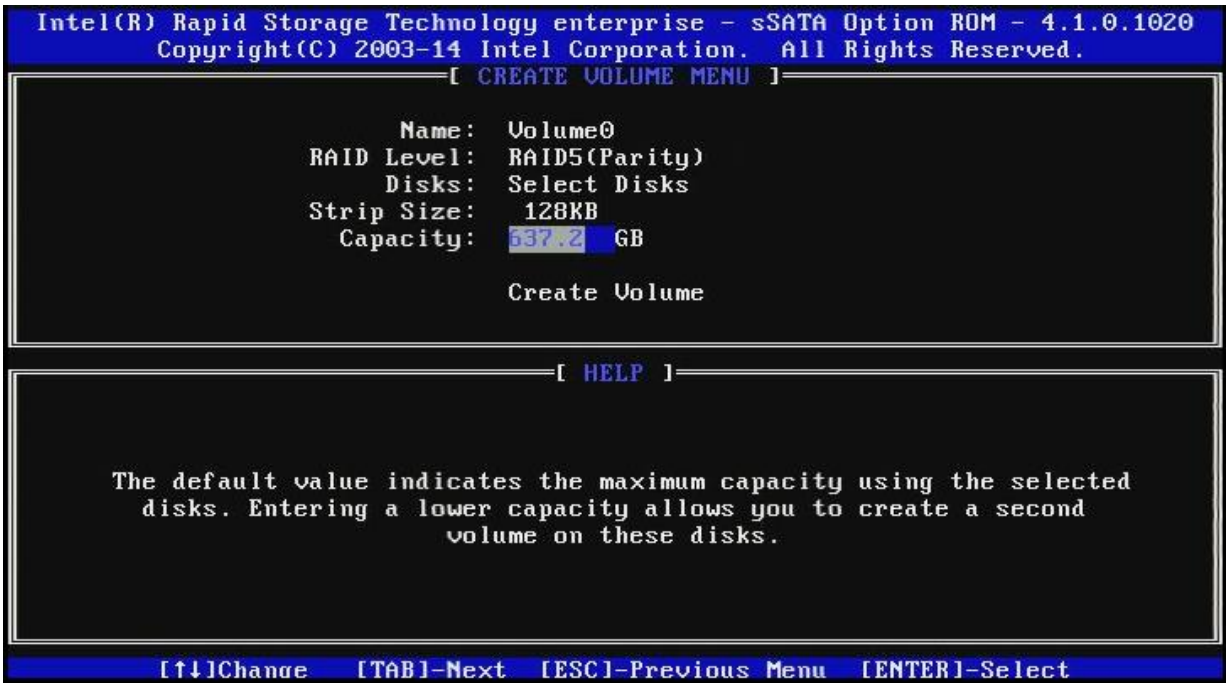




5. Select the strip size (128 KB is the default for RAID 5) by using the arrow keys, then press Enter when done.



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. When done select arrow down to “Create Volumne” and press enter to complete.

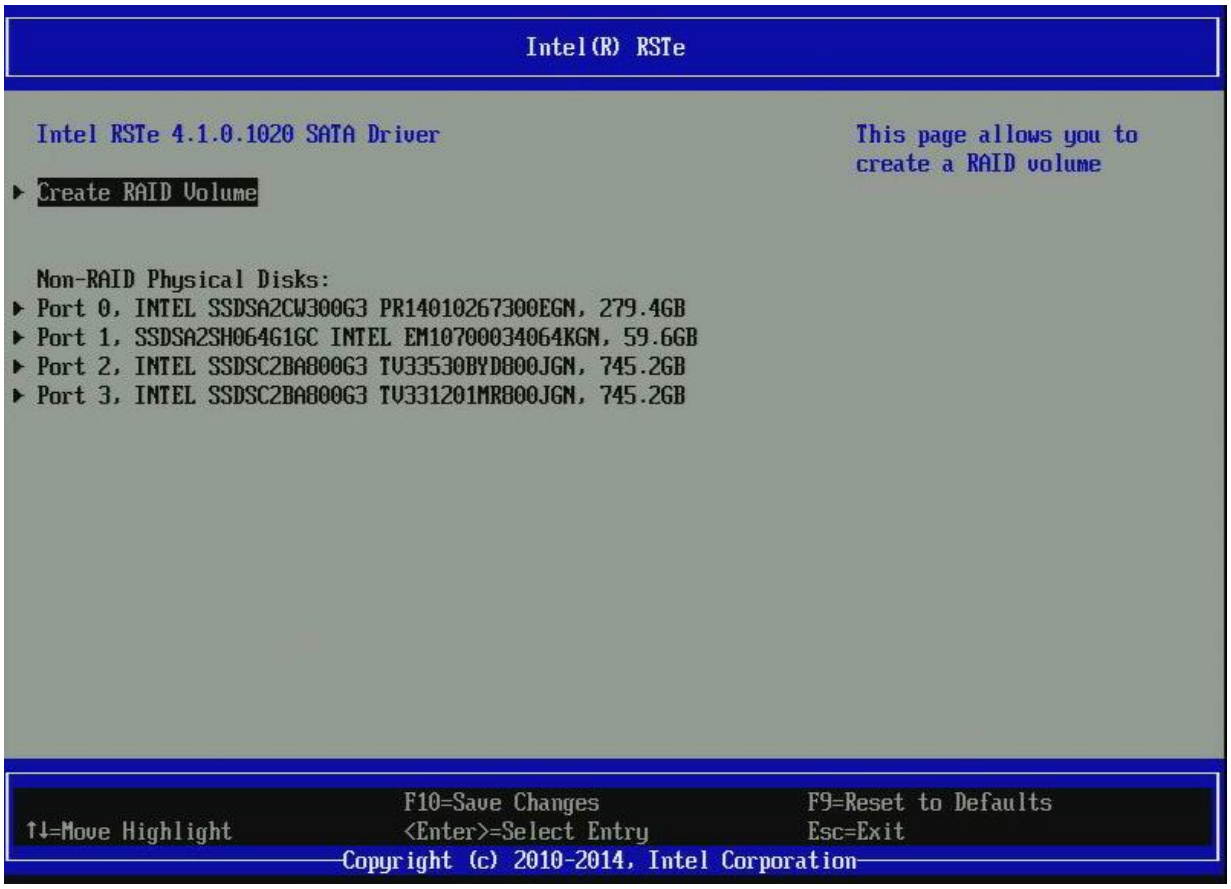


7. After this is done, exit the Option ROM user interface.

7.3 Using the Intel® Rapid Storage Technology enterprise 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, access the Intel® RSTe UEFI user interface (HII compliant) though the BIOS setup menus. Each BIOS vendor may have a different path to reach the RSTe UEFI HII.



The UEFI UI is divided into three main sections:

1. Section 1 is the main section.
    a. Displays RAID configuration and status information
    b. Displays RSTe UEFI driver version (in this example it is 4.1.0.1020)

c. Displays physical devices enumerated by the RSTe UEFI driver that are not part of the RAID volume

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

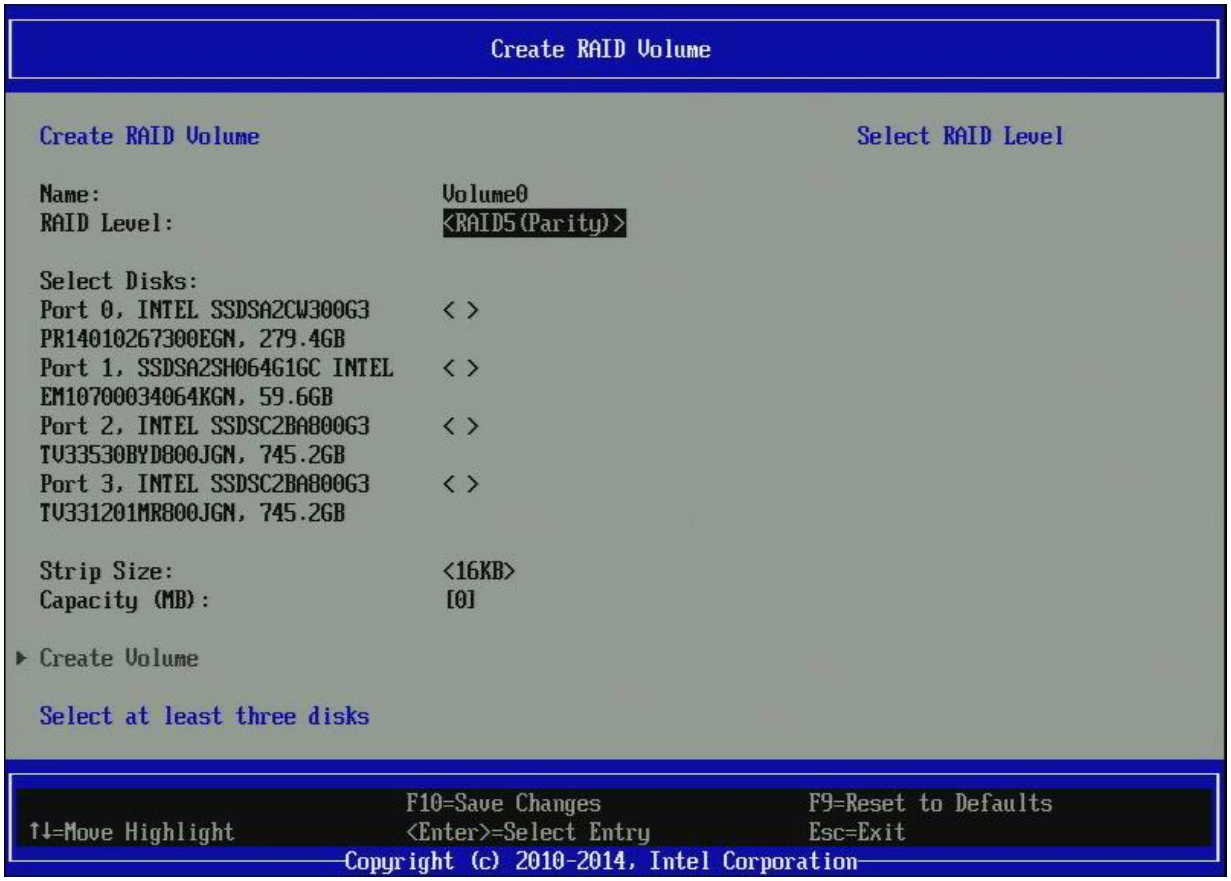
3. Section 3 gives information on how to navigate within the current page of the UEFI UI.
Note: this section is not implemented by the RSTe UEFI driver and is specific to the BIOS that was used for documentation purposes.

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

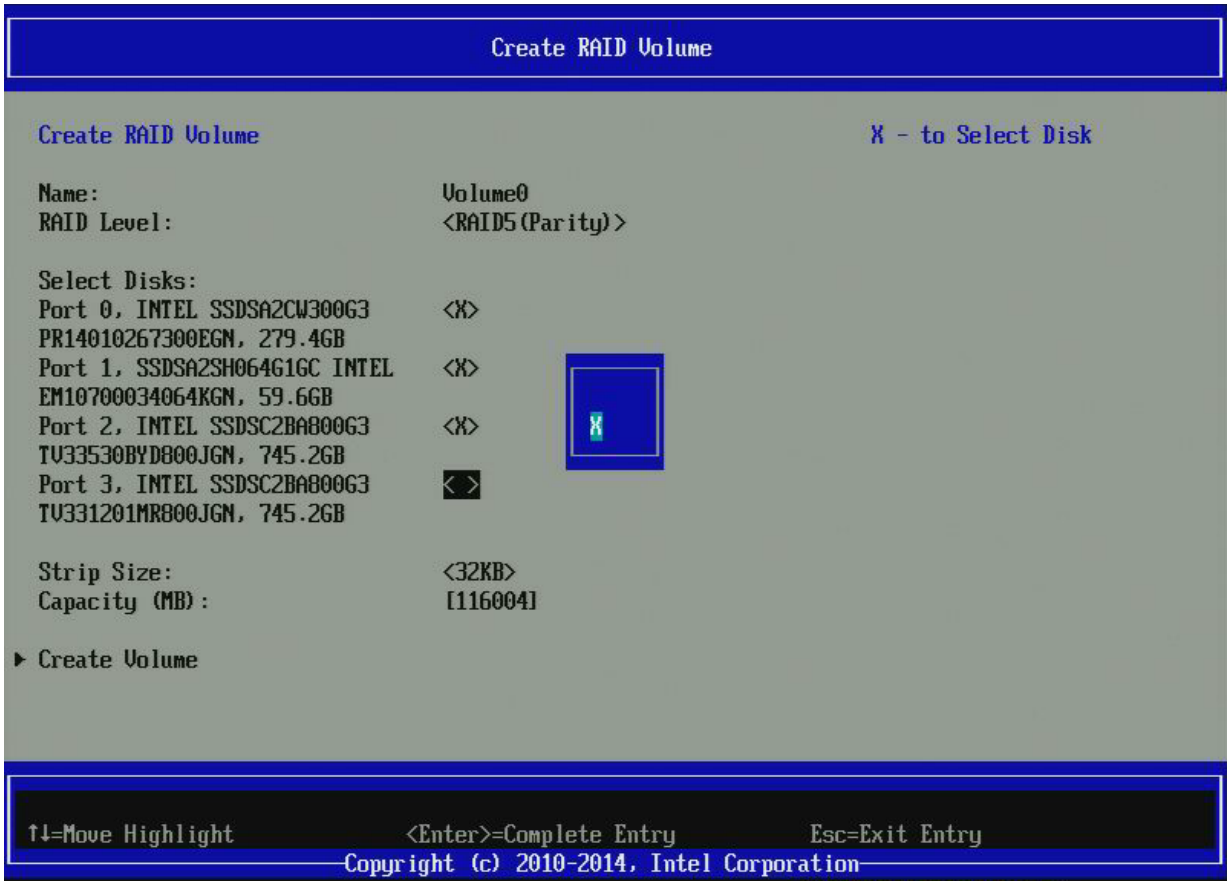
a. Enter the name you want to use for the RAID volume, then press <Enter>.



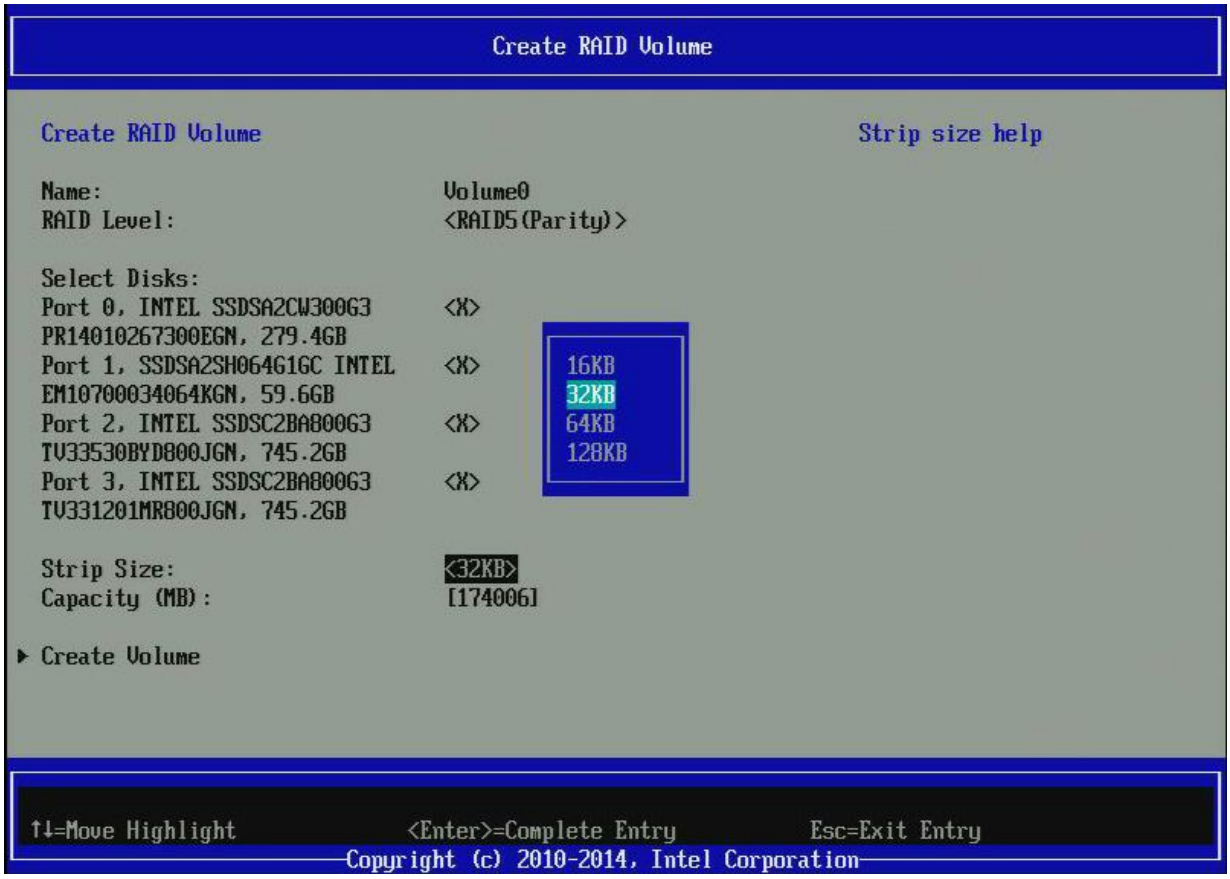
b. Scroll down to 'RAID Level" and press <Enter> to select a RAID level



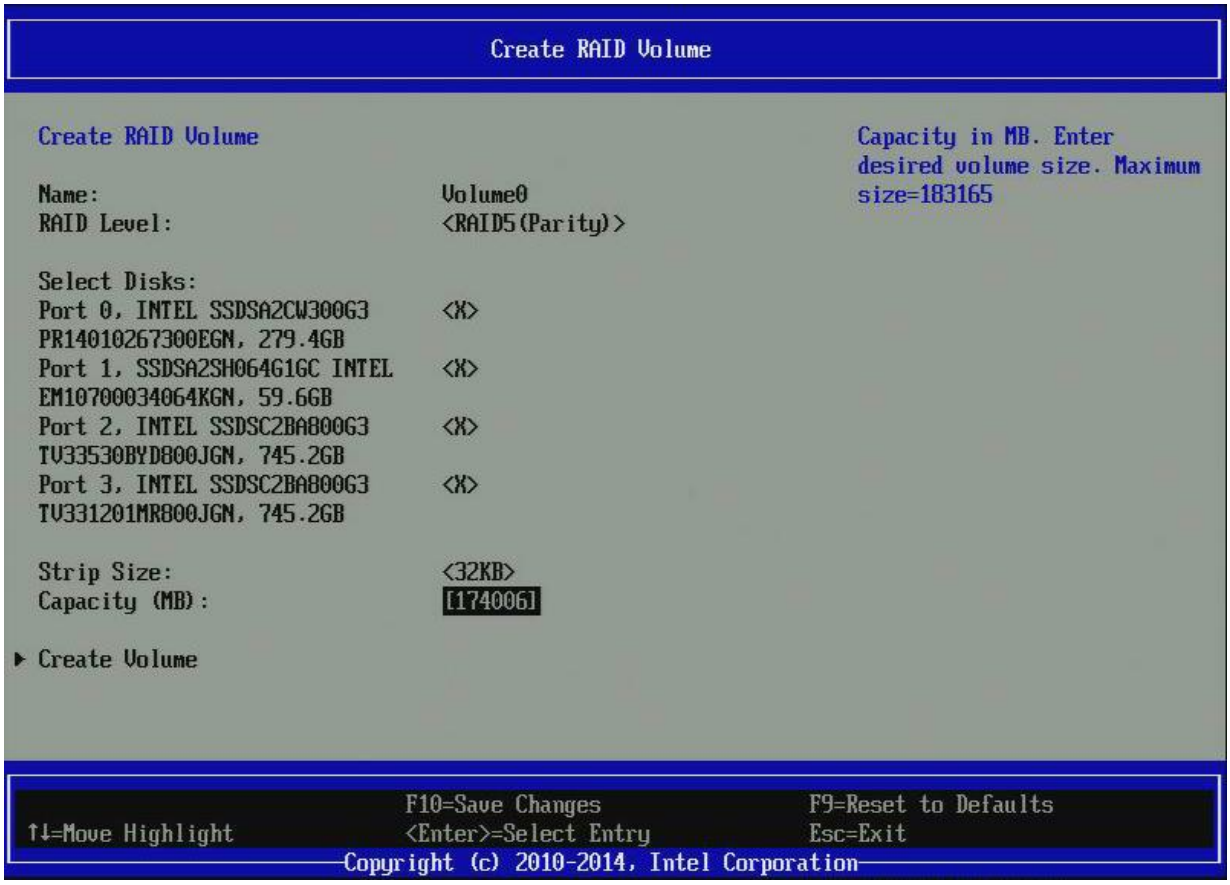
c. Scroll down to ‘Select Disks’ and at each disk that you wish to include in the RAID volume press <space bar>



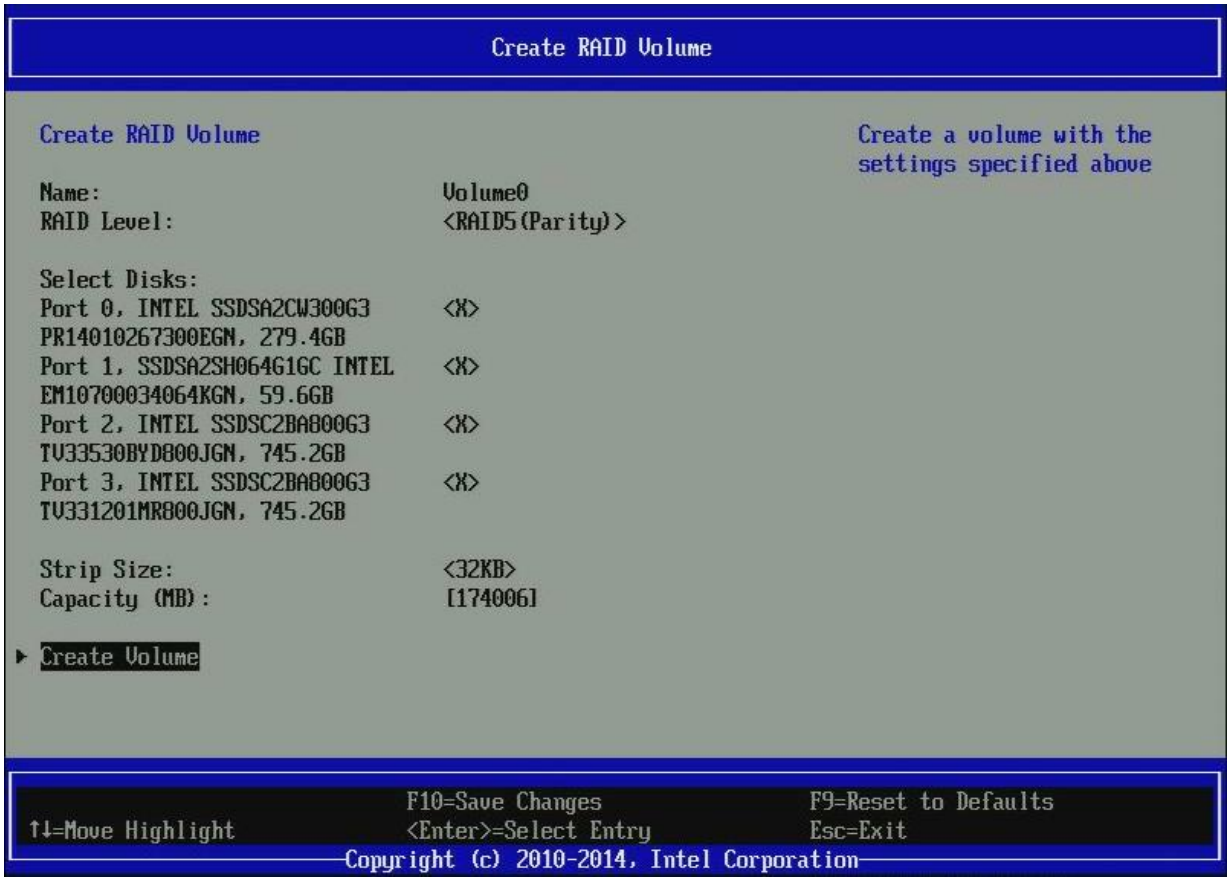
d. Next scroll down to 'Strip Size' and press <enter> to select a Strip size or continue if you wish to use the default strip size



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



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



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

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

7.4 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") in Windows environment with the following command line flags to create a RAID volume.

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 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.

Note: The Intel® RSTe 4.3.0.1219 release package contains the windows based tools to operate on the PCH. The Intel® RSTe 4.3.0.1681 release package contains the windows based tools to operate on Intel® NVMe SSDs. These tools are not interchangeable.



8 Define a Hot Spare

One to help automate RAID volume management it so be able to define an extra drive as a “Hot Spare”. This allows a RAID volume that become degraded to automantically replace the failed/missing drive with the drive marked as a “Hot Spare”. This can be can be accomplished in three ways.

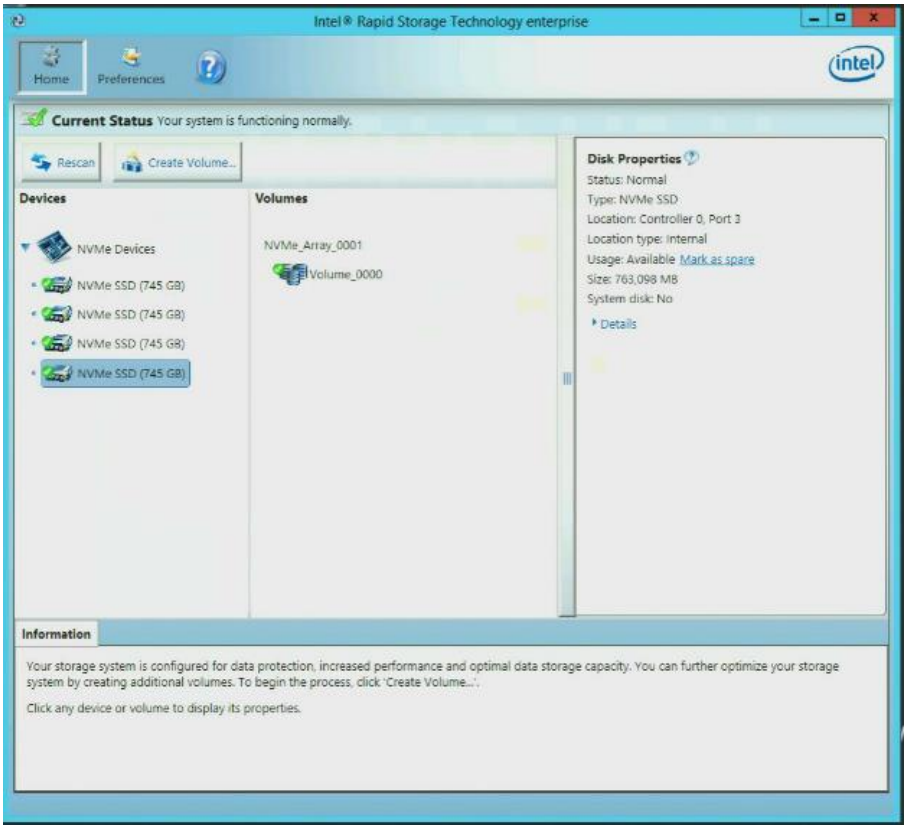
1. The Intel® RSTe GUI.
2. The Intel® RSTe 4.3.0.1219 provides Pre-OS environemt (LegacyOption ROM or UEFI HII) user interface support.
3. The Intel® RSTe CLI tools

8.1 Using the Windows User Interface Utility

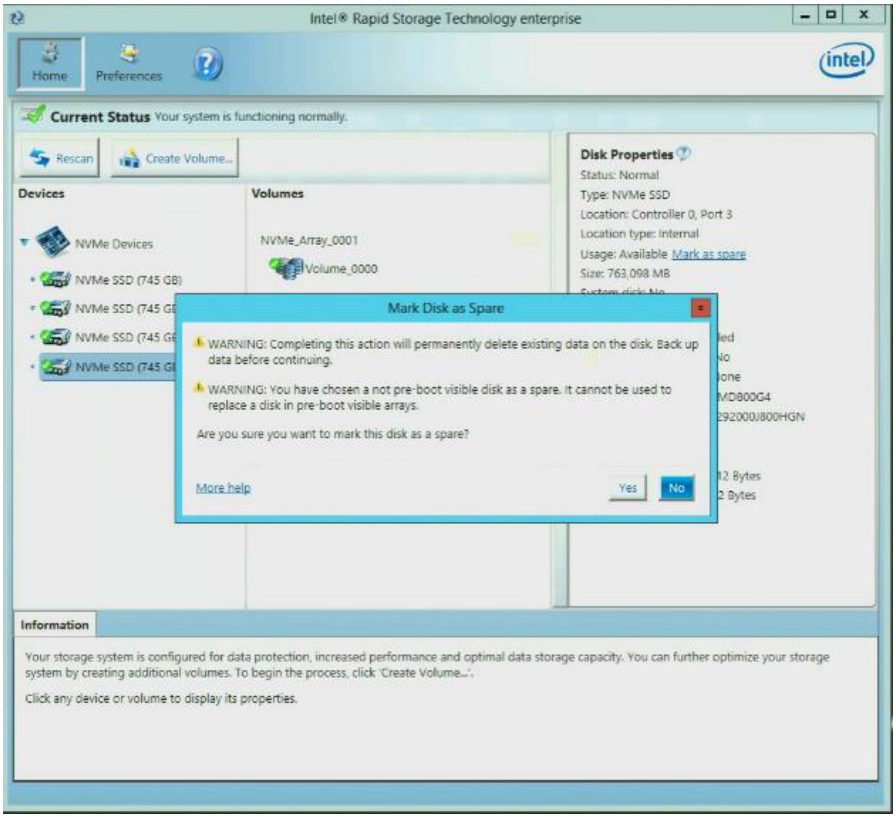
1. Run the Intel® RSTe GUI from the following Start menu link:

   Start→All Programs→Intel® Rapid Storage Technology enterprise →Intel Rapid Storage Technology enterprise GUI

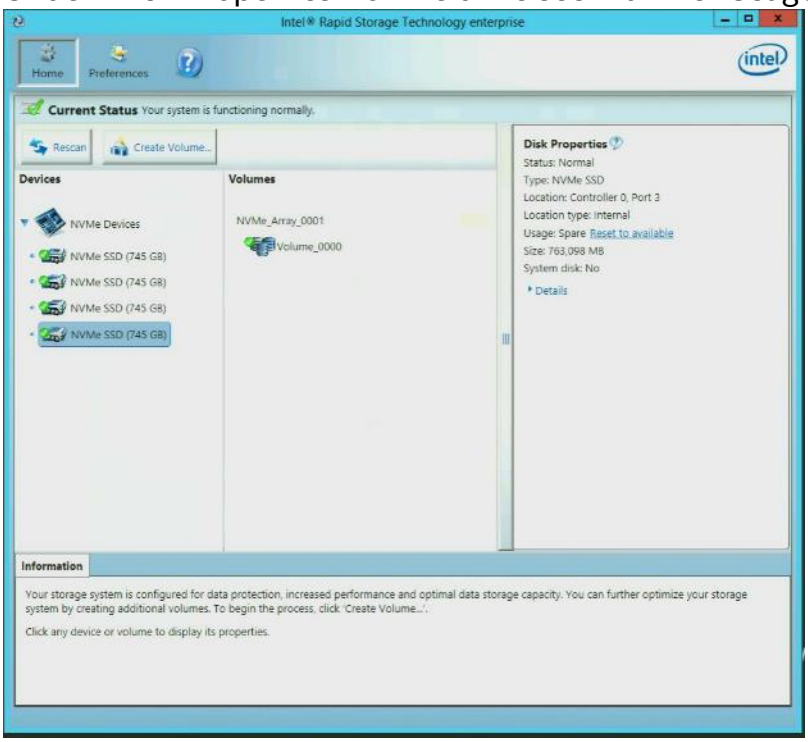
2. In the “Devices” pane, select the driver that is to marked as a spare drive. Then, from the Disck Properties pane, select “Mark as spare”



3. The warning window will appear to inform the user that all of data on the drive will be deleted. Click on “Yes”



4. Under “Disk Properties” for the drive see that the “Usage” is not set to “Spare”.





9 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 GUI. The second method is to use the Intel Rapid Storage Technology enterpris Pre-OS environemt (LegacyOption ROM or UEFI HII) user interface. The third way, used by OEMs only, uses the RAID command line configuration utility.

9.1 Using the Intel® RSTe GUI

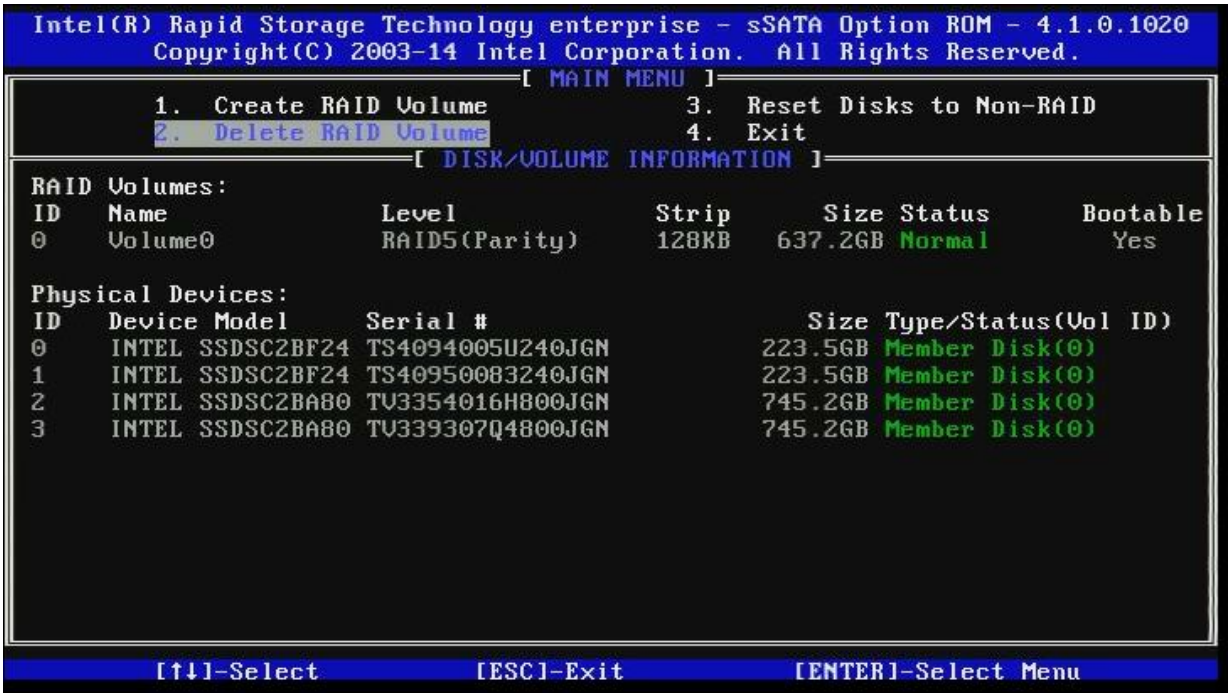
1. Run the Intel Rapid Storage Technology enterprise GUI from the following Start menu link:

   Start→All Programs→Intel® Rapid Storage Technology enterprise →Intel Rapid Storage Technology enterprise GUI

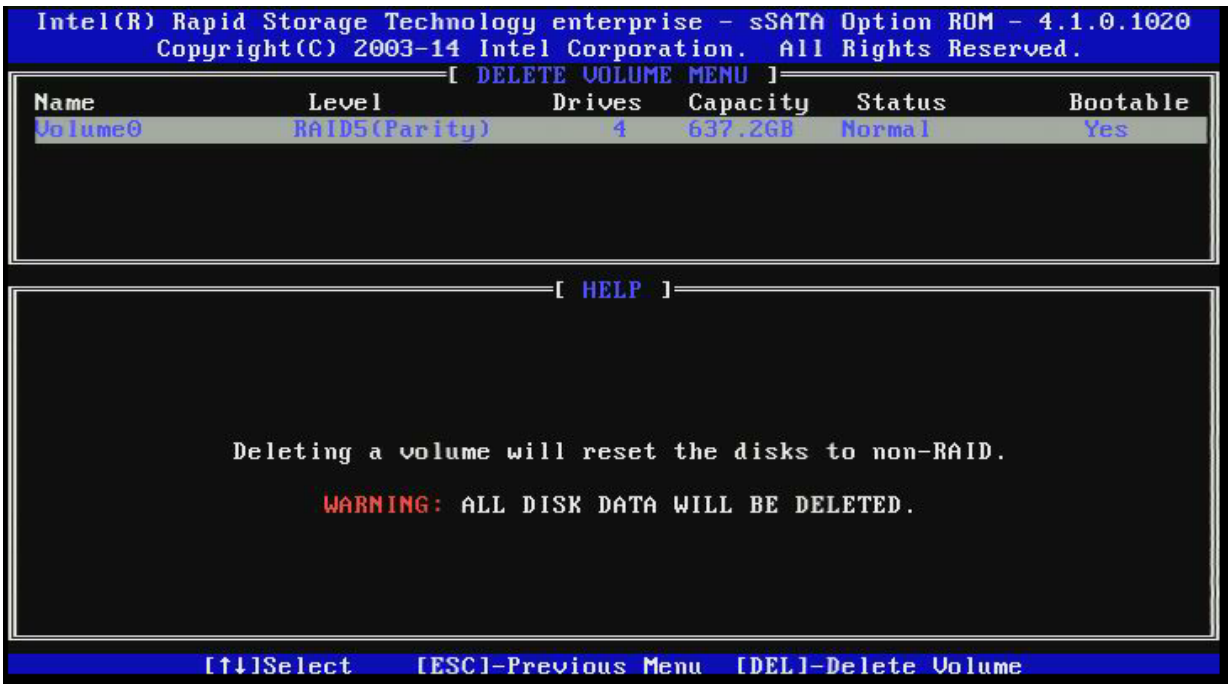
2. Under 'Home' Click on the volume you want to delete. The user will be presented with the volume properties on the right.
3. Click on 'Delete volume'
4. Review the warning message, and click 'Yes' to delete the volume.
5. The 'Status' page refreshes and displays the resulting available space in the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

9.2 Using the Intel® RSTe Legacy Option ROM User Interface

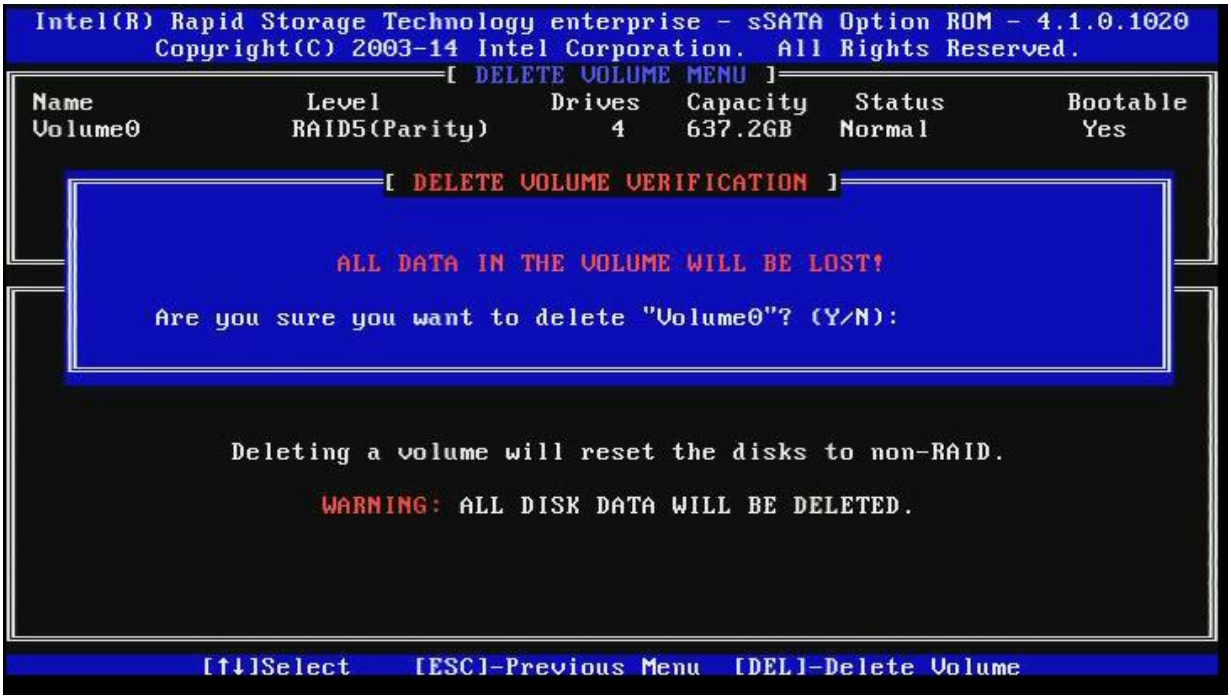
1. Upon re-boot, you will see the Intel Rapid Storage Technology enterprise option ROM status message on the screen – press CTRL-I to enter the option ROM user interface.
2. Within this UI, select option #2 'Delete RAID volume'.



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 affect 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).

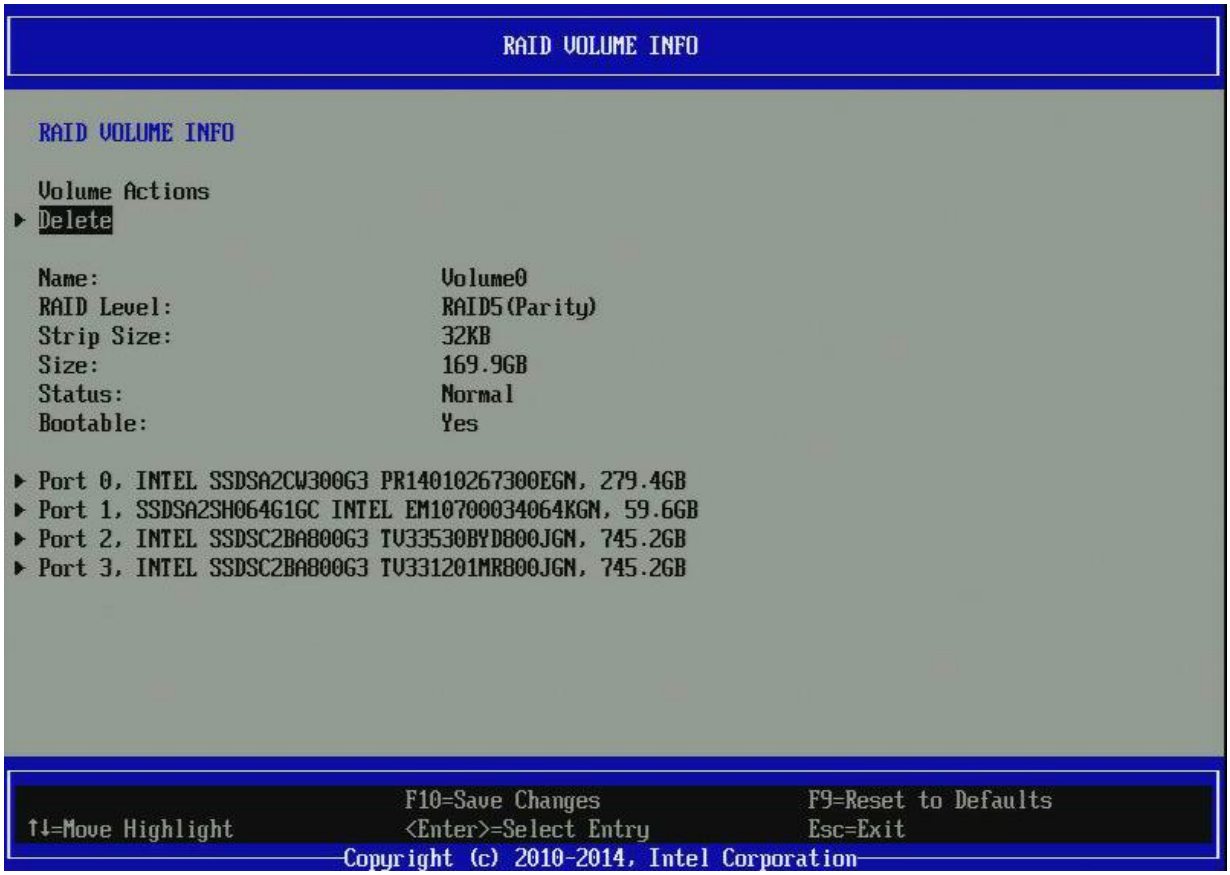
9.3 Using the Intel® RSTe UEFI HII User Interface

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.

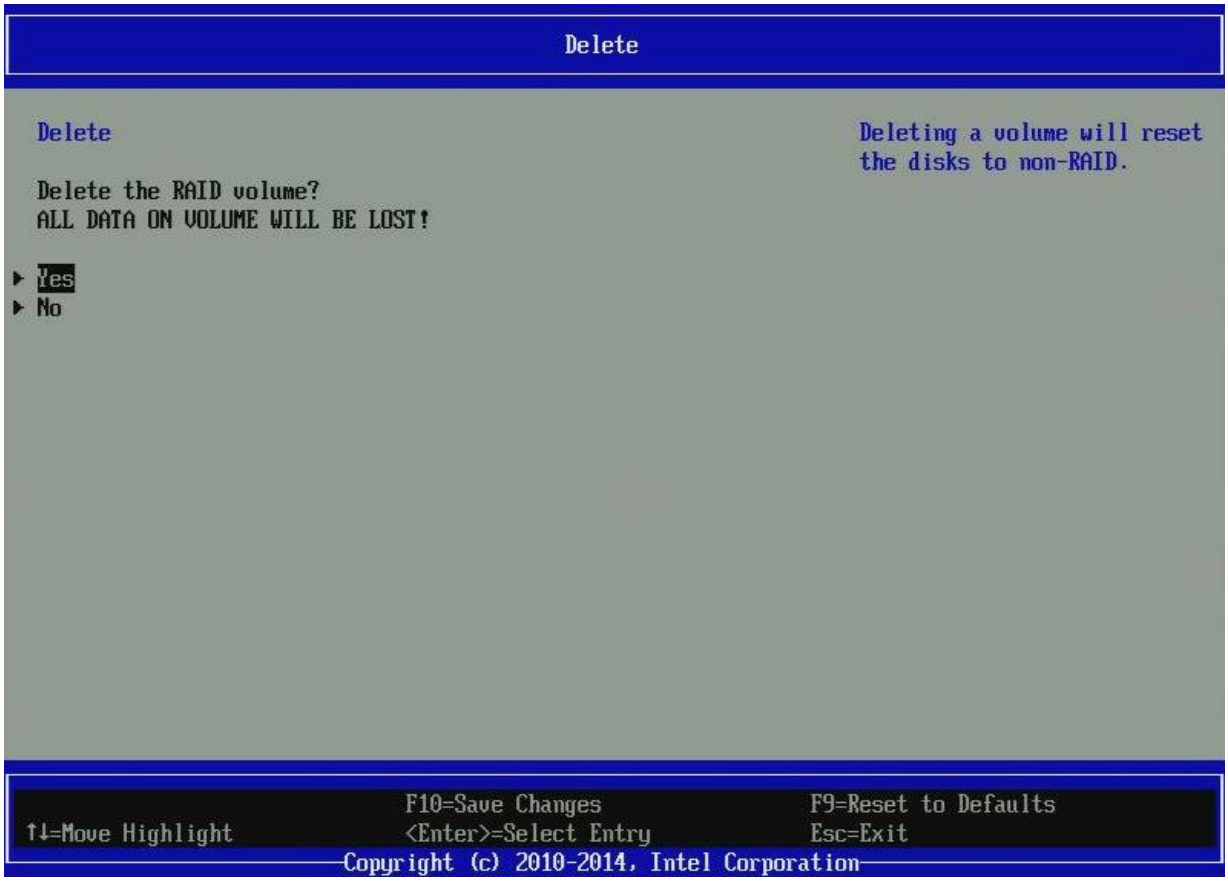
1. Upon re-boot, enter the system BIOS and select the Intel® RSTe menu for the UEFI user interface

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® RSTe UEFI UI.

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

To delete RAID volumes managed by Intel® RSTe PCH/SATA, 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



10 Common RAID Setup Procedures

10.1 Build a SATA RAID 0, 1, 5 or 10 System (OS Installation)

For platforms managed by Intel® RSTe PCH/SATA, these are the different RAID volume configurations that will support operating system installation. The first is the simple the striped RAID 0. The second is the mirrored RAID 1. The third is the striped with parity RAID 5 nd the final is the mirrored and striped across four drives for RAID 10. All OS bootable RAID member drives must be attached to the same controller of the PCH (spanning is not supported). 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® RSTe 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 on the SATA/sSATA controller. This setting may be different for each motherboard manufacturer. Consult the manufacturer's user manual if necessary. When done, exit Setup.

10.1.1 Using the Legacy OROM User Interface

3. 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® RSTe Legacy Option ROM user interface.
4. 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.

5. After this is done, exit the Intel® RSTe Legacy option ROM user interface by pressing the Esc key or Option #4.

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

7. Installation procedures as follows: Use the ‘load driver’ mechanism when prompted. Insert a USB key with the Intel® RSTe 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.

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

9. Install the Intel® RSTe 4.3.0.1219 software package obtained from the Intel VIP website. This will add the *Intel® RSTe* GUI that can be used to manage the RAID configuration.

10.1.2 Using the UEFI HII User Interface

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

4. 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.

5. After this is done, exit the Intel® RSTe menu HII user interface by pressing to save changes and the Esc key.

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

7. Installation procedures as follows: Use the 'load driver' mechanism when prompted. Insert a USB key with the Intel® RSTe 4.3.0.1219 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.

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

9. Install the Intel® RSTe *4.3.0.1219* software package obtained from the Intel VIP website. This will add the *Intel®* RSTe UI that can be used to manage the RAID configuration.



10.2 Build a "RAID Ready" System

For platforms managed by Inte® RSTe PCH/SATA 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® RSTe. Intel® RSTe 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® RSTe with Intel® RSTe Legacy OROM integrated into the BIOS and attach one SATA hard drive.

2. Enter System BIOS Setup; ensure that RAID mode is enabled on the SATA controller. 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® RSTe 4.3.0.1219 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® RSTe 4.3.0.1219 software package obtained from the Intel VIP website. This will add the *Intel®RSTe* GUI that can be used to manage the RAID configuration.

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

For platforms that are managed by Intel® RSTe PCH/SATA, if there is an existing "RAID Ready" system as defined in above section, 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 10.1 Build a SATA RAID 0, 1, 5 or 10 System (OS Installation). 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® RSTe 4.3.0.1219 software, if not already installed, using the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel® RSTe GUI and start menu links.
4. Open the Intel® RSTe GUI 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.

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

For system that are managed by Intel® RSTe PCH/SATA and/or Intel® RSTe NVMe, if there is an existing drive that contains important data, you can use Intel® RSTe to migrate that single drive into a 2 drive RAID 0 or RAID 1 volume. After the migration is completed, you will have a two hard drive RAID 0 volume where data is striped or a two drive RAID 1 volume where the data is mirrored across the two drives. For the drives attached to the SATA Controller, the BIOS must have the controller configured for RAID mode and you must have the Intel® RSTe software installed.

1. Note the serial number of the data 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 drive to the available SATA port (PCH/SATA) or PCI exress slot managed by the CPU (NVMe).
3. Boot to Windows, install the Intel® RSTe software, if not already installed, using the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel® RSTe GUI and start menu links.
4. Open the Intel® RSTe GUI from the Start Menu.
5. Follow steps 4 to 7 in section 10.3 Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System.

10.5 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 please Intel® RSTe 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 drives for the RAID level have to be met.

Please follow the procedure illustrated below

1. Start the Intel Rapid Storage Technology enterprise GUI application:



Start Menu ->All Programs -> Intel Rapid Storage Technology enterprise-> Intel Rapid Storage Technology enterprise GUI

2. From the Volumes pane, click the array or volume to which you want to modify. The volume properties now display on the right.
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.



11 Intel® RSTe NVMe LED Mangement

Intel® RSTe NVMe support provides a mechanism to manage the status LEDs in an NVMe drive backplane

11.1 An Overview of LED Managment Support in PCIe Enclosures

In order to manage status LEDs in PCIe NVMe Enclosures, several SW/HW combinations have been considered but none of them have been widely implemented yet.

Unlike the standard PCH architecture, where there is a single controller independent of the drives attached, every NVMe SSDs has its own built-in storage controller. This integration of the NVMe controller with each drive forces the LED management functionality reside outside of the NVMe controller so that LED management can be maintained even when one of the NVMe drives fails. As a result, LED management must stay outside of the NVMe subsystem. While NVMe management standards descript how to get the status from the NVMe subsystem, translating this status information into a visual indication is not and is therefore vendor specific.

To support LED management for in a system running Intel® RSTe NVMe, a vendor specific plug-in must be created. This vendor specific plug-in provide the interface between Intel® RSTe NVMe and the NVMe enclosure.

11.2 LED Blinking Support in RSTe 4.3

In order for Intel® RSTe NVMe to deliver support for managing LEDs on all enclosures, a plug-in based architecture has been implemented. It allows a platform vendor to write a small portion of code, compile it into a DLL and put it into the directory accessed by RSTe. This plug-in provides the communication with the NVMe enclosure.

11.2.1 LED Management Information Flow

The Intel® RSTe System Service provides the information about RAID member disks states. The information is passed to the plug-in DLL in the form of a callback.

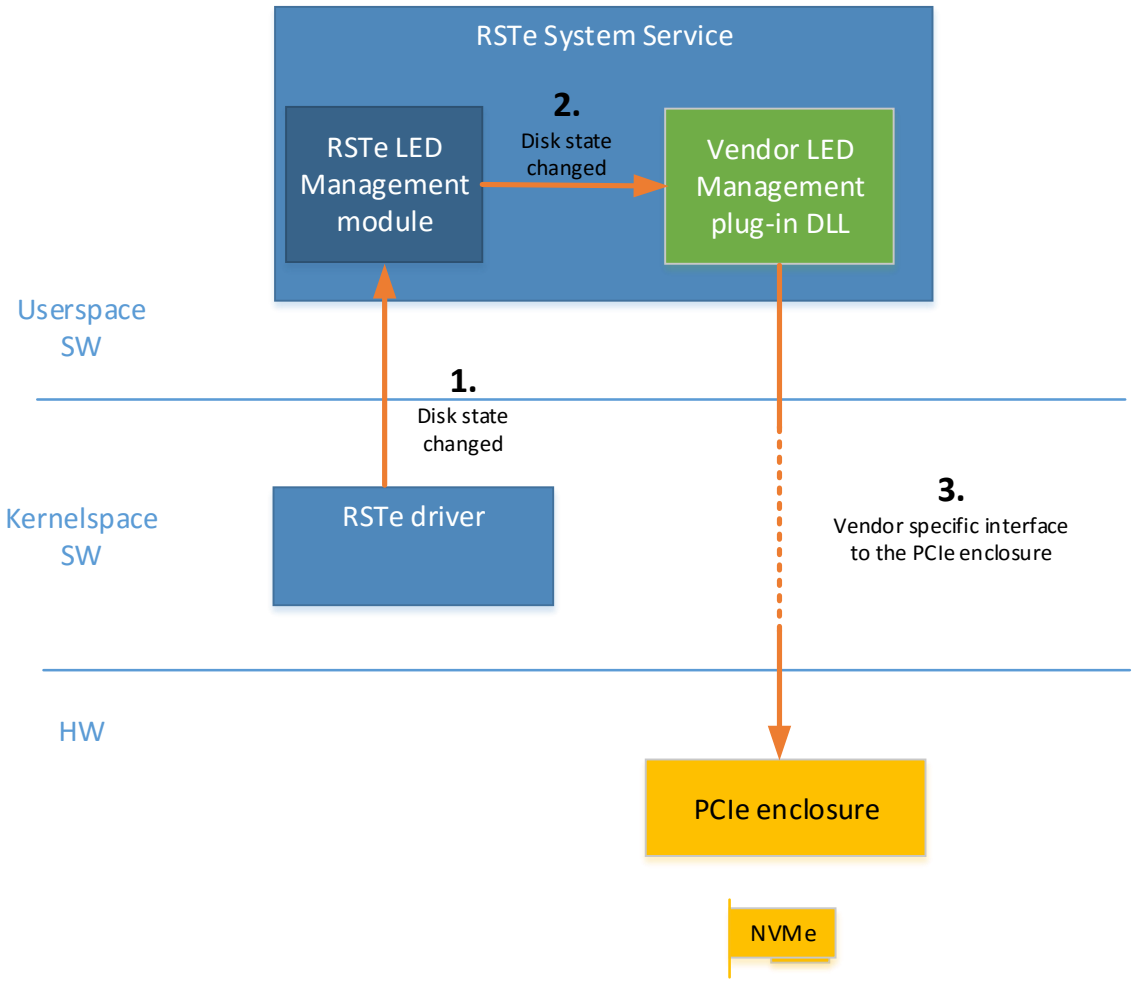


Figure 1: The flow of disk state changes resulting in LED Management

11.3 LED Blinking Plug-in

11.3.1 The interface

The interface which has to be implemented by the Vendor plug-in:

```
public interface LedEnclosurePlugin
{
  string Name { get; }
  bool Blink(
    BlinkState State,
```



```
    byte Bus,
    byte Device,
    byte Function
    );
}
```

11.3.2 'Name' Property

Defines the name of the plug-in. This property is used for debug purposes.

11.3.3 'Blink' Method

Arguments:

Argument name	Description
BlinkState	Enum defining the disk state, in regard to LED blinking (See BlinkState values mapping to SFF-8489 specification table)
Bus	Bus number of PCI device (part of BDF)
Device	Device number of PCI device (part of BDF)
Function	Function number of PCI device (part of BDF)

BlinkState values mapping to SFF-8489 specification:

SFF specification value	BlinState value
Locate (identify)	eStateLocate eStateLocateOff*
Fail	eStateFailed
Rebuild	eStateRebuild
PFA	eStatePFA
Hotspare	eStateHotSpare



In A Critical Array	eStateCriticalArray
In A Failed Array	eStateFailedArray

*eStateLocateOff is sent after LedLocateTimeout expires (see 11.4)

Return values:

Value	Description
True	The blinking related information was successfully passed to the enclosure
False	The Blink function failed

11.4 Plug-in Deployment

When Intel® RSTe NVMe System Service is started, the RSTe LED Mangement Module searches for the Vendor plug-in DLL. The search is performed in the directory defined in system registry in the following location:

HKLM\SOFTWARE\Intel\RSTe (for 32 bit OS)

or

HKLM\SOFTWARE\Wow6432Node\Intel\RSTe (for 64 bit OS)

1. Add two registry entries in HKLM\SOFTWARE\Intel\RSTe or HKLM\SOFTWARE\Wow6432Node\Intel\RSTe (depending on whether it is 32 bit or 64 bit architecture):
    a. First entry defining full plugin path (without name i.e: C:\Program Files (x86)\Intel\Intel(R) Rapid Storage Technology enterprise\ Enclosure LED plugins):
        i. name: EnclosurePluginPath
        ii. type: String value
    b. Second entry defining led locate timeout value:
        i. name: LedLocateTimeout
        ii. type: dword
2. Copy plugin file to path specified in point 1a (i.e: put file SampleEnclosurePlugin.dll in C:\Program Files (x86)\Intel\Intel(R) Rapid Storage Technology enterprise\LedPlugin):



a. Plugin must implement interface LedEnclosurePlugin from library IAStorEnclosurePlugin.dll

The Intel® RSTe installer default path is:

C:\Program Files (x86)\Intel\Intel(R) Rapid Storage Technology enterprise\Enclosure LED plugins

11.4.1 Potential Deployment Issues

1. If the plug-in location contains more than one plug-in, only one will be loaded, but there is no guarantee which one.
2. If the plug-in location contains DLLs, which do not meet the requirements of a plug-in, they will be ignored.
3. If the plug-in code performs an illegal operation (e.g. illegal memory access), RSTe components do not crash, but the LED blinking functionality may stop working until the system restart.
4. If the DLL is not loaded properly, the “Locate LED” option will not appear in Disk Properties in the Intel® RSTe GUI.

11.5 Example Enclousre Plugin Code

These is example source code to help start the DLL plugin effort included in this package.

NOTE: This sample code an not a fully functional solution. It does contain the necessary files to begine the process of building the customer DLL to support NVMe LED management.

11.6 Example of the Build Envronment

This is an example of a successful DLL software build environement:

- Microsoft Visual Studio Professional 2013
  Version 12.0.31101.00 Update 4
  Microsoft .NET Framework
  Version 4.5.51650
- Installed Version: Professional
- LightSwitch for Visual Studio 2013 06177-004-0447006-02393
  Microsoft LightSwitch for Visual Studio 2013



- Team Explorer for Visual Studio 2013 06177-004-0447006-02393
  Microsoft Team Explorer for Visual Studio 2013
- Visual Basic 2013 06177-004-0447006-02393
  Microsoft Visual Basic 2013
- Visual C# 2013 06177-004-0447006-02393
  Microsoft Visual C# 2013
- Visual C++ 2013 06177-004-0447006-02393
  Microsoft Visual C++ 2013
- Visual F# 2013 06177-004-0447006-02393
  Microsoft Visual F# 2013
- Visual Studio 2013 Code Analysis Spell Checker 06177-004-0447006-02393
  Microsoft® Visual Studio® 2013 Code Analysis Spell Checker
- Portions of International CorrectSpell™ spelling correction system © 1993 by Lernout & Hauspie Speech Products N.V. All rights reserved.
- The American Heritage® Dictionary of the English Language, Third Edition Copyright © 1992 Houghton Mifflin Company. Electronic version licensed from Lernout & Hauspie Speech Products N.V. All rights reserved.
- Windows Phone SDK 8.0 - ENU 06177-004-0447006-02393
  Windows Phone SDK 8.0 – ENU
- Application Insights Tools for Visual Studio Package 1.0
  Application Insights Tools for Visual Studio
- ASP.NET and Web Tools 12.4.51016.0
  Microsoft Web Developer Tools contains the following components:
  - Support for creating and opening ASP.NET web projects
    Browser Link: A communication channel between Visual Studio and browsers
    Editor extensions for HTML, CSS, and JavaScript
    Page Inspector: Inspection tool for ASP.NET web projects
    Scaffolding: A framework for building and running code generators
    Server Explorer extensions for Microsoft Azure Websites
    Web publishing: Extensions for publishing ASP.NET web projects to hosting providers, on-premises servers, or Microsoft Azure
- ASP.NET Web Frameworks and Tools 2012.2 4.1.21001.0
  For additional information, visit
  http://go.microsoft.com/fwlink/?LinkID=309563
- ASP.NET Web Frameworks and Tools 2013 5.2.21010.0
  For additional information, visit http://www.asp.net/



- Common Azure Tools 1.3
  Provides common services for use by Azure Mobile Services and Microsoft Azure Tools.
- Debugging Tools for Windows 6.3.9600.17029
  Integrates the Windows Debugger functionality (http://go.microsoft.com/fwlink/?linkid=223405) in Visual Studio.
- Microsoft Advertising SDK for Windows Phone
  Microsoft Advertising SDK for Windows Phone Build
- Microsoft Azure Mobile Services Tools 1.3
  Microsoft Azure Mobile Services Tools
- NuGet Package Manager 2.8.50926.663
  NuGet Package Manager in Visual Studio. For more information about NuGet, visit http://docs.nuget.org/.
- Office Developer Tools for Visual Studio 2013 ENU 12.0.30626
  Microsoft Office Developer Tools for Visual Studio 2013 ENU
- PowerShell Tools 1.3
  Provides file classification services using PowerShell
- PreEmptive Analytics Visualizer 1.2
  Microsoft Visual Studio extension to visualize aggregated summaries from the PreEmptive Analytics product.
- SQL Server Data Tools 12.0.41012.0
  Microsoft SQL Server Data Tools
- Windows Driver Kit 6.3.9600.17246
  Headers, libraries, and tools needed to develop, debug, and test Windows drivers (msdn.microsoft.com/en-us/windows/hardware/gg487428.aspx)
- Windows Phone 8.1 SDK Integration 1.0
  This package integrates the tools for the Windows Phone 8.1 SDK into the menus and controls of Visual Studio.
- Workflow Manager Tools 1.0 1.0
  This package contains the necessary Visual Studio integration components for Workflow Manager.



12 Intel® RSTe PCH/SATA SGPIO

Intel® RSTe PCH/SATA provides support for enclosure LED management via SGPIO (Serial General Purpose I/O Bus) and complies with SFF-8485 as well as SFF-8489. To enable this feature, the SATA and sSATA controllers must be configured for RAID mode, otherwise SGPIO LED management is not supported. This feature is only supported with SGPIO compliant backplanes and enclosures.

Intel® RSTe supports SGPIO LED management by default. This is contolled by the following registry key:

HKLM, System\CurrentControlSet\Services\iaStorA\Parameters\Device, LED_Supported, %REG_TYPE_DWORD%, 0x00000001

To disable SGPIO support set this DWORD to 0.

Intel® RSTe supports both passive and active backplanes. An active backpane contains an SGPIO controller that manages the LED blink patterns. RSTe default behavior assumes that the SGPIO backplane is active. For platforms that contain passive backplanes the following registry key must be added:

HKR, Parameters\Device, BackplaneIsPassive, %REG_TYPE_DWORD%, 0x00000001

The following are the supported patterns, which follow the IBPI patterns described in SFF-8489:

- Drive not present
- Drive present, no Activity
- Drive present, Activity
- Locate (Identify)
- Fail
- Rebuild

The Bit layout is as follows:

Pattern	Bit0 (Activity)	Bit1 (Locate)	Bit2 (Fail)
Drive not present	X	X	X

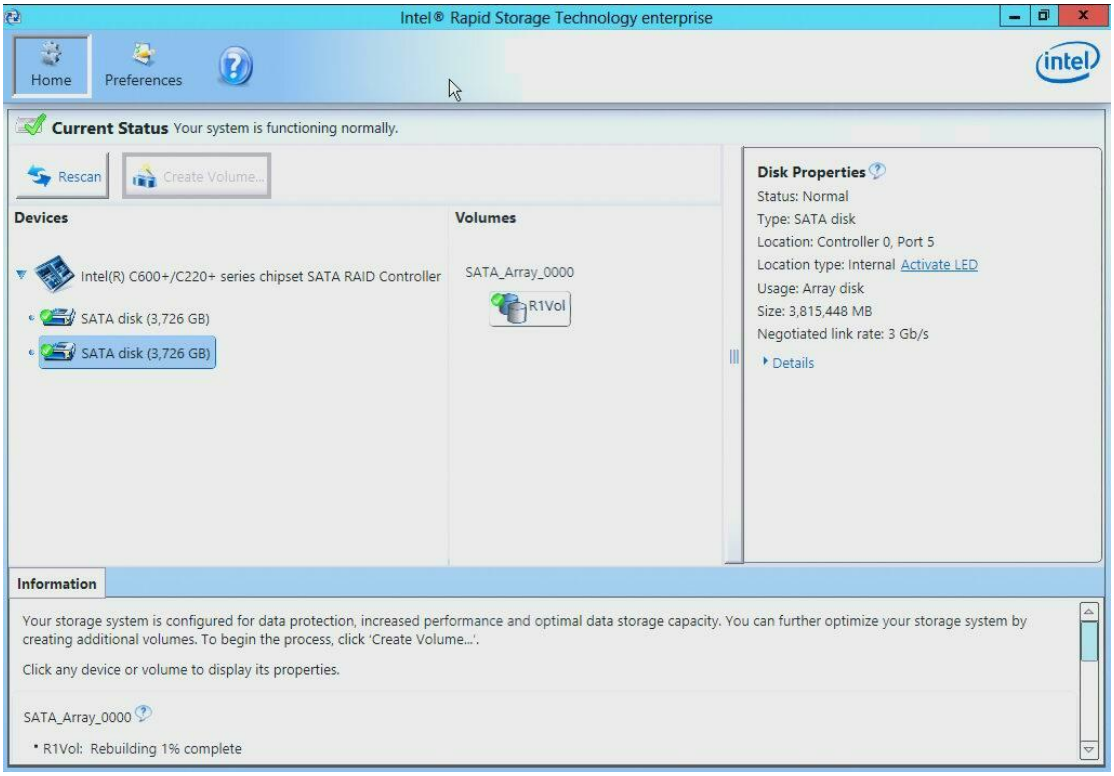


Drive present, no Activity	0	X	X
Drive present, Activity	1	X	X
Locate (Identify)	X	1	0
Fail	X	0	1
Rebuild	X	1	1

X=discard

12.1 Activate LED

The only SGPIO LED management pattern that is initiated by the user is “Locate”. This is initiated manually from theIntel® RSTe GUI. The “Activate LED” command is activated from the Disk Properties pane as shown below:





12.2 FAIL LED

When a disk encounters a failure condition that the causes the Intel® Rapid Storage Technology enterprise driver to mark it as FAILED, it will illuminate the SGPIO LED associated with the drive so that it is “on”.

The user will be able to go into the Intel® Rapid Storage Technology enterprise application to see which drive in the array has been marked as “FAILED” and will be able to manage accordingly.

12.3 SMART LED

When a disk encounters a SMART event condition causes that the Intel® Rapid Storage Technology enterprise driver to mark it as such, it will illuminate the SGPIO LED associated with the drive so that it is “on”.

The user will be able to go into the Intel® Rapid Storage Technology enterprise application to see which drive in the array has been marked as “SMART Event” and will be able to manage accordingly.

12.4 Rebuild LED

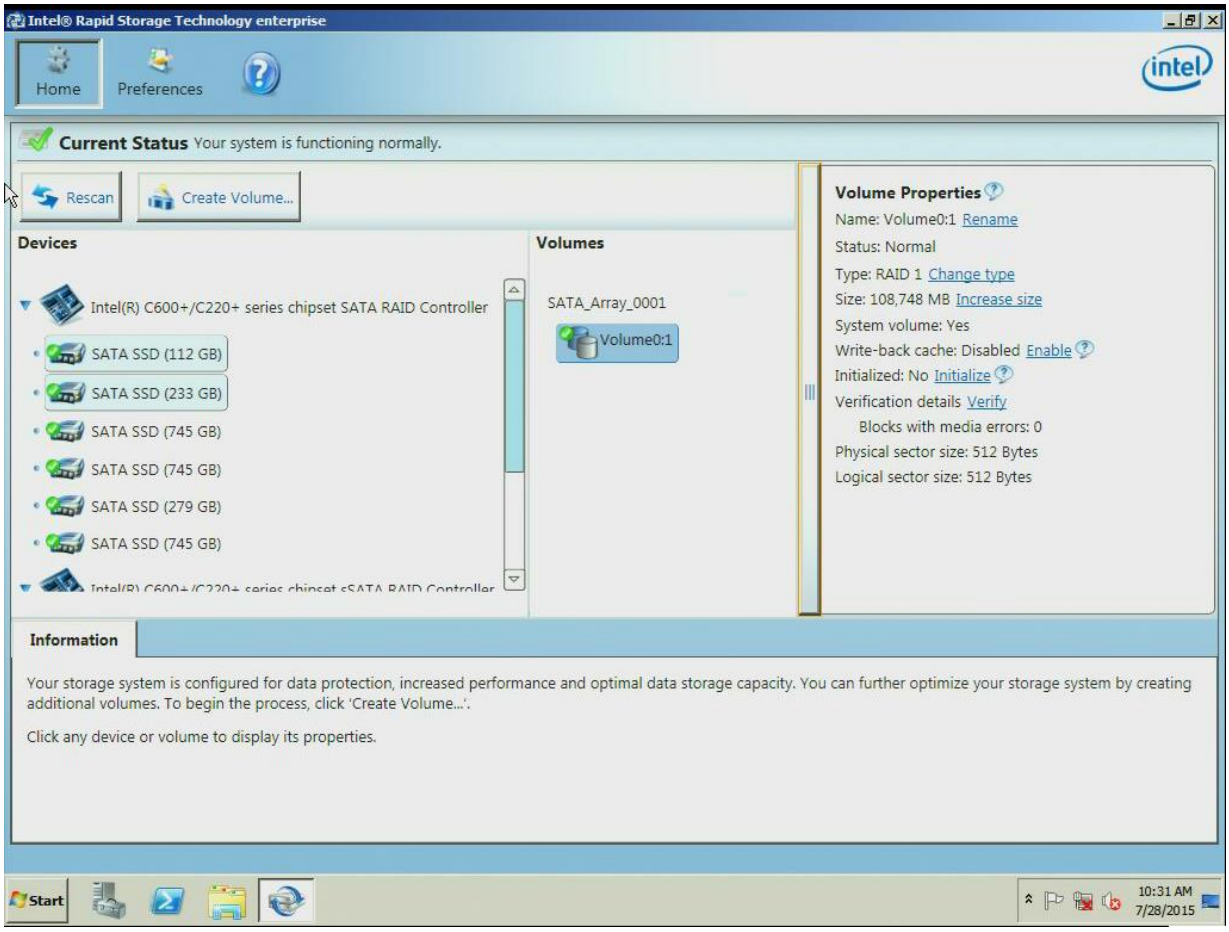
When a condition is encountered that results in a RAID Volume Rebuild, the Intel® Rapid Storage Technology enterprise driver will illuminate the SGPIO LED, associated with the drive being rebuilt so that it will “blink”. Once the rebuild has completed, the Intel® Rapid Storage Technology enterprise driver will set the SGPIO LED so that it will stop blinking (i.e. turn it “off”).

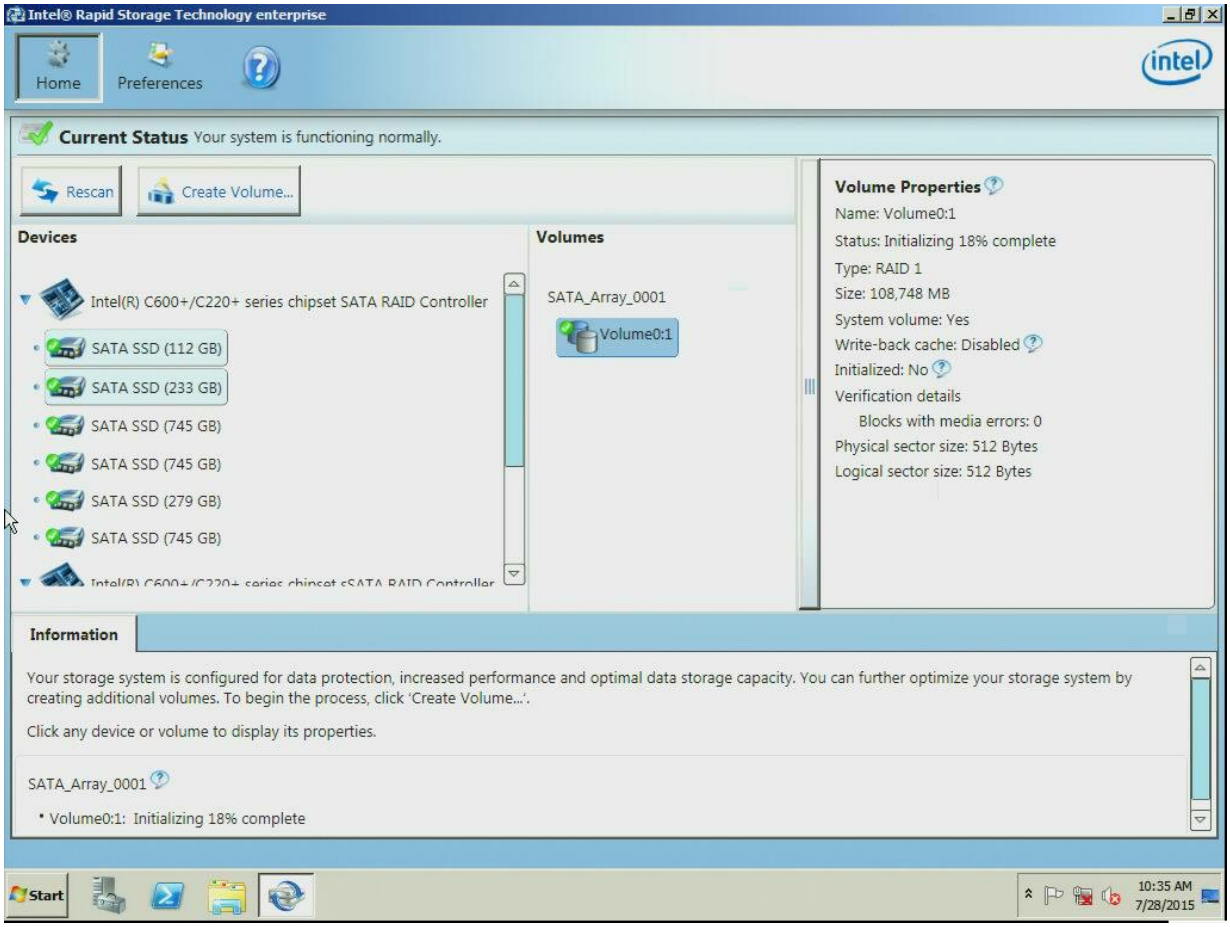
12.5 Migrate LED

When a condition is encountered that results in a RAID Volume Migration, the Intel® Rapid Storage Technology enterprise driver will illuminate the SGPIO LED, associated with the drive being migrated so that it will “blink”. Once the migrations process has completed, the Intel® Rapid Storage Technology enterprise driver will set the SGPIO LED so that it will stop blinking (i.e. turn it “off”).

13 Intel® RSTe RAID Volume Initialization

The RAID volume initialization process ensures that the entire RAID volume has all of its RAID metadata writing. Under most conditions the initialization occurs int the background while data is being written to the volume To ensure that all of RAID volume is properly setup, it is recommended that the volume initialization be manually initiated. This is done through the Intel® RSTe GUI. Under the “Volume” pane, select the RAID volume to be initialized. Then under the “Volume Properties” pane select “Initialize”





For RAID 5 volumes that contain 5 or more drives, the initialization process will begin automantically. The is done to inprove the operational performance of the RAID 5 voluem with lots of drives.



14 RAID Volume Verify and Repair Feature

The RAID volume verification and repair feature identifies any RAID metadata 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. Those inconsistencies are automatically repaired. The Verify and Repair process operates on the entire RAID volume, and can only be run on one RAID volume at a time.

14.1 Starting Verify and Repair Process

The following outlines the different methods that the Verify and Repair process is either automatically started or manually initiated.

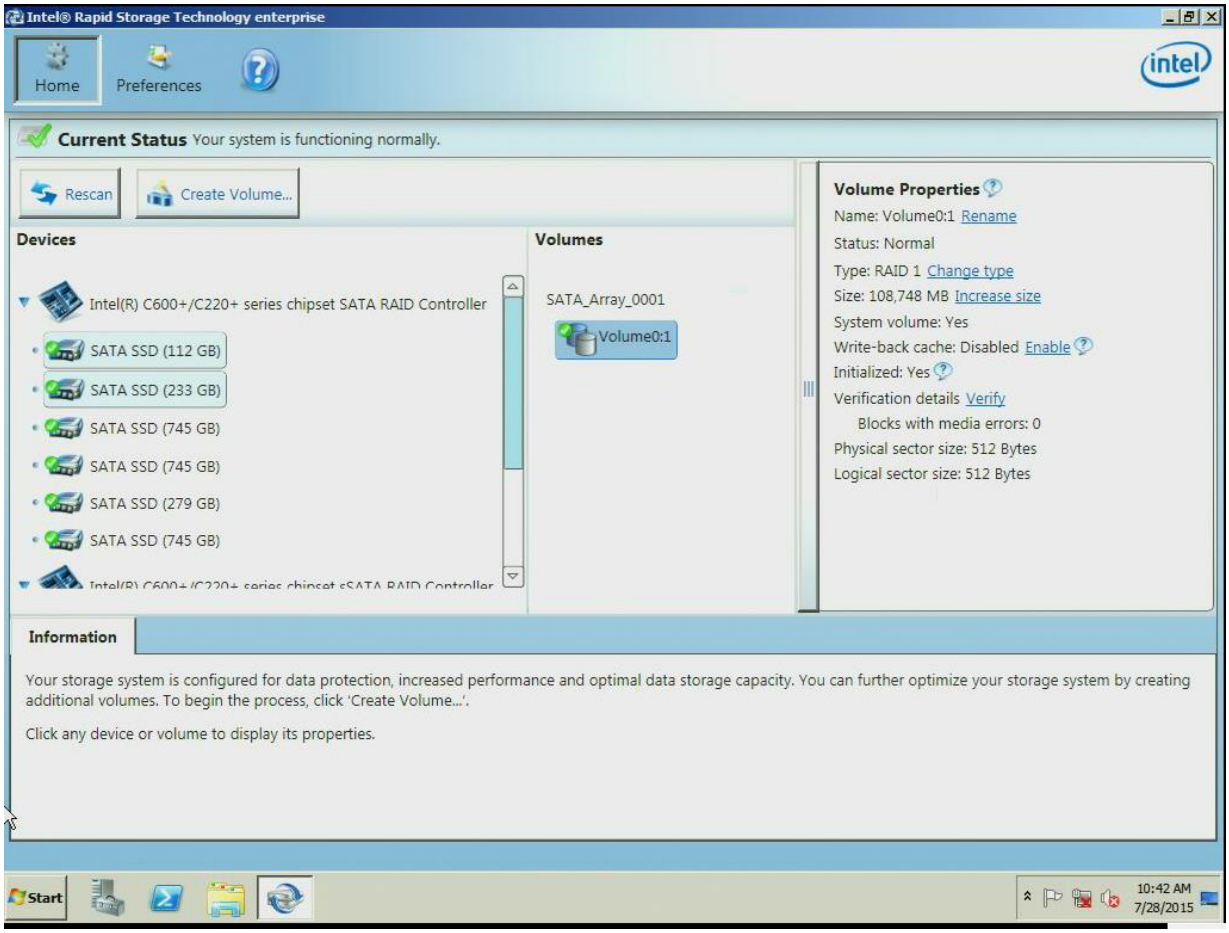
14.1.1 Dirty Shutdown

The Verify and Repair process will automatically begin after the platform encounters a unplanned/unexpected power outage (Diry Shutdown).

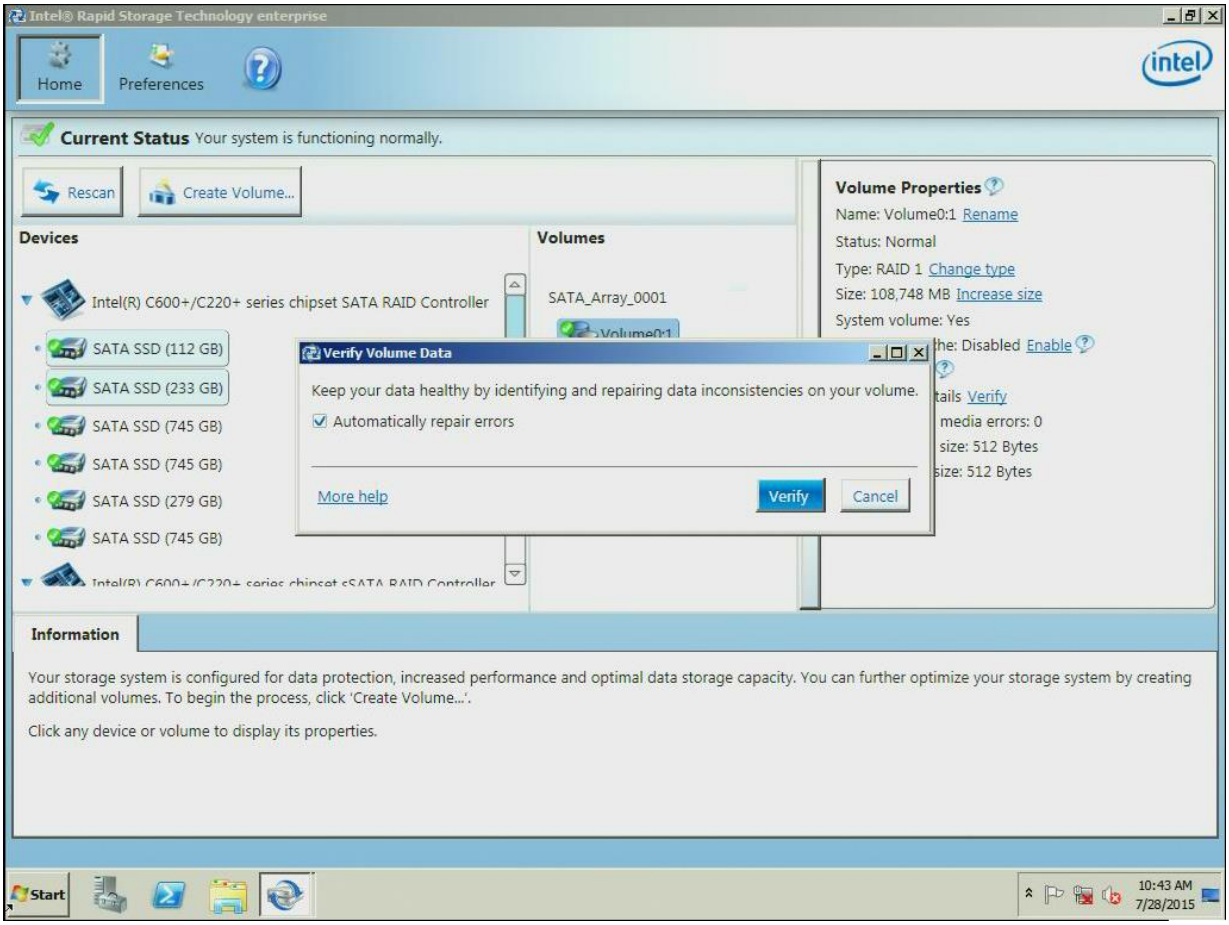
14.1.2 Manual Initiation

Intel® RSTe requires that a RAID volume be initialized prior manually initializing the verify and repair process. Follow the below steps to start RAID volume data verification and repair process.

1. From the Intel® RSTe GUI, in the “Volumes” pane, select the Volume to run Verify and Repair on. In the “Volume Properties” pane, click on “Verify” by “Verification details”.

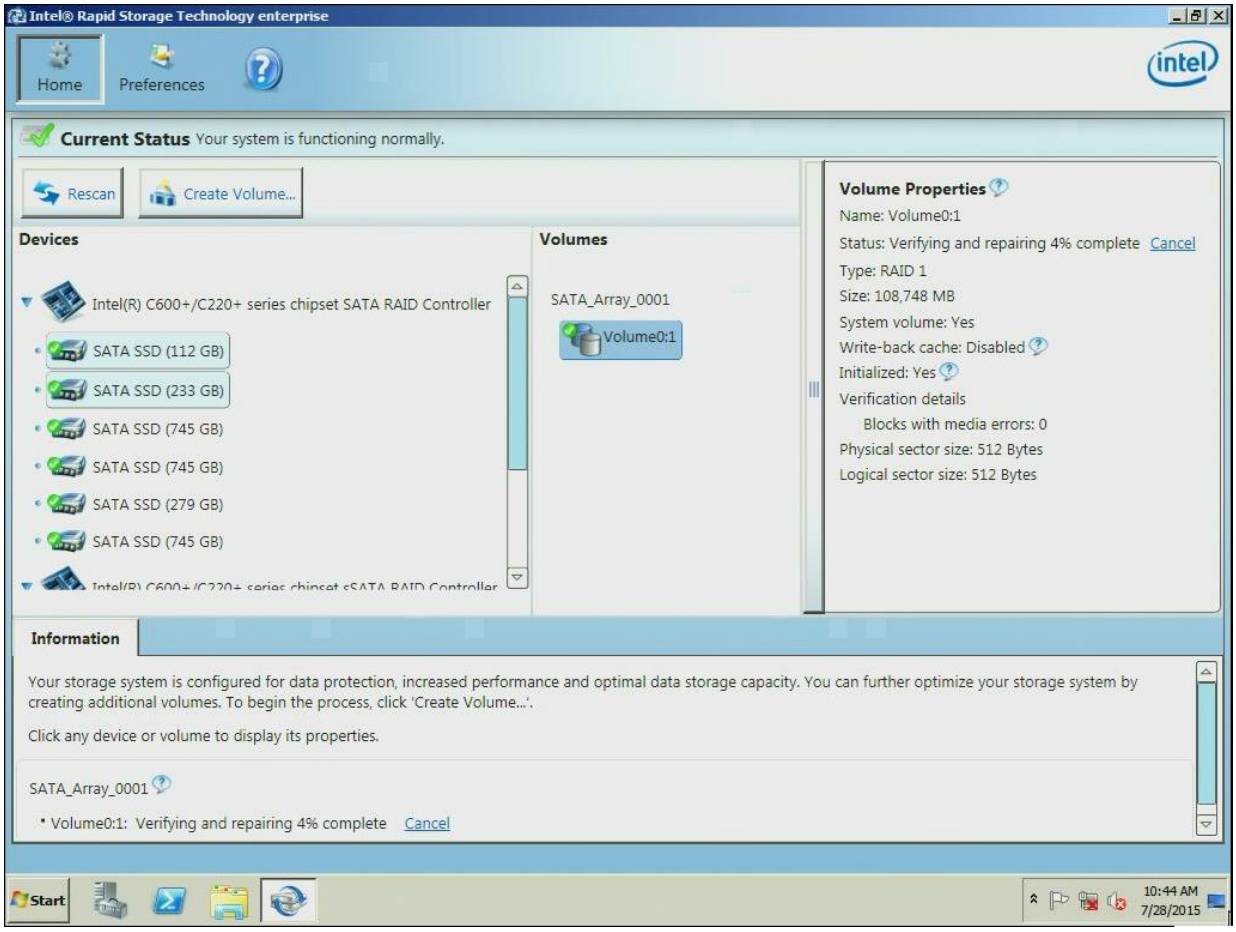


2. Click on 'Advanced' and then Click on 'Verify'



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.



15 Pre-OS Installation of the Intel® RSTe Driver

The Intel® RSTe PCH/SATA 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® RSTe driver be installed and loaded prior to the OS installation. On those OS versions the Intel® RSTe driver can be loaded post OS installation. The Intel® RSTe AHCI driver can be installed over Window's native AHCI driver.

15.1 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® RSTe storage AHCI/RAID driver.

2. When prompted, insert the media with the Intel® RSTe 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.

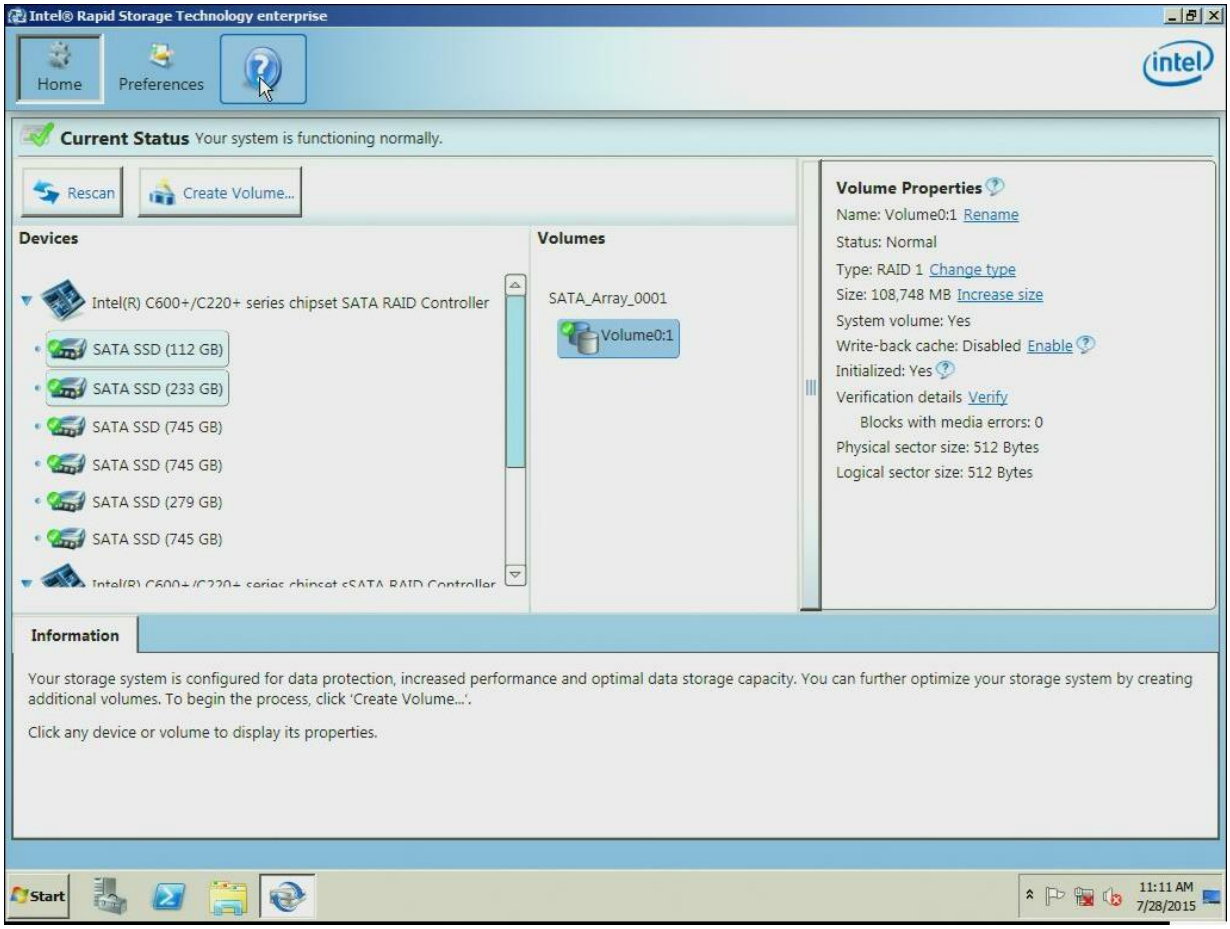
5. Continue the installation.

16 Intel® RSTe Component Versions

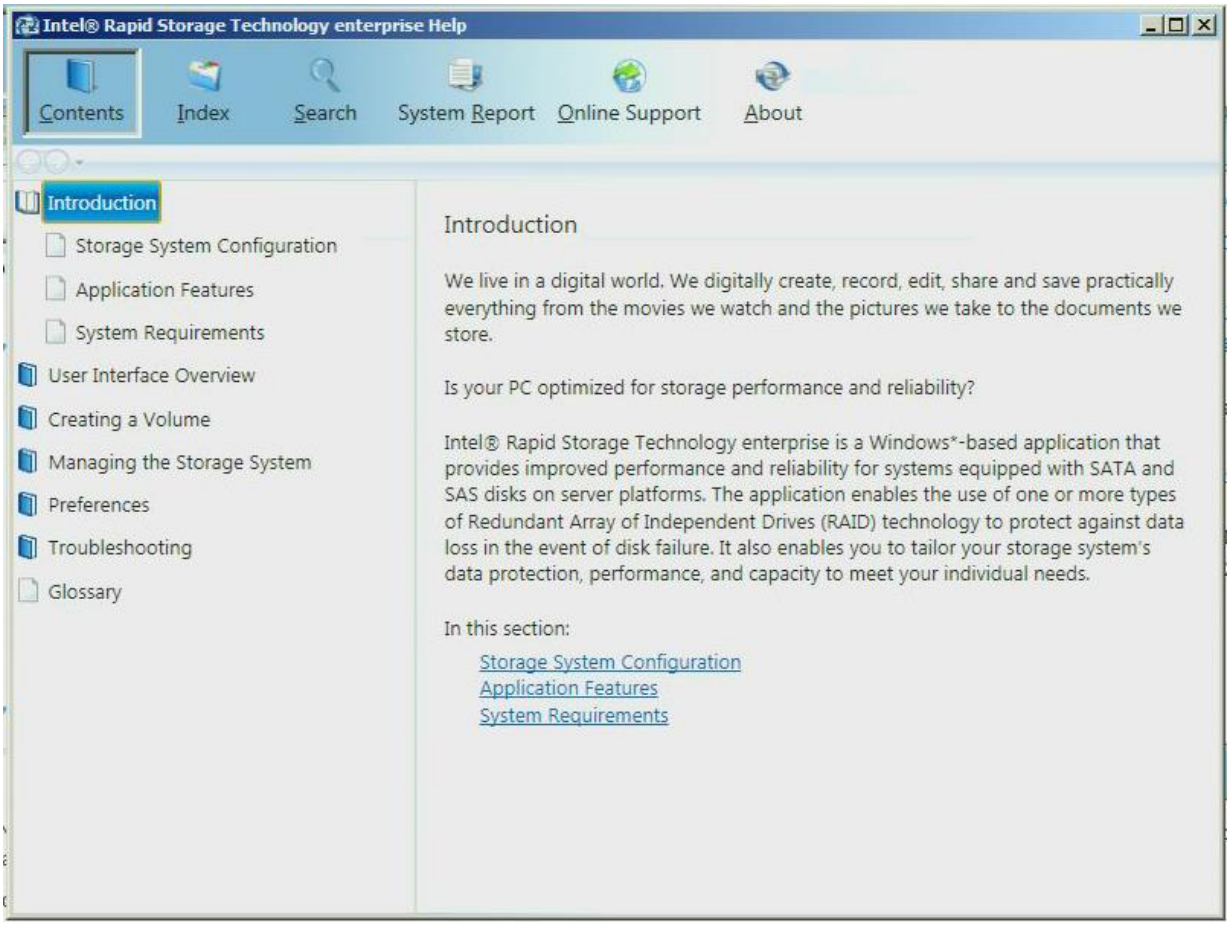
There are multiple ways to identify version(s) the different Intel® RSTe components installed. The first is to use the Intel® RSTe GUI. The second alternate method is to locate the specific version number on each individual component.

16.1 Using System Report

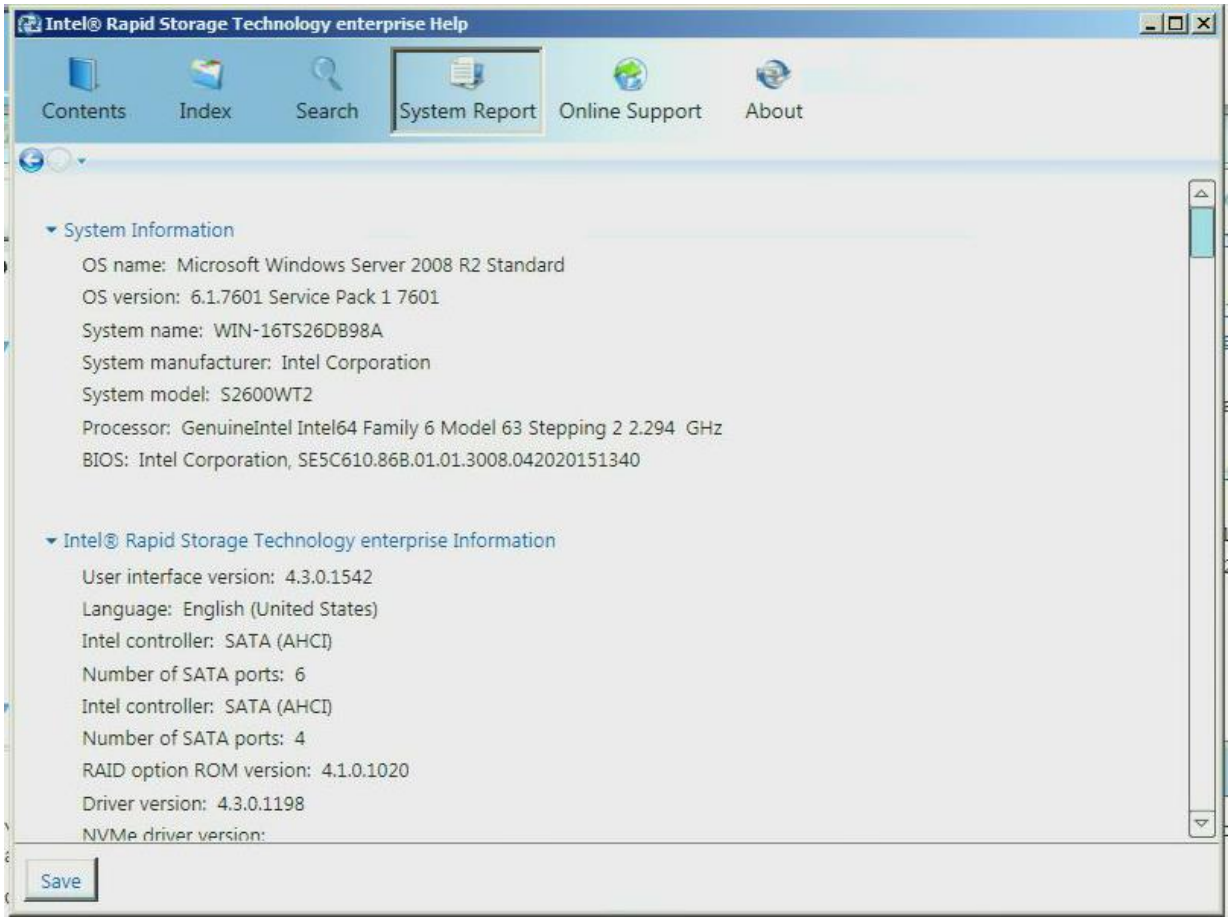
The System Report from the Intel® RSTe GUI will provide the detailed information about the system. It contains platform information , RSTe component information and information on all the attached devices. To obtain this information opent the Intel® RSTe GUI, and select the "?" button to open the Help window.



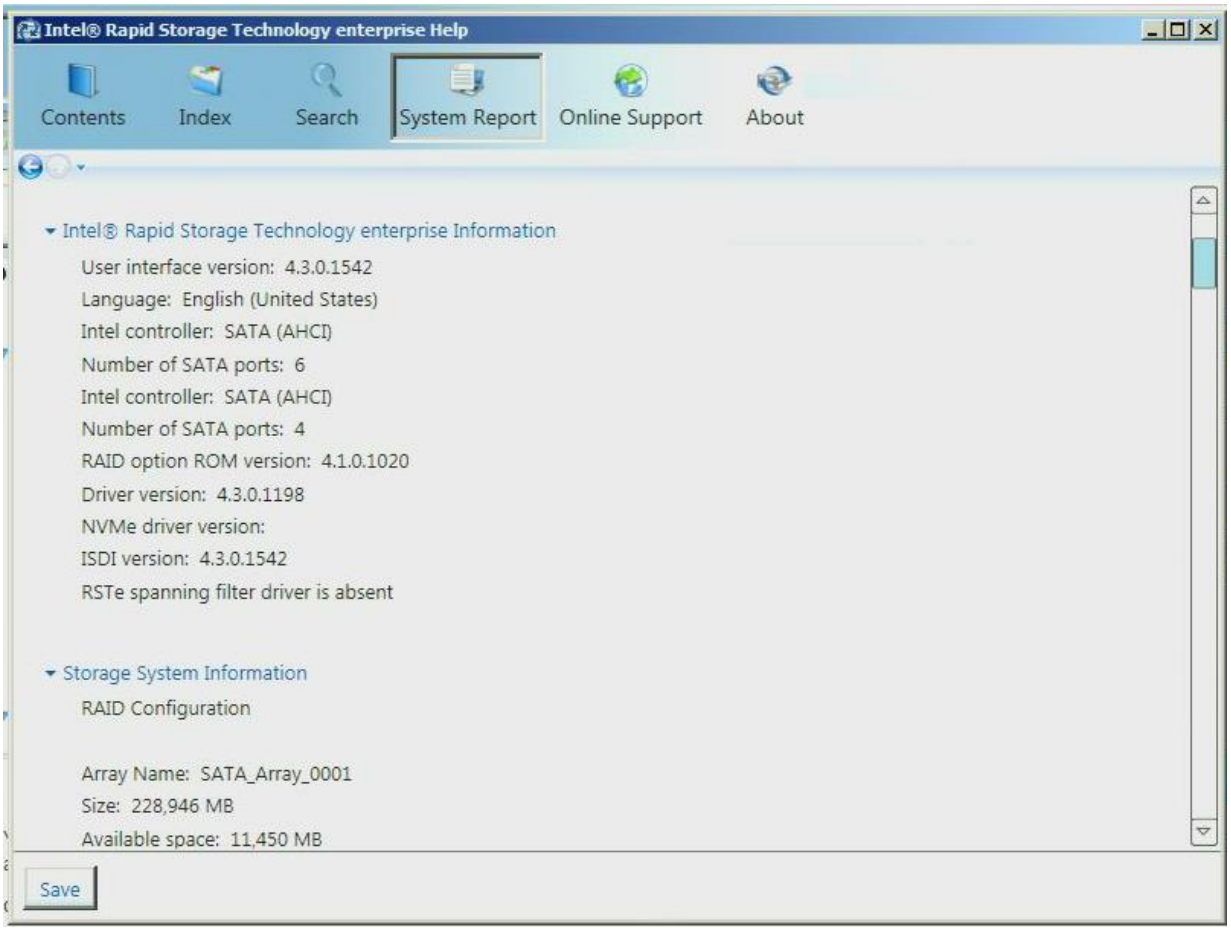
Select the "System Report" button.



The System Report windows provides a way to scroll through the information and an option to save the system report information to a file. This file can be provided to Intel to support debugging issues. To save the file, click on “Save” and follow the instructions.

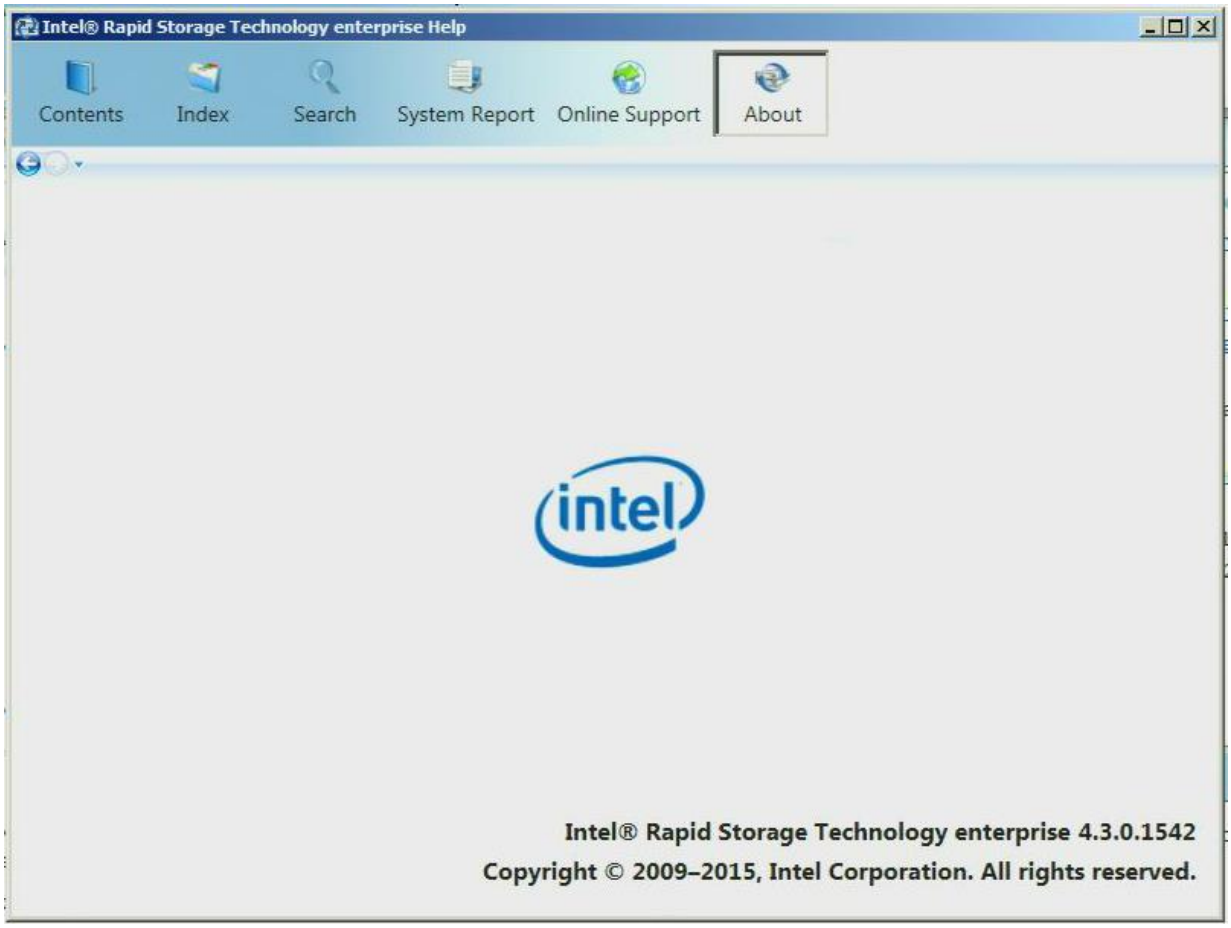


To view the versioning information for the RSTe components, scroll down to make all entries visable.



16.2 Intel® RSTe Graphical User Interface (GUI) Version

The specific version of the Intel® RSTe GUI is obtained from the Help window by clicking on “About”



16.3 Using Intel® RSTe File Properties to Obtain Driver Version

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® RSTe GUI.

16.4 Determining the Version of the Intel® RSTe Option ROM

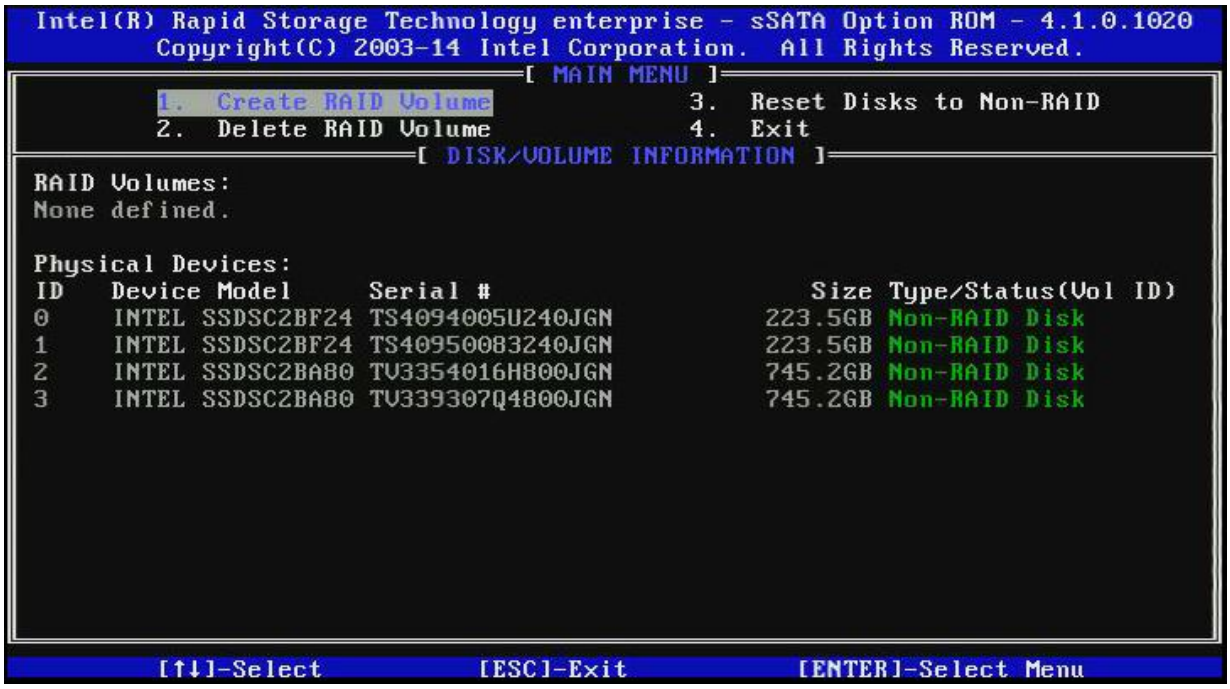
There re two ways to determine the version of the Intel® RSTe Legacy option ROM integated into the system BIOS. Use the following procedure to determine the version.

16.4.1 Using the Intel® RSTe GUI System Report

1. Follow the procedure illustrated in section 11716.1 Using System Report
2. Look for the parameter RAID Option ROM version.

16.4.2 Using the Intel® RSTe 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:



16.5 Determining the Version of the Intel® RSTe UEFI Driver

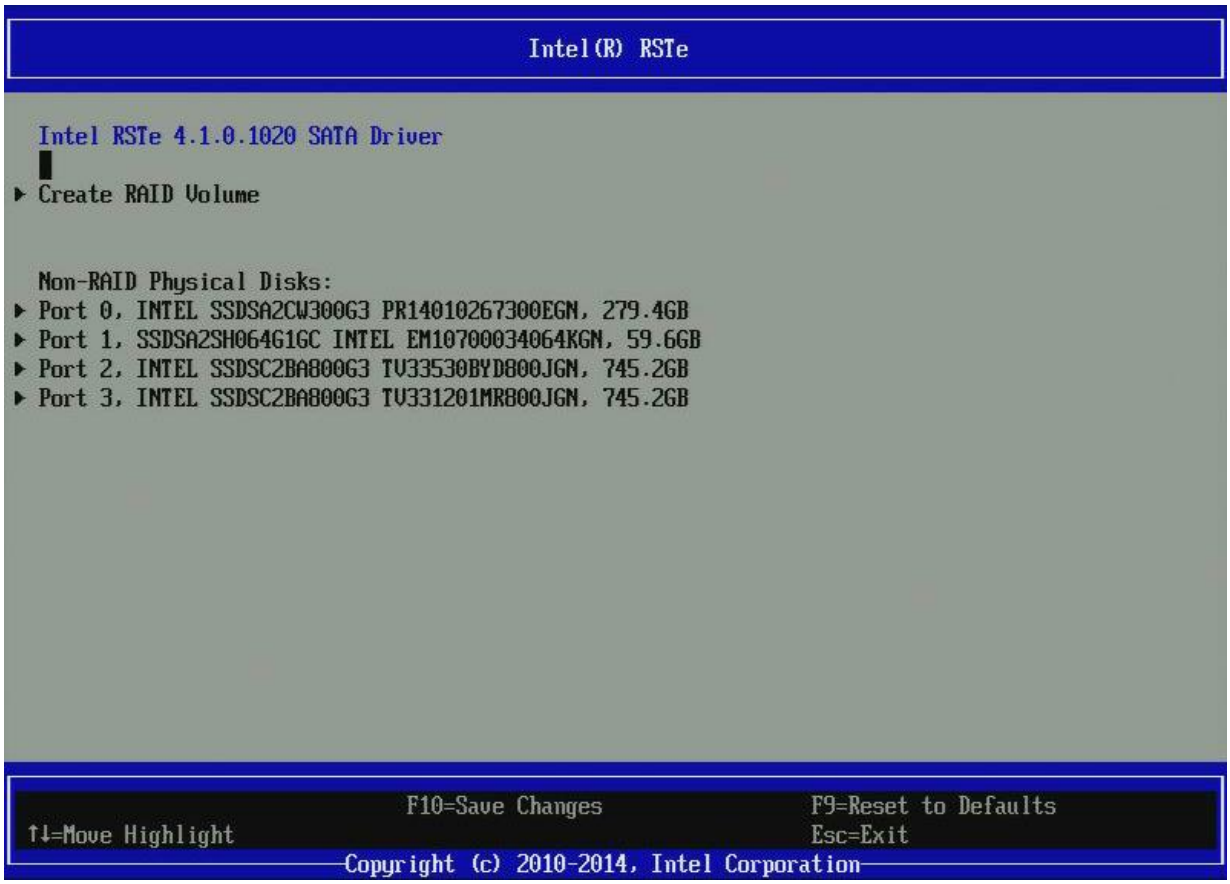
There are three ways to determine the version of the Intel® RSTe UEFI driver(s) integrated into the system BIOS. Use the following procedure to determine the version.

16.5.1 Using the Intel® RSTe GUI System Report

1. Follow the procedure illustrated in section 11716.1 Using System Report
2. Look for the parameter RAID UEFI driver version.

16.5.2 Using the UEFI HII

To obtain the Intel® RSTe UEFI driver version from the UEFI HII, enter into the BIOS Setup and navigate to the RSTe UEFI HII menus. The Version is at the top of the window.





16.5.3 Using the UEFI Shell

When the plarform BIOS is configured to boot from the UEFI Intel® RSTe UEFI, to obtain the driver version or to verify that the UEFI driver is loaded just enter into the BIOS setup or press the hot key to enter into the Boot Option menu. Boot into a UEFI Shell environment.

Shell:>Drivers

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

“CD 0000000B B - - 1 2 Intel® RSTe xx.x.x.xxxx SATA Driver”



17 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 uninstall application that is included with the Intel® Rapid Storage Technology enterprise 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 enterprise software or to disable the SATA RAID controller:

17.1 Uninstalling the Intel® RSTe Software (except the RAID Driver)

1. Run the Uninstall program from the following start menu link:
2. Start→All Programs→Intel® Rapid Storage Technology enterprise →Uninstall
3. The first dialog box that appears gives you the option of un-installing all components of the Intel® Rapid Storage Technology enterprise 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 enterprise. However, the RAID configuration should still function normally.

17.2 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.

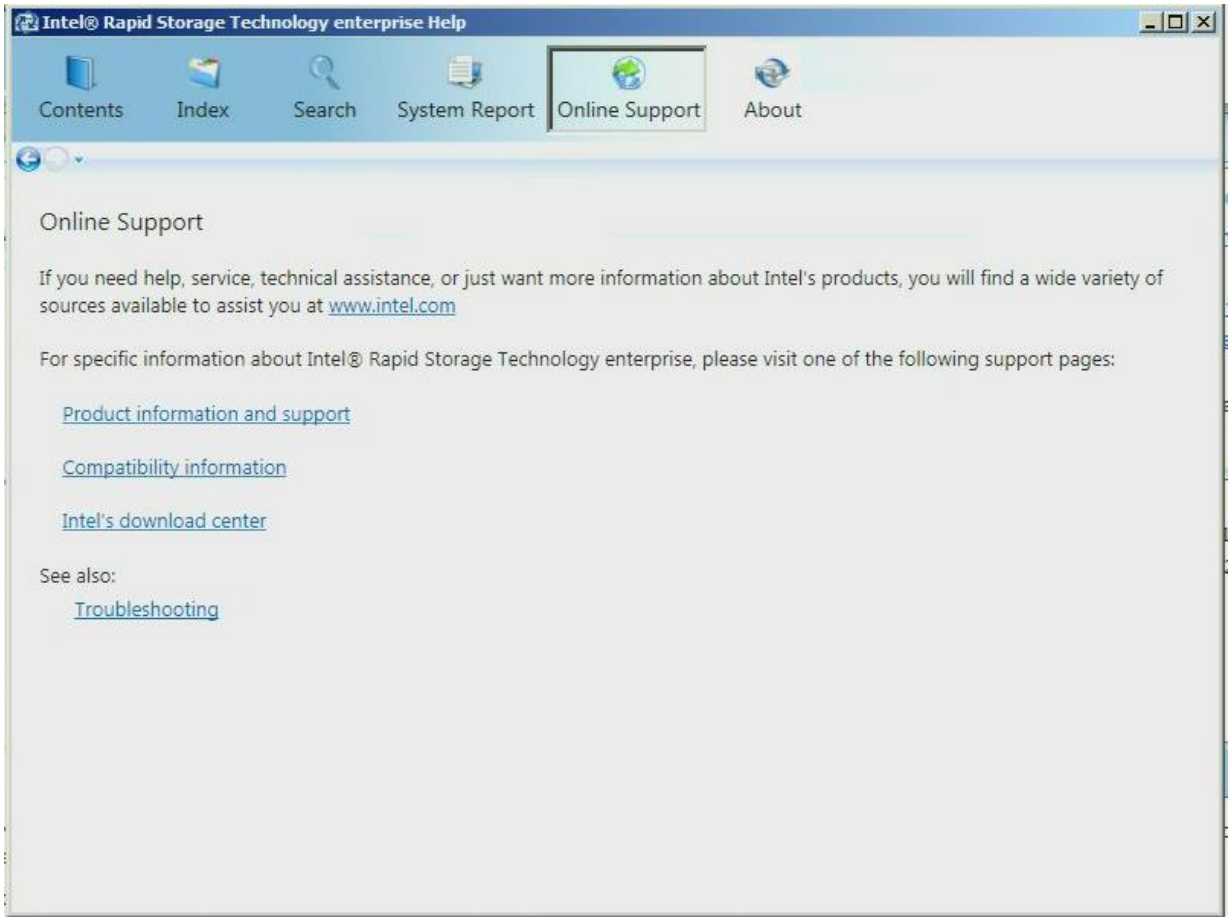
Uninstall Note: End-users can use this same procedure to disable the Intel® SATA RAID Controller if necessary. Run the Uninstall Program, click 'Cancel' when presented with the first dialog box, then click 'Yes' at the second dialog box to read the text document containing the procedure.

18 Registry Customizations

Note: Windows registry changes require reboot to take effect.

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

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



18.1 E-mail Notification UI Visible Enable/Disable

By default the e-mail notification feature is visible in the Intel® RSTe GUI 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\RSTe

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.

18.1.1 Email Alerting and Notification[1]

Intel® RSTe will support email notification of certain storage events (see **Error! Reference source not found.** for supported events). The Intel® RSTe UI will provide the interface for enabling/disabling and configuring the email notification feature. **The default setting in the UI is 'disabled'**.

The email notification feature allows the user to configure the platform to send alert / notification emails for each storage subsystem event that gets reported by the Intel® RSTe monitor service.

[1] This feature has a platform specific limitation. It is supported only on Intel® C600 series chipset based platforms, not supported on legacy platforms/chipsets.



18.1.1.1 Configuration

The Intel® RSTe user application will provide the interface to allow the user to configure the email alert notification feature via the 'Preferences' tab of the UI (***user must be logged on with administrative privileges***).

- User can enable/disable the email notification feature
- User can configure the level of storage system events to be sent via email notification (Storage system Information, Warning, and/or Error). Any combination of the three alert levels can be configured to trigger an email notification
- User can configure the email settings:
  - SMTP host (required)
    - Port (required)
    - Return email address (required)
    - Recipient email addresses (one address required, up to 3 maximum)
  - User can configure the Email alert / notifications to send test emails to all addresses specified

18.1.1.2 Email Message Format

- Message header:
  - Item1. Return email address: email address of the originating computer
  - Item2. Recipient email address: email address of computer receiving the email notification
  - Item3. Subject: system formatted subject content with product name, the storage system event level and the hostname of the originating computer
- Message body:
  - Item4. Log file text: contains the text of the event as it is displayed in the event log
  - Item5. System report: contains the system configuration information of the originating computer as seen in the Intel® RSTe GUI Preferences page.
- Optional text:
  - Item6. This section is blank unless the originating computer's OS is in a language other than English. If the originating computer sends items 4 and 5 in non-English, the English translation of those two items will appear in this section (for test emails, only item 4 will be translated here)

Figure 3 Email Format shows detail of how an email alert is formatted.

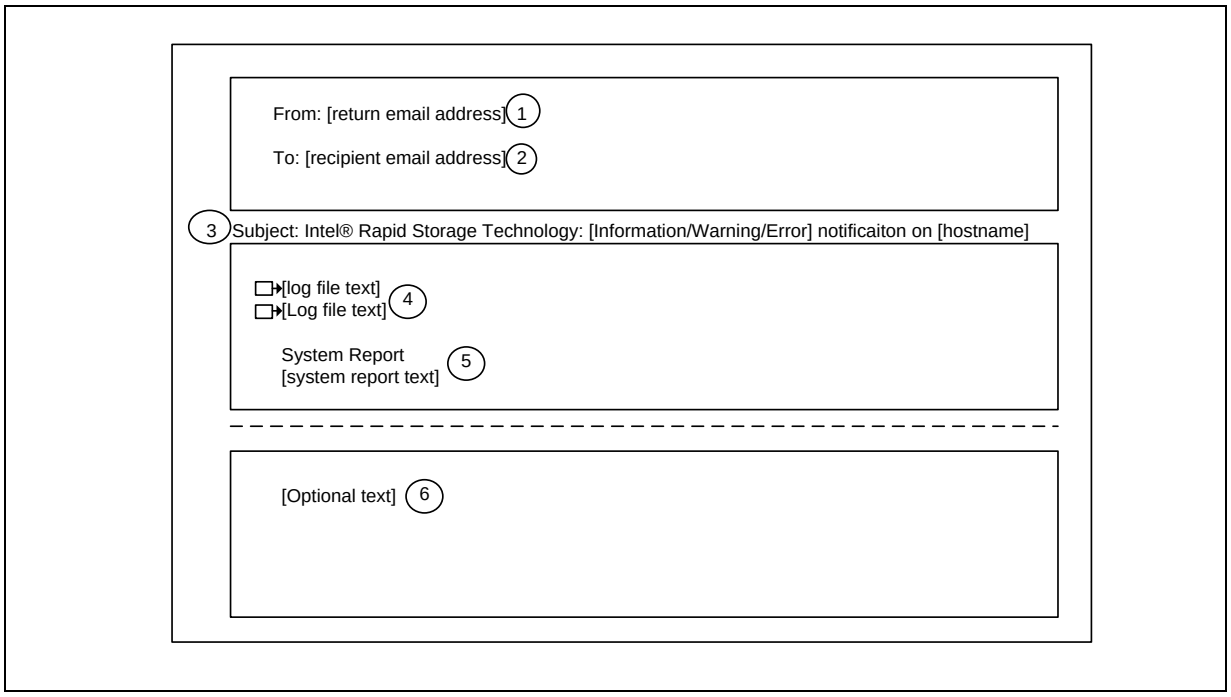


Figure 2 Email Format

18.1.1.3 Protocol Support

Email alert shall support SMTP host & SMTP port.

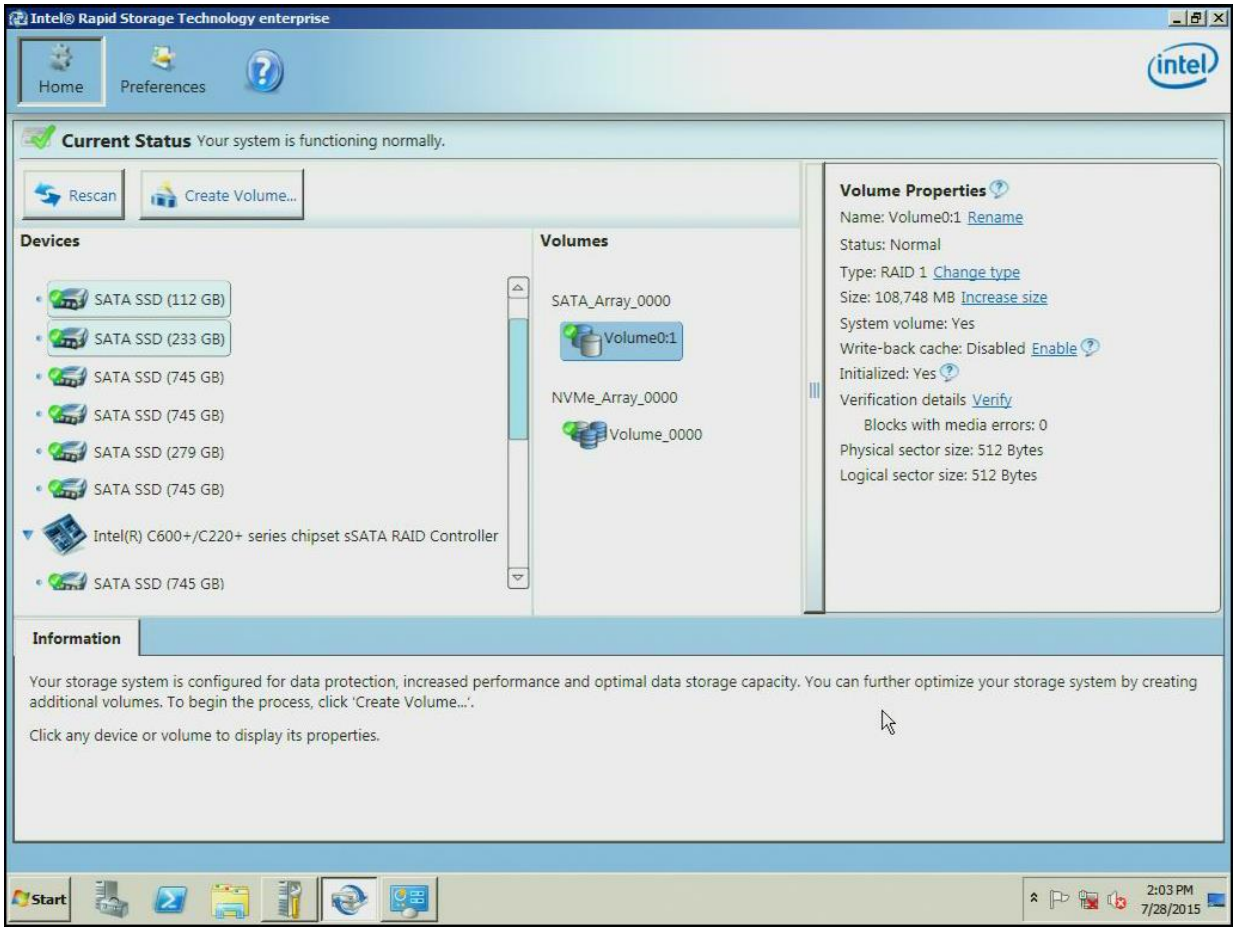
18.1.1.4 Error Conditions

Error! Reference source not found.and their notification mechanism:

- In the event of an SMTP server failure, the system will immediately attempt 3 retries. If the retries are unsuccessful, the system will discard the message without further attempts. The unsuccessful attempt will be written to the NT Event log.
- In the event of an improperly formatted email address in the “To’ field, the alert will fail and the failure will be written to the NT Event log.
- In the event of an improperly formatted email address in the “From” field, the alert will fail and the failure written to the NT Event log.
- If the SMTP name entered during configuration is an invalid format, the alert will fail and the failure written to the NT Event log.

19 Intel® Rapid Storage Technology enterprise GUI

NOTE: ***The Intel® Rapid Storage Technology enterprise GUI is not required to be installed when the storage subsystem is operating in AHCI mode. The main benefit of the Intel® RSTe GUI is in its management and monitoring of the Intel® RSTe RAID storage subsystem.***



19.1 Introduction

The Intel® Rapid Storage Technology enterprise GUI is a Windows*-based application that provides users monitoring and management capabilities for the Intel® RSTe storage subsystem. It offers a wide range of monitoring and



management activities for the Intel® RSTe RAID subsystem (***In AHCI mode there are no management or monitoring capabilities offered by the UI application***).
The RAID subsystem is defined as drives attached to the PCH or Intel NVMe SSD attached to the PCI express ports managed by the CPU.

19.1.1 Getting Started

The Intel® Rapid Storage Technology enterprise software package provides high-performance SATA/sSATA AHCI, SATA/sSATA RAID and Intel NVMe SSD RAID capabilities for supported operating systems.

The Intel® Rapid Storage Technology enterprise (RSTe) GUI requires (at a minimum) Microsoft .NET 3.5.

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 enterprise is to improve storage system performance, data protection, and increase fault-tolerance.

This Technology enterprise provides support for the following features:

- **Intel® Rapid Storage Technology enterprise RAID**
  This Technology enterprise 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 enterprise 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.
- **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.

- **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. TRIM is supported on pass-through SSDs and SSDs in RAID0, 1, 5, and 10.

- **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.

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 RSTe UEFI pre-OS driver or legacy OptionROM)**

19.1.2 Understanding the Application

The Intel® Rapid Storage Technology enterprise 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.

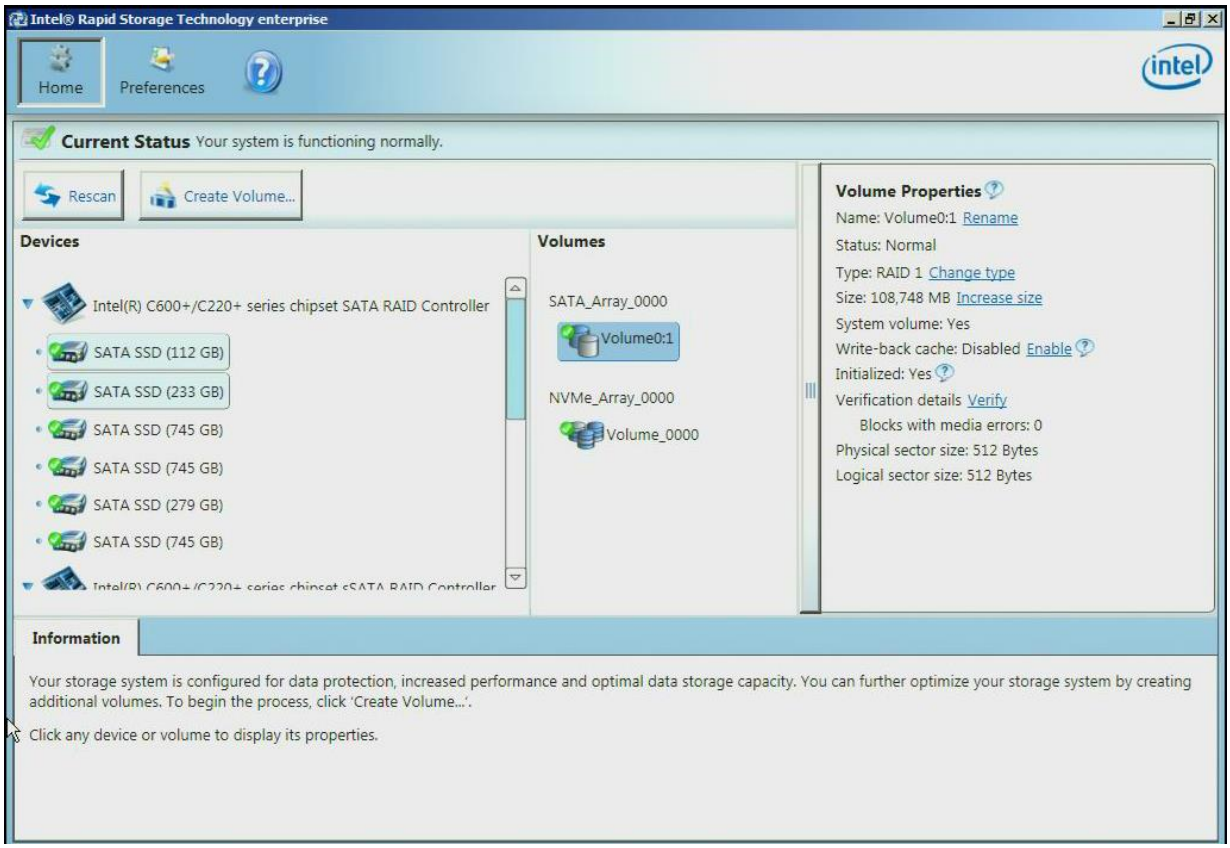


19.1.2.1 Storage System Configuration

The storage system combines hardware capabilities with RAID Technology enterprise to provide flexible data storage units on your computer. Each data storage unit, or RAID configuration, consists of three elements that include

1. The physical disk:
    i. SATA disks
    ii. Intel NVMe SSDs
2. The RAID array
3. The RAID volume

When at least one RAID volume is present on the system, these elements are represented in the storage system. The drives on controller information is represented in the “Devices” pane to the left. The arrays and volumes are represented in the “Volumes” pane in the center. The properties pane to the right show the properties information of what ever it selected (i.e. Volume0.1).



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, solid state drive or Intel NVMe SSD) 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.

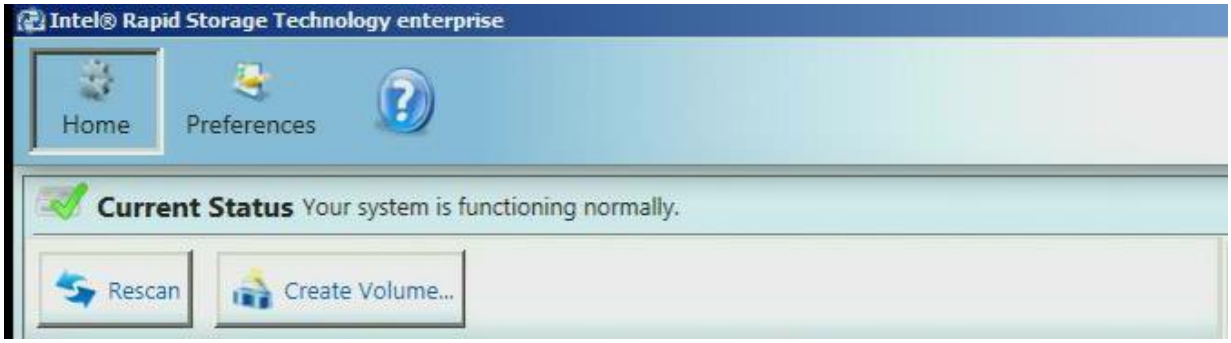
19.1.2.2 Navigation

The application is organized into three main areas depicted by the top navigation buttons: Home, Preferences and Help. Depending on your computer's configuration and available hardware, Create and Accelerate may not be available.

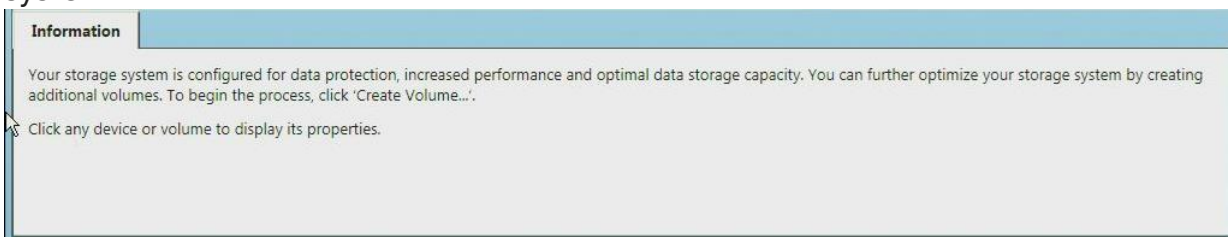


Home

The Home is the main area of operation. This is where the user can select to “Scan” for hardware changes or “Create Volume”

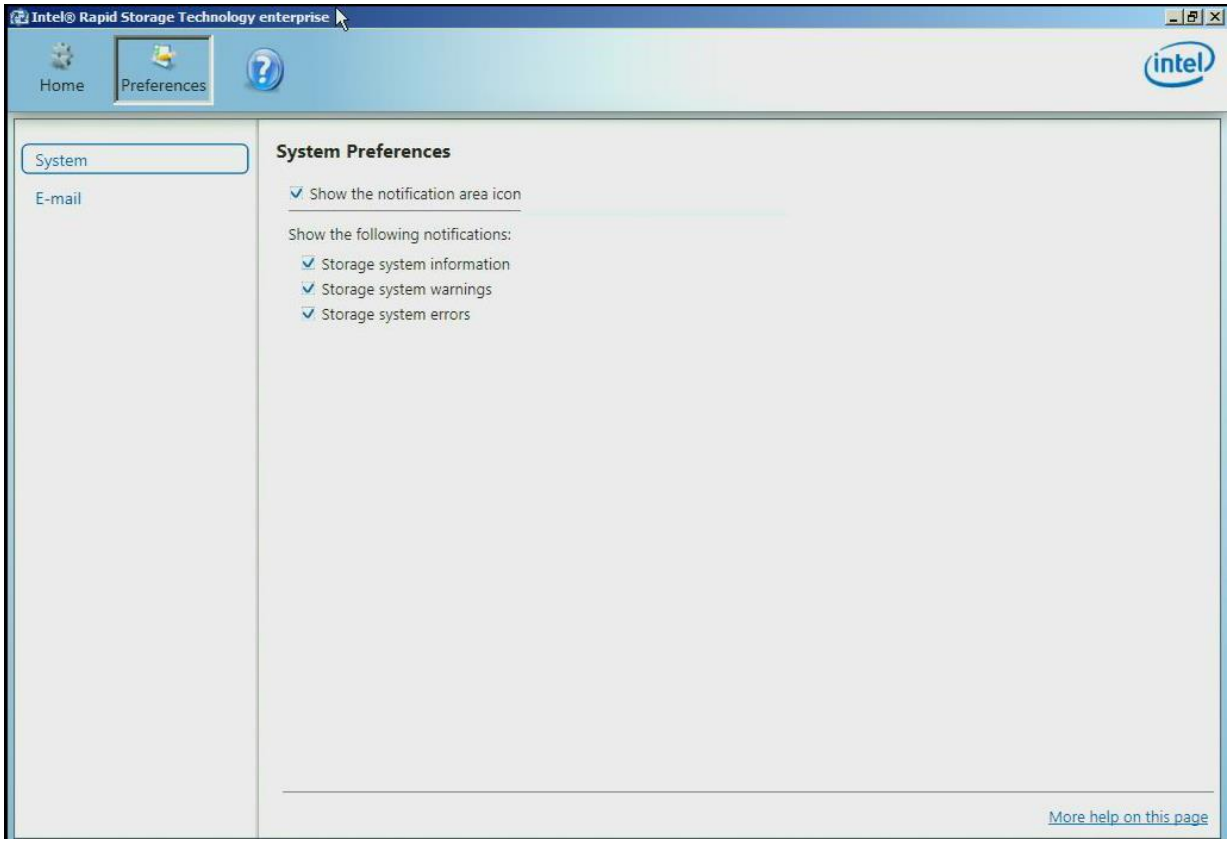


This is also the working area that shows all of the controllers, drives, volumes, arrays and properties. At the bottom of the page is the Information or current status of the system.



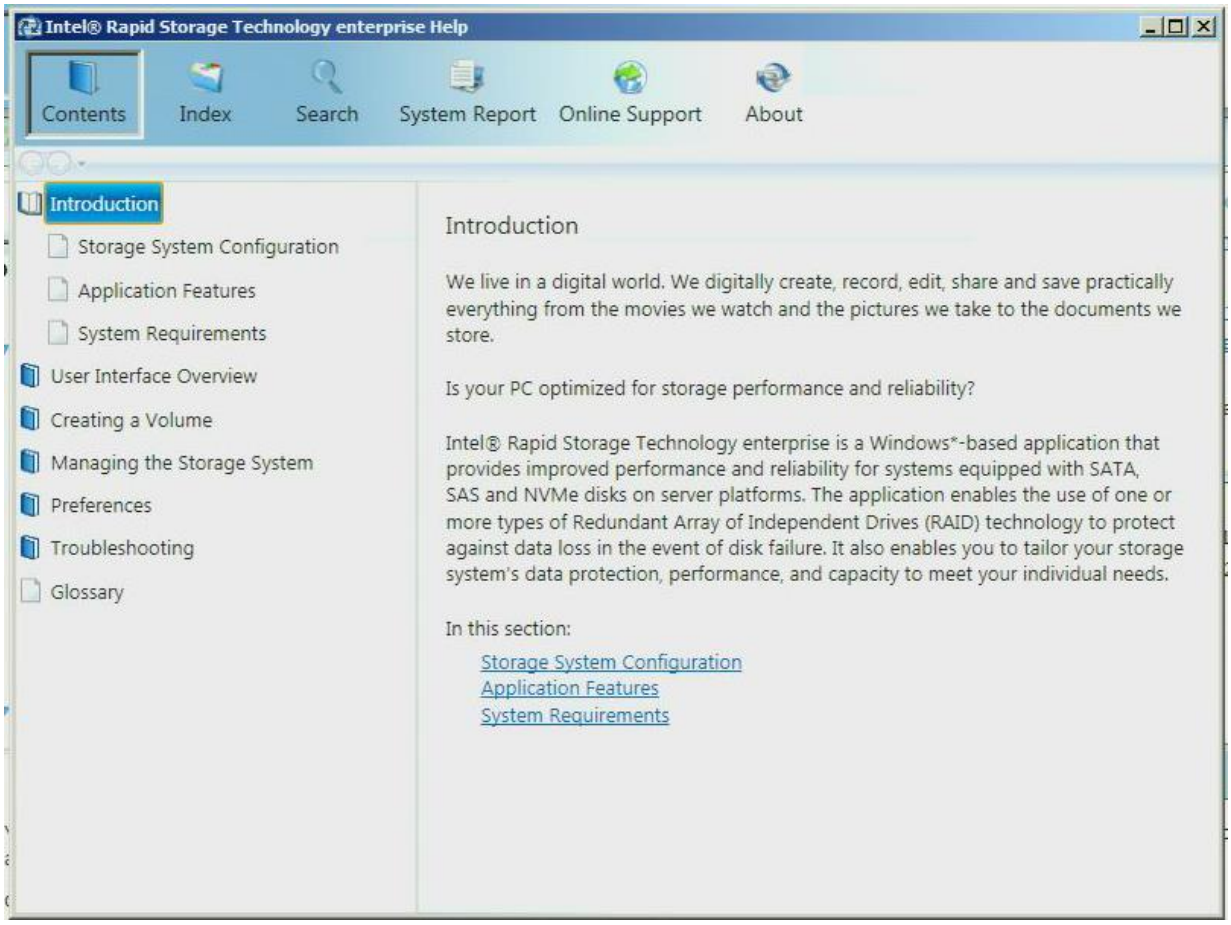
Preferences

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.




Help

The “Help” area is where the detailed information on how to use the Intel® RSTe GUI and provides key system information



19.1.3 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 enterprise. 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 enterprise 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

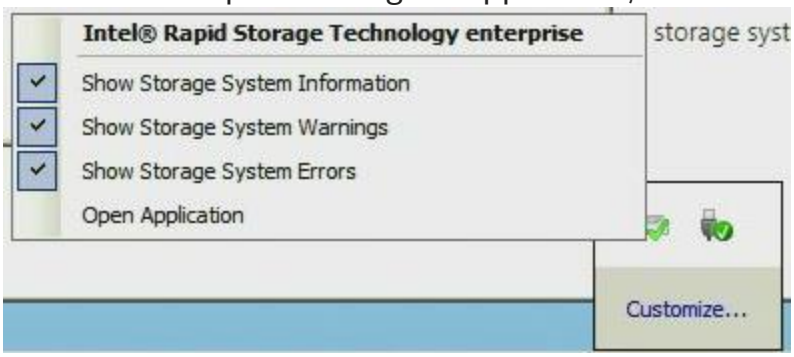
1. Right-click the icon.
2. Click 'Open Application'.

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 enterprise 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.

19.2 Storage System Status

Anytime Intel® Rapid Storage Technology enterprise 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.

19.2.1 Understanding the Status

To get the full benefits of what Intel® Rapid Storage Technology enterprise 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.

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 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

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 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

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		

19.3 Creating a Volume

SATA disks that reside on the same PCH controller (SATA/sSATA) or Intel NVMe SSDs plugged into PCI express slots controlled by the CPU can be used to create a volume in order to enhance your storage system. This can be accomplished 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.

19.3.1 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 (or Intel NVMe SSDs), including the operating system disk are connected, in a normal state, and unlocked (only applies to password-protected disks).
- Each of the 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 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.
    - Intel NVMe SSDs cannot be part of a System/OS RAID 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 (SATA/sSATA/NVMe). An SATA disk shall not be paired with an sSATA disk to create a volume.

19.3.2 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.

19.3.2.1 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 19.3.3 Creating Additional Volumes for more information on this type of configuration. Below is an overview of the five volume types that you can create.

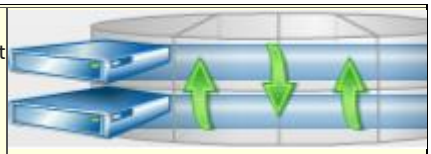
Creating a two-disk volume from 'Status'

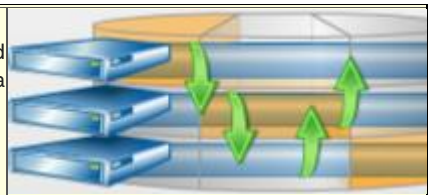
This option displays if only two disks are available, one has data such as system files, the second one doesn't, and the latter has a size that is equal or greater than the other. Based on this simple configuration, you can create a volume to protect data or optimize disk performance by clicking one 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

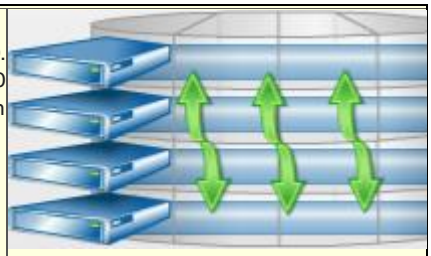
1. Click ‘Create’.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click ‘Next’.

RAID 1: **Real-time data protection** Combines two 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 eight disks and breaks down data into units that are spread across the array disks.	
Disks required	2 to 8
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 eight n a rotating sequence. Parity is a mathematical method for recreating lost data to a single disk.	

Disks required	3 to 8
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 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.

19.3.2.2 Configuring the Volume

Once the volume type is selected, you are ready to configure your volume.

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.

19.3.2.3 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 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.

19.3.3 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 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

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.

19.4 Managing the Storage System

The Home 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 area which is accessible by clicking any element displayed in the storage system..

The 'Home' area also provides the actions available for the selected element, such as renaming a volume or changing the volume type.

19.4.1 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 “Volume” pame. 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.

19.4.1.1 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.

19.4.1.2 Adding a Disk to an Array

You can add one or more 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.

⚠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.



1. Under “Home”, 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

19.4.1.3 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 ‘Home”, 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
5-disk RAID 0	5-disk RAID 0 5-disk RAID 5
6-disk RAID 0	6-disk RAID 0 6-disk RAID 5

19.4.1.4 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 ‘Home”, in the storage system view, click the array you want to manage. The array properties are now displayed on the right.
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 ‘Home’, in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the right.
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

19.4.1.5 Enabling Disk Data Cache

Enaling the disk data cache for all 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 by selecting the Array in the ‘Volumes” pane. The array must be selected because the data cache must be in the same state across all disks that are part of a single array.

Under Array Properties, the disk data cache is reported as enabled or disabled for all SATA disks in the array. Under Disk Properties, 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 Array Properties.

⚠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 ‘Home”, in the storage system view, click the array you want to manage. The element properties are now displayed on the right.
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.

19.4.2 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 ‘Home’. This allows you to review the volume properties and access all actions associated with that volume, such as renaming, changing type, and deleting.



19.4.2.1 Volume Properties

A volume is an area of storage on one or more 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 right of the storage system view under 'Home' 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.
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.



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.

General parameters table

Parameter	Value
Details	Provides detailed information if a volume is in a state other than normal.
Type	Reports the volume type.
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.

19.4.2.2 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 ‘Home’, in the storage system view, click the volume that you want to rename. The volume properties are now displayed on the right.
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 “\”.



19.4.2.3 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 'Home", in the storage system view, click the volume you want to rebuild. The element properties are now displayed on the right.
2. Click 'Rebuild to another disk', and then follow the procedure described above.

19.4.2.4 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 'Home', in the disk “Properties” pane, click 'Reset volume to normal'.
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.

19.4.2.5 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 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 8-disk RAID 0 3, 4, 5 or 8-disk RAID 5
2-disk RAID 0	3, 4, 5 or 8-disk RAID 5
3-disk RAID 0	4, 5 or 8-disk RAID 5
4-disk RAID 0	5 or 8-disk RAID 5
4-disk RAID 10	4, 5 or 8-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 disks are connected. The Intel® chipset provides support for the creation of all RAID volume types and for up to 8 SATA ports and 8 Intel NVMe SSD connected to the CPU. Changing volume type does not require re-installation of the operating system

1. Under ‘Home’, in the storage system view, click the volume that you want to modify. The volume properties are now displayed on the right.
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.

19.4.2.6 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.

19.4.2.6.1 Increasing the volume size from Manage Array

1. Under 'Home', in the storage system view, click the array you want to manage. The array properties are now displayed on the right.
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.

19.4.2.6.2 Increasing the volume size from Manage Volume

1. Under 'Home', in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the right.
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

19.4.2.7 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 ‘Home’, in the storage system view, click the volume to which you want to add a disk. The element properties are now displayed on the right.
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

19.4.2.8 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 ‘Home’, in the storage system view, click the volume you want to delete. The volume properties are now displayed on the right.
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.



19.4.2.9 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 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 'Home', 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 'Home', in the storage system view, click the RAID volume that you want to modify. The volume properties are now displayed on the right.
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.

19.4.2.10 Enabling Volume Write-back Cache

You can improve the read/write performance of a RAID 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 'Home', 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 right.



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 'Home', 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 right.
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.

19.4.2.11 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 'Home', in the storage system view, click the volume that you want to initialize. The volume properties are now displayed on the right.
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 'Home', in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the right.
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 enterprise icon, or in the application under the Home pane.

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.

19.4.2.12 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 ‘Home’, in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the right.
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.

19.4.3 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 ‘Home’. This allows you to review the properties and access all actions associated



with that disk, such as marking a disk as spare. Depending on their usage or status, some actions may not be available.

19.4.3.1 Disk Properties

The disk properties listed below display to the left of the storage system view under ‘Home” 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.
	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.
Status	Normal: the disk is present, functioning as expected.
	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
Status Size Serial number Model Firmware System disk	Reports the total capacity of the disk in megabytes (MB) in the disk properties and in gigabytes (GB) in the storage system view.
	Reports the manufacturer's serial number for the disk.
	Reports the model number of the disk.
	Reports the version of the firmware found in the disk.
	Reports whether the disk contains system files that are required to start and run the operating system.
	Reports whether the disk is protected with a password.
Disk data cache	Reports whether the disk supports this feature.
Native command queuing	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).

19.4.3.2 Marking a Disk as Spare

This action is only available for non-system disks in a normal state.

Marking a disk as a spare allows you to designate an available 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 ‘Home’, in the storage system view, click the disk that you want to mark as a spare. The volume properties are now displayed on the right.
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 Intel® RSTe Legacy 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.

19.4.3.3 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 ‘Home’, in the storage system view, click the disk that you want to reset to available. The volume properties are now displayed on the right.
2. Click 'Reset to available'.
3. The page refreshes and the disk usage is now reported as available.

19.4.3.4 Resetting a Disk to Normal

You can reset a 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 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 'Home', in the "Device" pane, locate the disk reported as at risk or failed. You can also perform this action from Disk Properties to the right..
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.

19.4.3.5 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 enterprise provides SATA drive 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.

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 'Home', in the storage system view, that the new disk displays. You can then access management options by clicking that disk.

19.4.4 Managing Ports

A port is a connection point on your computer where you can physically connect a device, such as a SATA disk, an ATAPI device or an Intel NVMe SSD. 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 SATA devices is eight (per controller).



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.

19.4.5 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.

19.5 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.

19.6 Email Alerting and Notification[2]

Intel® RSTe will support email notification of certain storage events. Please see section 26 Storage System Events Detected by Monitor Service (IAStorDataMgrSvc) for a full list of supported events. The Intel® RSTe GUI will provide the interface for enabling/disabling and configuring the email notification feature. **The default setting in the GUI is '<u>disabled</u>'**.

The email notification feature allows the user to configure the platform to send alert / notification emails for each storage subsystem event that gets reported by the Intel® RSTe monitor service.

19.6.1 Configuration

The Intel® RSTe user application will provide the interface to allow the user to configure the email alert notification feature via the 'Preferences' tab of the UI (***user must be logged on with administrative privileges***).

- User can enable/disable the email notification feature
- User can configure the level of storage system events to be sent via email notification (Storage system Information, Warning, and/or Error). Any combination of the three alert levels can be configured to trigger an email notification
- User can configure the email settings:
  - SMTP host (required)
    - Port (required)
    - Return email address (required)
    - Recipient email addresses (one address required, up to 3 maximum)

[2] This feature has a platform specific limitation. It is supported only on Intel® C600 series chipset based platforms, not supported on legacy platforms/chipsets.



- User can configure the Email alert / notifications to send test emails to all addresses specified

19.6.2 Email Message Format

- Message header:
  - Item1. Return email address: email address of the originating computer
  - Item2. Recipient email address: email address of computer receiving the email notification
  - Item3. Subject: system formatted subject content with product name, the storage system event level and the hostname of the originating computer
- Message body:
  - Item4. Log file text: contains the text of the event as it is displayed in the event log
  - Item5. System report: contains the system configuration information of the originating computer as seen in the Intel® RSTe GUI Preferences page.
- Optional text:
  - Item6. This section is blank unless the originating computer's OS is in a language other than English. If the originating computer sends items 4 and 5 in non-English, the English translation of those two items will appear in this section (for test emails, only item 4 will be translated here)

Figure 3 Email Format shows detail of how an email alert is formatted.

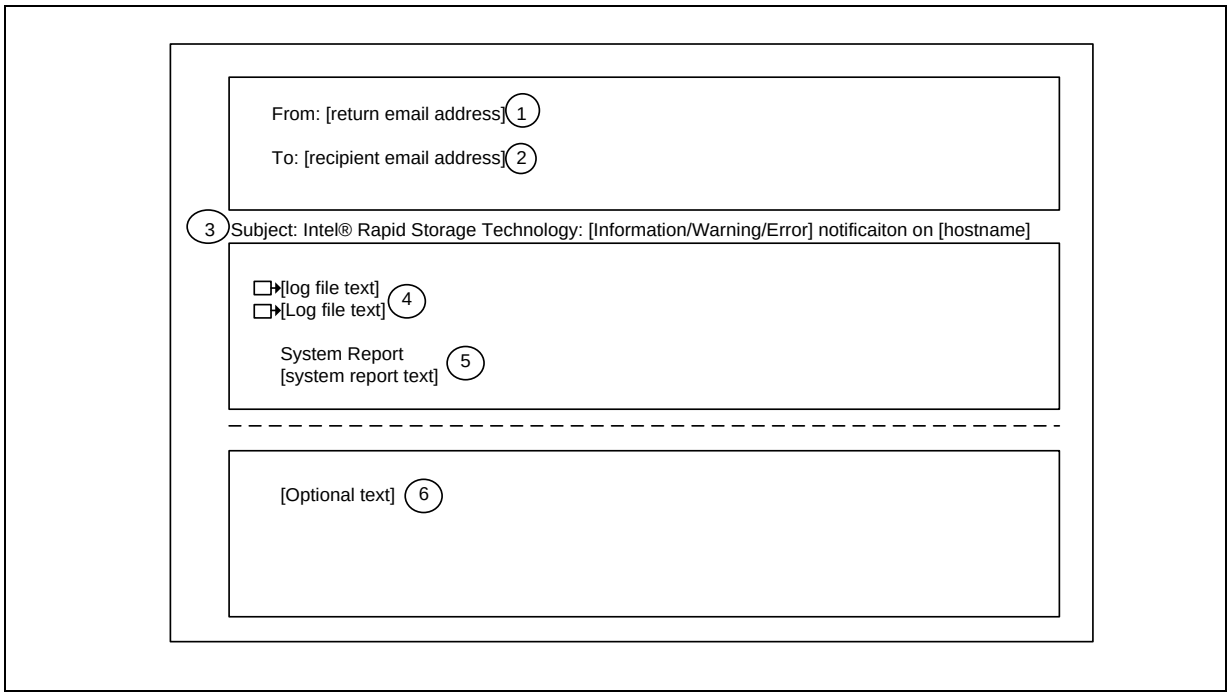


Figure 3 Email Format

19.6.3 Protocol Support

Email alert shall support SMTP host & SMTP port.

19.6.4 Error Conditions

See **Error! Reference source not found.** and their notification mechanism:

- In the event of an SMTP server failure, the system will immediately attempt 3 retries. If the retries are unsuccessful, the system will discard the message without further attempts. The unsuccessful attempt will be written to the NT Event log.
- In the event of an improperly formatted email address in the “To” field, the alert will fail and the failure will be written to the NT Event log.
- In the event of an improperly formatted email address in the “From” field, the alert will fail and the failure written to the NT Event log.
- If the SMTP name entered during configuration is an invalid format, the alert will fail and the failure written to the NT Event log.

20 Using the BCFS to Differentiate Platform SKUs

Intel® RSTe has support for BIOS Control Feature Set (BCFS) to enabled OEMs the opportunity to customize the Intel® RSTe feature offerings. OEMs no longer need special Intel® RSTe 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 RSTe Feature Capabilities***' register in the Intel chipset's SATA controller MMIO space, OEMs now have greater flexibility in determining what Intel® RSTe 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 please contact your platform's BIOS vendor or your Intel field representative to put you in contact with the appropriate Intel BIOS support personnel.

20.1 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 RSTe 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' is an unsupported configuration. The Intel® RSTe OROM will ignore the BIOS settings and enable all RAID levels.

20.1.1 BCFS Bit Setting

BCFS bit number	Bit meaning	Values		Additional info
0	Enable/disable Raid0	Raid type disabled	0x0	If you disable all raid levels – all BCFS settings will be set back to default and OROM UI
		Raid type enabled	0x1	
1		Raid type disabled	0x0	



	Enable/disable Raid1	Raid type enabled	0x1	delay will be set to 2 seconds
2	Enable/disable Raid10	Raid type disabled	0x0	
		Raid type enabled	0x1	
3	Enable/disable Raid5	Raid type disabled	0x0	
		Raid type enabled	0x1	
4	RESERVED			
5	Enable/disable UI	Feature disabled	0x0	If you set it to 0 (disable) bits 10-14 are ignored
		Feature enabled	0x1	
6	Enable/disable unlock of HDD in OS	OS can unlock HDDs OS cannot unlock HDDs	0x1 0x0	
7	Enable/disable LED SGPIO	Feature disabled Feature enabled	0x0 0x1	
8	eSATA RAID usage	Allow all volume types to span internal and external ports Allow only R1 volume type to span internal and external ports. Other volume types – allow all eSATA-internal disks only if there are enough of them to create the volume	0x0 0x1	
9	reserved			



10 - 12	Delay on UI splash screen	2 seconds	0x000	Default setting is 0x000: 2 seconds.
		4 seconds	0x001	
		6 seconds	0x010	
		8 seconds	0x011	
		10 seconds	0x100	
		15 seconds	0x101	
		30 seconds	0x110	
		60 seconds	0x111	
13 - 14	Mode of showing UI	Show if error or >=2 disks	0x00	Default setting 0x00: show if there are 2 or more disks connected or error occured
		Show only if error	0x01	
		Never show UI	0x10	
		Show always	0x11	
15	RESERVED			

20.1.1.1 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



20.1.2 Configuring the Behavior of the OROM UI and Banner

There are three (3) bits that control the behavior of the Intel® RSTe 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 RSTe 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.

20.1.2.1 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

20.1.3 Configuring Intel® RSTe UI Capabilities

There are a few capabilities within the Intel® RSTe UI that is controlled by the BCFS bits.

HDD unlock and Delay on UI splash screen are controlled with the following bits:

Bits	Type	Reset/ Default	Description
6	RWO	0h	**HDD Unlock (HDDLK):** If set to ‘1’, then HDD password unlock is enabled in the OS.
10:12	RW	0x000	**Delay on UI splash screen:** 2 seconds 0x000 4 seconds 0x001 6 seconds 0x010 8 seconds 0x011 10 seconds 0x100 15 seconds 0x101 30 seconds 0x110 60 seconds 0x111

20.1.3.1 Example Configuration

To configure a platform SKU to not allow unlocking passwords from the Intel® RSTe UI and to allow the UI to activate the disk/port LEDs, the bits must be configured as follows:

Bit 6 == 0 (default)

20.1.4 BIOS Control Feature Set

Bits	Type	Reset	Description
15:14	RO	0h	Reserved.



13:12	RWO	0h	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	**Reserved**
8	RWO	0h	**eSATA RAID usage:** If set to '1', allow only R1 volume type to span internal and external ports. Other volume types – allow all esata-internal disks only if there are enough of them to create the volume If cleared to '0', then allow all volume types to span internal and external ports.
7	RWO	0h	**Enable/disable LED SGPIO:** If set to "1", then LED SGPIO support is enabled.
6	RWO	0h	**Enable/disable unlock of HDD in OS**: If set to "1", then unlock of HDD in OS is enabled.
5	RWO	1h	**Intel® RSTe 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	0h	**Reserved**



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



21 Testing, Certification Notes

21.1 Correcting Microsoft* Windows 7 (Win7) WHQL test failure

For Internal SATA ports with interlock switches, the RSTe 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, RSTe 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 RSTe 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 RSTe whether to set Removable bit or not. If _EJ0 ACPI method is defined in the system BIOS by the manufacturers, they can tell RSTe 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 RSTe will set Removable bit. The default value is 0.

22 Glossary

Term	Definition
ATA	Advanced Technology enterprise 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 & 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
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



Term	Definition
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 enterprise 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 of a member that 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 enterprise 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 enterprise 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 enterprise 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 enterprise 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 enterprise primarily designed for transfer of data between storage devices such as hard drives or optical devices, and a computer. Benefits of this



Technology enterprise 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 enterprise 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.



23 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.

23.1 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 enterprise option ROM user interface. 3. Once the operating system is running, open Intel Rapid Storage Technology enterprise from the Start menu or click the Intel Rapid Storage Technology enterprise 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 enterprise 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 enterprise 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 enterprise 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 enterprise 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.

23.2 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. **Note** 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. To reset the failed master disk and the volume to normal, follow this procedure: 1. 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. 2. 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. 3. Click 'Yes' to confirm. 4. 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'. 5. 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: 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. **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 failed. 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: 1. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 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.
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.

23.3 Other Volume States

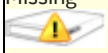
Incompatible	
Cause	Solution
Indicates that the volume was moved to another system that does not support the volume type and configuration.	In this situation, volume data is accessible to the operating system and can be backed up, but the volume cannot operate because your system does not support its 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.
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** 1. 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. 2. Once the operation has completed, restart your computer. 3. 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.

23.4 Disk Events

State	Cause	Solution

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. 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.
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.

23.5 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. 3. 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.
5. The page refreshes and the storage system displays the solid state disk usage as available.
6. Power off your computer and replace the failed solid state disk with an operational one.
7. 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.

23.6 Software Errors

Message	Cause	Solution
An unknown error occurred while running this application. If the problem persists, please restart your computer or try reinstalling the application.	This error may be related to: ssing components rrupted application plication unable to connect to the service plication fails to start.	Restart your computer or try reinstalling the application.

Intel® Rapid Storage Technology enterprise 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 enterprise 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 enterprise 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. The Intel Rapid Storage Technology enterprise 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. Please try again later.	The Intel® Rapid Storage Technology enterprise 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. Please 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. Please 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'.



24 RSTe SATA Port Bitmap Implementation

24.1 Legacy OROM

RSTe 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



24.2 UEFI Driver

*The RSTe 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).*

24.3 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.

25 Intel® C600 series chipset Support

25.1 Intel® C600 series chipset Supported Platforms

25.1.1 Intel® C600 series chipset (Romley/Romley Refresh) Platforms

Intel® Server Series (Sandy Bridge) Chipset-based platform CRBs with the Intel® C600 series chipset (PCH)

- River City (Workstation, 2 CPUs/Sockets)
- Harbor City (EN Server, 2 CPUs/Sockets)
- Rose City (EP, 2 CPU)
- Potter City (EX, 4 CPUs/Sockets)

25.1.2 Intel® C600 series chipset SKU Breakdown

SKUs / Features	-A (SATA Only) C202	-B C204	-D C206	-T C208
SATA – SATA2 Ports (3Gb/s)	4	4	4	4
SATA – SATA3 ports (6Gb/s)	2	2	2	2
SCU ports (3Gb/s)	4	4	8	8
Intel® RSTe SATA RAID Support	0/1/10/5	0/1/10/5	0/1/10/5	0/1/10/5
Intel® RSTe SCU RAID Support	0/1/5/10	0/1/10*	0/1/10*	0/1/5/10

* Intel® RSTe RAID 5 supported on SATA Only configurations with specific BIOS enabled settings.

Please reference CDI/IBL Document No. 454672 for information on properly configuring the PCHSTRP16 strap in BIOS. It is Intel's recommendation that the PCHSTRP16 strapping be left to the default –T SKU value.



25.1.3 Intel® C600 series chipset Upgrade ROM Options

Intel® RSTe will provide support for the Intel® C600 series chipset Upgrade ROM feature. When the Upgrade ROM is implemented to modify the base SKU of the Intel® C600 series chipset, the Intel® RSTe driver will not be required to be changed, updated or reinstalled.

Upgrade ROM SKU #	SCU Ports	Protocol	RSTe SCU RAID 5	Equivalent SKU (Device ID)
Intel® C600 series chipset-A with no upgrade ROM	4 ports	SATA Only	Yes	-A (1D6Bh)
1	4 ports	SATA/SAS	No	-B (1D69h)
2	4 ports	SATA/SAS	Yes	(1D65h)
5	8 ports	SATA/SAS	No	-D (1D68h)
6	8 ports	SATA/SAS	Yes	-T (1D60h)
9	8 ports	SATA Only	Yes	(1D6Ah)

25.1.4 Drivers Supporting the Intel® C600 series chipset

To support the Intel® C600 series chipset based platforms,

- There are two drivers for the SATA (AHCI) Controller. The first driver supports the SATA controller in AHCI mode and a second driver supports RAID mode.

  NOTE: The system BIOS Configuration Setup utility is used to select either AHCI or RAID modes for the SATA controller

- The SCU Controller will have its own driver that supports both pass-through and RAID configurations. This driver will support/manage/control SAS and/or SATA devices connected to the SCU controller.

  NOTE: The SCU controller does not currently support a non-RAID mode of operation analogous to “AHCI Mode” on the AHCI/SATA controllers



- Graphical User Interface (Intel® RSTe GUI) – This is an application that can be used to manage RAID arrays and volumes on drives attached to the SATA and SCU controllers.

  NOTE: RAID spanning across controllers is not supported.

25.2 SCU OEM Parameters

25.2.1 Intel® C600 series chipset OEM Parameters

The Storage Controller Unit (SCU) complies with the Serial Attached SCSI (SAS) Specification's Physical and Link Layer definitions. However, when it comes to actually implementing motherboard designs, OEMs often encounter challenges that pit high-frequency design best-practices against form factor and trace lengths. Therefore, SCU provides a mechanism for OEMs to tweak its PHY parameters to find the most optimal settings for a given platform. Once determined, the OEM can capture these settings to a binary file that is placed into the platform's Serial Peripheral Interconnect (SPI) flash. If the OEM has multiple motherboard platforms using the Intel® C600 series chipset, each platform could have its own unique PHY parameter settings if needed.

SAS is a connection based protocol. Thus, the SCU requires that a valid SAS address be assigned to each of its controllers. The OEM registers these addresses through a proper naming authority. So, in addition to the PHY settings, SAS addresses are placed in the OEM parameter block as well.

SAS addresses are not unique to platforms as are the PHY settings. They are unique to each SCU controller in an individual Intel® C600 series chipset unit (i.e. SKU) similar to the way MAC addresses are unique to each LAN controller.

Since the OEM parameters are unique to platforms and chipset SKUs, they are set once at manufacturing time. OEM parameters will have been validated by OEMs with their respective platforms. Therefore, end users should not be allowed to change these settings for fear of system compatibility problems.

The remaining subsections discuss the details of the OEM parameter block and the requirements they place on the BIOS.



25.3 Structure of OEM Parameter Block

The OEM parameter block is 512 bytes total in size and an unpacked, unpadded binary block. It is partitioned into a header and up to two SCU controller element structures. Figure 4 shows this structure.

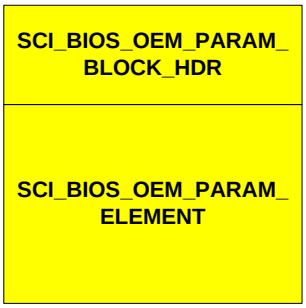


Figure 4: Structure of OEM Parameter Block

The following is the C programming language definition of the OEM parameter block. The comments give a full description of each field.

```
// For Intel Storage Controller Unit OEM Block
#define SCI_OEM_PARAM_SIGNATURE    "ISCUOEMB"

#define SCI_PREBOOT_SOURCE_INIT      (0x00)
#define SCI_PREBOOT_SOURCE_OROM      (0x80)
#define SCI_PREBOOT_SOURCE_EFI       (0x81)

#define SCI_OEM_PARAM_VER_1_0        (0x10)
#define SCI_OEM_PARAM_VER_1_1        (0x11)
#define SCI_OEM_PARAM_VER_1_4        (0x14)


// current version
#define SCI_OEM_PARAM_VER_CUR     SCI_OEM_PARAM_VER_1_4

// port configuration mode
#define SCI_BIOS_MODE_MPC   (0)
#define SCI_BIOS_MODE_APC   (1)


#ifndef SCI_MAX_PHYS
#define SCI_MAX_PHYS (4)
#endif
```



```
#ifndef SCI_MAX_PORTS
#define SCI_MAX_PORTS (4)
#endif


/**
 * @struct SCI_BIOS_OEM_PARAM_BLOCK_HDR
 *
 * @brief This structure defines the OEM Parameter block header.
 */
typedef struct SCI_BIOS_OEM_PARAM_BLOCK_HDR
{
   /**
    * This field contains the OEM Parameter Block Signature which is
    * used by BIOS and driver software to identify that the memory location
    * contains valid OEM Parameter data.  The value must be set to
    * SCI_OEM_PARAM_SIGNATURE which is the string "ISCUOEMB" which
    * stands for Intel Storage Controller Unit OEM Block.
    */
   U8 signature[8];
   /**
    * This field contains the size in bytes of the complete OEM
    * Parameter Block, both header and payload hdr_length +
    * (num_elements * element_length).
    */
   U16 total_block_length;
   /**
    * This field contains the size in bytes of the
    * SCI_BIOS_OEM_PARAM_BLOCK_HDR. It also indicates the offset from
    * the beginning of this data structure to where the actual
    * parameter data payload begins.
    */
   U8 hdr_length;
   /**
    * This field contains the version info defining the structure
    * of the OEM Parameter block.
    */
   U8  version;
   /**
    * This field contains a value indicating the preboot initialization
    * method (Option ROM or UEFI driver) so that after OS transition,
    * the OS driver can know the preboot method. OEMs who build a single
    * flash image where the preboot method is unknown at manufacturing
    * time should set this field to SCI_PREBOOT_SOURCE_INIT.  Then
    * after the block is retrieved into host memory and under preboot
    * driver control, the OROM or UEFI driver can set this field
    * appropriately (SCI_PREBOOT_SOURCE_OROM and SCI_PREBOOT_SOURCE_EFI,
```



```
   * respectively).
   */
  U8 preboot_source;
  /**
   * This field contains the number of parameter descriptor elements
   * (i.e. controller_elements) following this header.  The number of
   * elements corresponds to the number of SCU controller units contained
   * in the platform:
   *  controller_element[0] = SCU0
   *  controller_element[1] = SCU1
   */
  U8 num_elements;
  /**
   * This field contains the size in bytes of the descriptor element(s)
   * in the block.
   */
  U16 element_length;

  /**
  * This field contains the BIOS Controlled Feature Set and allows
  * BIOS vendors to enable/disable specific RSTe product features.
   */
  U8 BCFS;
```

NOTE: Please see the section Error! Reference source not found. Error! Reference source not found. **below for details on the allowed settings.**

```
  /**
   * Reserve fields for future use.
   */
  U8 reserved[6];

} SCI_BIOS_OEM_PARAM_BLOCK_HDR_T;


/**
 * @struct SCIC_SDS_OEM_PARAMETERS
 *
 * @brief This structure delineates the various OEM parameters that must
 * be set for the Intel SAS Storage Controller Unit (SCU).
 */
typedef struct SCI_BIOS_OEM_PARAM_ELEMENT
{
  /**
   * Per SCU Controller Data
   */
  struct
```



```
{
  /**
   * This field indicates the port configuration mode for
   * this controller:
   *   Automatic Port Configuration(APC) or
   *   Manual Port Configuration (MPC).
   *
   * APC means the Platform OEM expects SCI to configure
   * SAS Ports automatically according to the discovered SAS
   * Address pairs of the endpoints, wide and/or narrow.
   *
   * MPC means the Platform OEM manually defines wide or narrow
   * connectors by apriori assigning PHYs to SAS Ports.
   *
   * By default, the mode type is APC
   * in APC mode, if ANY of the phy mask is non-zero,
   *    SCI_FAILURE_INVALID_PARAMETER_VALUE will be returned
   *      from scic_oem_parameters_set AND the default oem
   *      configuration will be applied
   * in MPC mode, if ALL of the phy masks are zero,
   *    SCI_FAILURE_INVALID_PARAMETER_VALUE will be returned
   *      from scic_oem_parameters_set AND the default oem
   *      configuration will be applied
   */
  U8  mode_type;

  /**
   * This field specifies the maximum number of direct attached
   * devices the OEM will allow to have powered up simultaneously
   * on this controller.  This allows the OEM to avoid exceeding
   * power supply limits for this platform.  A value of zero
   * indicates there are no restrictions.
   */
  U8  max_number_concurrent_device_spin_up;

   /**
   * This bitfield indicates the OEM's desired default Tx
   * Spread Spectrum Clocking (SSC) settings for SATA and SAS.
   * NOTE: Default SSC Modulation Frequency is 31.5KHz.
   *-----------------------------------------------------------*/
  /**
   * NOTE: Max spread for SATA is +0 / -5000 PPM.
   * Down-spreading SSC (only method allowed for SATA):
   *  SATA SSC Tx Disabled                 = 0x0
   *  SATA SSC Tx at +0 / -1419 PPM Spread    = 0x2
   *  SATA SSC Tx at +0 / -2129 PPM Spread    = 0x3
   *  SATA SSC Tx at +0 / -4257 PPM Spread    = 0x6
```



```
     *  SATA SSC Tx at +0 / -4967 PPM Spread    = 0x7
     */
    U8 ssc_sata_tx_spread_level : 4;

    /**
     * SAS SSC Tx Disabled                  = 0x0
     *
     * NOTE: Max spread for SAS down-spreading +0 / -2300 PPM
     * Down-spreading SSC:
     *  SAS SSC Tx at +0 / -1419 PPM Spread     = 0x2
     *  SAS SSC Tx at +0 / -2129 PPM Spread     = 0x3
     *
     * NOTE: Max spread for SAS center-spreading +2300 / -2300 PPM
     * Center-spreading SSC:
     *  SAS SSC Tx at +1064 / -1064 PPM Spread  = 0x3
     *  SAS SSC Tx at +2129 / -2129 PPM Spread  = 0x6
     */
    U8 ssc_sas_tx_spread_level : 3;
    /**
     * NOTE: Refer to the SSC section of the SAS 2.x Specification
     * for proper setting of this field. For standard SAS Initiator
     * SAS PHY operation it should be 0 for Down-spreading.
     * SAS SSC Tx spread type:
     *  Down-spreading SSC      = 0
     *  Center-spreading SSC    = 1
     */
    U8 ssc_sas_tx_type : 1;
    /*---------------------------------------------------------------------*/


    U8 reserved;

} controller;

/**
 * Per SAS Port data.
 */
struct
{
    /**
     * This field specifies the phys to be contained inside a port.
     * The bit position in the mask specifies the index of the phy
     * to be contained in the port.  Multiple bits (i.e. phys)
     * can be contained in a single port:
     *      Bit 0 = This controller's PHY index 0     (0x01)
     *      Bit 1 = This controller's PHY index 1     (0x02)
     *      Bit 2 = This controller's PHY index 2     (0x04)
```



```
     *     Bit 3 = This controller's PHY index 3     (0x08)
     *
     * Refer to the mode_type field for rules regarding APC and MPC mode.
     * General rule: For APC mode phy_mask = 0
     */
     U8 phy_mask;

  } ports[SCI_MAX_PORTS]; // Up to 4 Ports per SCU controller unit

  /**
   * Per PHY Parameter data.
   */
  struct
  {
     /**
      * This field indicates the SAS Address that will be transmitted on
      * this PHY index.  The field is defined as a union, however, the
      * OEM should use the U8 array definition when encoding it to ensure
      * correct byte ordering.
      *
      * NOTE: If using APC MODE, along with phy_mask being set to ZERO, the
      * SAS Addresses for all PHYs within a controller group SHALL be the
      * same.
      */
     union
     {
        /**
         * The array should be stored in little endian order.  For example,
         * if the desired SAS Address is 0x50010B90_0003538D, then it
         * should be stored in the following manner:
         *    array[0] = 0x90
         *    array[1] = 0x0B
         *    array[2] = 0x01
         *    array[3] = 0x50
         *    array[4] = 0x8D
         *    array[5] = 0x53
         *    array[6] = 0x03
         *    array[7] = 0x00
         */
        U8 array[8];
        /**
         * This is the typedef'd version of the SAS Address used in
         * the SCI Library.
         */
        SCI_SAS_ADDRESS_T  sci_format;

     } sas_address;
```



```
        /**
         * These are the per PHY equalization settings associated with the the
         * AFE XCVR Tx Amplitude and Equalization Control Register Set
         * (0 thru 3).
         *
         * Operational Note: The following Look-Up-Table registers are engaged
         * by the AFE block after the following:
         *  - Software programs the Link Layer AFE Look Up Table Control
         *    Registers (AFE_LUTCR).
         *  - Software sets AFE XCVR Tx Control Register Tx Equalization
         *    Enable bit.
         */
        /**
         * AFE_TX_AMP_CTRL0.  This register is associated with AFE_LUTCR
         * LUTSel=00b. It contains the Tx Equalization settings that will be
         * used if a SATA 1.5Gbs or SATA 3.0Gbs device is direct-attached.
         */
        U32 afe_tx_amp_control0;

        /**
         * AFE_TX_AMP_CTRL1.  This register is associated with AFE_LUTCR
         * LUTSel=01b. It contains the Tx Equalization settings that will
         * be used if a SATA 6.0Gbs device is direct-attached.
         */
        U32 afe_tx_amp_control1;

        /**
         * AFE_TX_AMP_CTRL2.  This register is associated with AFE_LUTCR
         * LUTSel=10b. It contains the Tx Equalization settings that will
         * be used if a SAS 1.5Gbs or SAS 3.0Gbs device is direct-attached.
         */
        U32 afe_tx_amp_control2;

        /**
         * AFE_TX_AMP_CTRL3.  This register is associated with AFE_LUTCR
         * LUTSel=11b. It contains the Tx Equalization settings that will
         * be used if a SAS 6.0Gbs device is direct-attached (which will only run
        * at SAS 3.0Gbs ) .
         */
        U32 afe_tx_amp_control3;

    } phys[SCI_MAX_PHYS];  // 4 PHYs per SCU controller unit

} SCI_BIOS_OEM_PARAM_ELEMENT_T;
```



```
/**
 *  @struct SCI_BIOS_OEM_PARAM_BLOCK
 *
 * @brief This structure defines the OEM Parameter block as it will be stored
 *  in the last 512 bytes of the PDR region in the SPI flash.  It must be
 *  unpacked or pack(1).
 */
typedef struct SCI_BIOS_OEM_PARAM_BLOCK
{
   /**
    * OEM Parameter Block header.
    */
   SCI_BIOS_OEM_PARAM_BLOCK_HDR_T  header;

   /**
    * Per controller element descriptor containing the controller's
    * parameter data. The prototype defines just one of these descriptors,
    * however, the actual runtime number is determined by the num_elements
    * field in the header.
    */
   SCI_BIOS_OEM_PARAM_ELEMENT_T controller_element[1];

} SCI_BIOS_OEM_PARAM_BLOCK_T;
```

25.3.1 Recommended Location in SPI Flash

Figure 5 shows the recommended location where the OEM parameter block should be placed in SPI flash. Intel reference images will be formatted this way, and the FITC tool supplied by Intel for constructing SPI flash images is configured to take a formatted OEM parameter block and place it appropriately in the Platform Data Region (PDR) as shown in the figure. The OEM Parameter Block should be placed in the last 512 bytes of the PDR region aligned on a 32-bit boundary. There will be more discussion on this in section *25.3.2 OEM Parameter and SPI Flash Tools*.

Note: The recommended SPI flash format is descriptor mode.

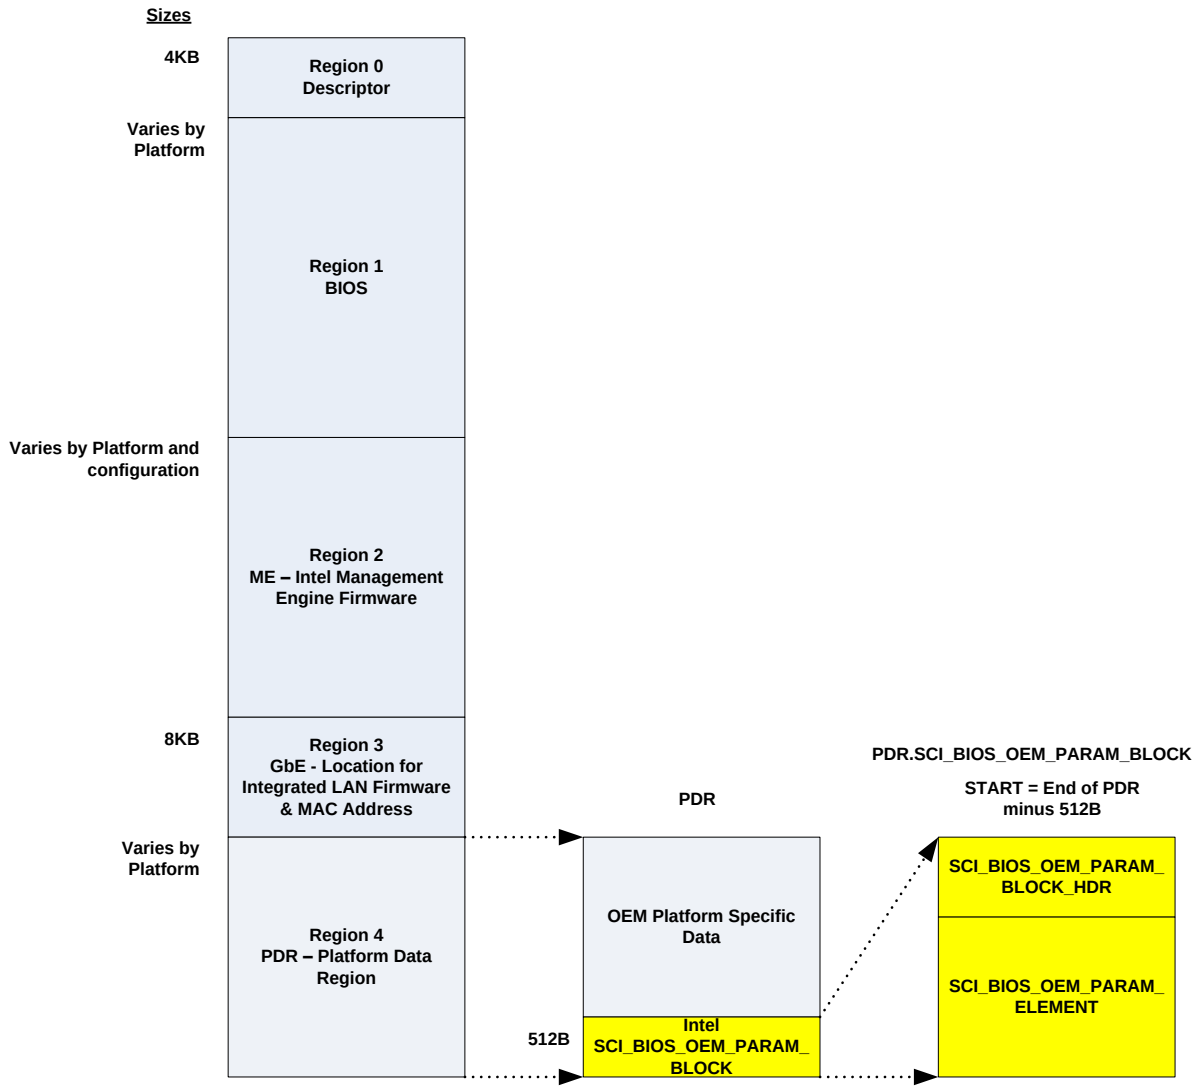


Figure 5 Recommended OEM Param Location in SPI Flash

Third party BIOS vendors are not obligated to follow this recommendation. For the Romley platform, if 3rd parties want Intel provided pre-OS SCU drivers (legacy OROM and UEFI drivers) to boot their systems, they are required to provide a valid OEM Parameter block in flash somewhere. Further, they are required to implement the access services defined in section *25.3.2.4 Required BIOS Services* so that legacy OROM and/or a UEFI driver can retrieve these parameters and configure the SCU.



These BIOS services abstract the Intel provided pre-OS drivers from having to know where the parameters actually reside in flash.

25.3.2 OEM Parameter and SPI Flash Tools

As part of the SCU development kit, tools are provided to help Original Equipment Manufacturers (OEMs), System Integrators, and Value Added Resellers (VARs) tune SCU PHY parameters to a platform's motherboard design and mass produce SPI flash images for that platform. The names of the tools are as follows:

- PHY Tune Tool (part of SCU development kit)
- SAS Address Tool (part of SCU development kit)
- FITC Tool (part of SPI Flash Programming kit)

These tools can be found by contacting your Intel FAE. Detailed user guides are provided in each of the respective kits. However, the following subsections do provide a brief overview of each tool.

Except for the PHY SAS Addresses, the OEM parameters are tuned and fixed for a given platform by the OEM. SAS Addresses, however, are unique to individual SCU SKUs similar to the MAC Address for LAN chips. The system OEM must provide valid SAS addresses that are registered to them. Refer back to section *25.3 Structure of OEM Parameter Block* for a detailed description of the PHY parameters.

25.3.2.1 PHY Tune Tool

PHY Tune is a Windows .Net 4.0 application. It provides a GUI front-end to display and configure SCU PHY information on the console. The tool communicates with a special SCU driver packaged with the tool. The tool is designed to run on Windows Server 2008 R2.

Note: The driver that comes with the PHY Tune Tool is part of an Intel Confidential toolset. That toolset is strictly for the use of the entity with an NDA with Intel and no part may be distributed.

In addition, PHY Tune retrieves and displays special diagnostic information from the SCU. From this information, OEMs can determine which settings provide the most optimal PHY performance for that platform.

Once the OEM determines the appropriate parameter settings for the platform, they can use a PHY Tune menu option to capture these settings. PHY Tune also allows the OEM to set initial values for the other fields in the OEM parameter block



including the header fields and the PHY SAS Addresses. PHY Tune provides another menu option to export these settings to a properly formatted binary file. This block can then be loaded into the Intel FITC tool recognizes where to place this block in constructing a SPI flash image. If the OEM uses its own tools other than FITC to construct the flash image, then it just operates on the "raw" binary file.

Note: If FITC is used for flash image construction, it will place the block in the PDR region as described in section 25.3.1 Recommended Location in SPI Flash.

OEM fields initialized by PHY Tune that need to change per SKU (e.g. SAS Addresses) can be modified later by other manufacturing tools.

25.3.2.2 SAS Address Tool (sasaddresses.efi)

The developer kit includes a UEFI application providing a command line interface that can be used by a manufacturing script to update per SKU SAS addresses. The name of the application is ***sasaddresses.efi*** . Based on input parameters, the application will either work directly on an OEM block already in the PDR of SPI flash or it will operate on a properly formatted binary input file. In the case of the binary input file, it will modify the SAS Addresses in the binary image according to the SAS addresses on the command line and then overwrite the OEM section of PDR with the binary image.

25.3.2.3 FITC Tool

The FITC tool is part of the SPI Flash Programming kit. It is used to construct flash images and modify data at manufacturing using an XML schema. The Intel FITC tool constructs a SPI flash image in descriptor mode containing the follow regions:

- Descriptor Region indicating how the flash is partitioned and where other regions begin and end.
- BIOS Region
- Intel Manufacturing Engine Region
- Gigabit Ethernet Region
- Platform Data Region

When using the FITC tool to construct a SPI flash image, FITC will place the OEM parameter block in the last 512 bytes of the Platform Data Region as described in section 25.3.1 Recommended Location in SPI Flash.



25.3.2.4 Required BIOS Services

System BIOS shall provide services to extract the OEM parameter block from SPI Flash. As mentioned earlier, the recommended location for the OEM parameter block is in the last 512 bytes of the PDR region as described in section 25.3.1 Recommended Location in SPI Flash.

The following subsections describe the necessary services according to the type of pre-boot driver environment whether legacy Option ROM or native UEFI.

25.3.2.4.1 Legacy OROM

Under legacy BIOS / OROM mode, the BIOS must provide the following software interrupt services.

25.3.2.4.2 Get RSTe OROM SCU OEM Parameter Block: INT 15, Function=F300h, Sub-Function=0001h

INT 15 / AX=F300h / BX=0001h (Get RSTe OROM SCU OEM Parameter Block)

Description:

Through this function, BIOS provides to the RSTe legacy OROM driver the OEM Parameter Block needed to initialize the SCU controller.

Inputs:

AX = F300h (Function)

BX = 0001h (Sub Function)

EAX = 0000F300h

EBX = 4F450001h ('OE') + Sub Function

EDX = 20534355h (' SCU')

ECX = size of data buffer in bytes

(512 Bytes = OEM header + Descriptor Elements)

EDI = 0000xxxxh (Upper 16bits are zero, lower 16bits defined below)

ES:DI = address of data buffer (Real mode address)



Normal Outputs:

CF = clear if successful

EAX = 20534355h (' SCU')

ES:DI = data buffer filled

ECX = actual transfer size in bytes (512 Bytes)

Data Format:

= char oem_params[512];

(512 Bytes of data as encoded by OEM or system integrator)

Error Outputs:

CF = set on error

AH = error code

= 86h Function Not Supported

= 87h OEM Block Not Present

25.3.2.4.3 INT 15 / AX=F300h / BX=0002h (Get RSTe OROM Boot Info)

Description:

Through this function, BIOS provides user-settable RSTe boot information to the RSTe legacy OROM driver. These values are visible to the user through the BIOS Setup menus. The menu options should be linked to legacy OROM selections in the PCH-IO section.

Inputs:

AX = F300h (Function)

BX = 0002h (Sub Function)

EAX = 0000F300h



EBX = 4F450002h ('OE') + Sub Function

EDX = 424F4F54h ('BOOT')

Normal Outputs:

CF = clear if successful

EAX = 424F4F54h ('BOOT')

BL = legacy_orom_boot_controller_selection:

Due to limited shadow RAM and EBDA space, and the fact that a platform may require multiple OROMs be loaded for other functions, there might not be enough runtime space for both the RSTe SATA RAID controller OROM and the RSTe SCU RAID controller OROM to provide int13h support simultaneously. Even so, each RSTe OROM still needs to initialize so that it can configure each controller based on platform dependencies and store data needed by the OS drivers in the Shadow RAM area even if it does not provide full int13h runtime support. Thus, through this setup option BIOS can avoid the runtime space conflict by allowing the user to select the boot controller according to the following values:

0000h = No runtime space restrictions. BIOS indicates that both RSTe SATA and SCU runtime code should provide full int13h support for RSTe devices.

(NOTE: The BIOS should allow this option if it knows that there is room in shadow RAM for both OROMs' runtime code. If the BIOS can always guarantee this condition, then it does NOT need to make Legacy OROM boot controller selection visible to the user in BIOS setup.)

0001h = The SCU controller is selected as boot controller. BIOS will load RSTe SATA OROM first, but the SATA OROM will only initialize and then leave pertinent RAID configuration information for the SATA OS RAID Driver in runtime space. The RSTe SCU OROM will initialize, relocate to runtime space, and provide full int13h support for SCU attached devices.

0002h = The SATA controller is selected as boot controller. BIOS will load RSTe SCUOROM first, but the SCU OROM will only initialize and then leave pertinent RAIDconfiguration and SCU OEM parameter



information for the SCU OS RAID Driver in runtime space. The RSTe SATA OROM will initialize, relocate to runtime space, and provide full int13h support for SATA controller attached devices.

0003h = Neither SATA nor SCU controller is selected as boot controller. Boot support is being provided by another device. BIOS will load both RSTe OROMs, but each will only initialize and leave pertinent RAID configuration and SCU OEM parameter information for the RSTe OS RAID Drivers in runtime space. Therewill NOT be int13h support for RSTe devices.

BH = scu_legacy_orom_max_disk_slots_enum:

For boot_controller_selection = {0000h or 0001h}, this option allows the user to set the maximum number of disk slots the SCU legacy OROM will enumerate. Its range will be between 1 and 8.

The default value shall be 8. For boot_controller_selection = {0002h or 0003h}, this field should not be settable by the user, and OROM will automatically assume a value of 0.

Error Outputs:

CF = set on error

AH = error code

= 86h Function Not Supported = (boot_controller_selection = 0000h assumed)

25.3.2.5 UEFI

Under native UEFI mode, the BIOS must provide the following protocol.

```
//
// Define SCU Parameters protocol GUID
//
// EDK and EDKII have different GUID formats
//
#if !defined(EDK_RELEASE_VERSION) || (EDK_RELEASE_VERSION < 0x00020000)
#define EFI_PCH_SCU_PARAMETERS_PROTOCOL_GUID \
```



```
 { \
   0xe165e866, 0x6643 , 0x40b3, 0xb4, 0x35, 0x52, 0x6b, 0x47, 0x3f, 0x75, 0xc2 \
 }

#else
#define EFI_PCH_SCU_PARAMETERS_PROTOCOL_GUID \
 { \
   0xe165e866, 0x6643 , 0x40b3, \
   {0xb4, 0x35, 0x52, 0x6b, 0x47, 0x3f, 0x75, 0xc2} \
 }
#endif
#define SCU_PARMS_SIZE 512
//
// Protocol definition
//
struct _PCH_SCU_PARAMETERS_PROTOCOL {
  U8 SCUParameters[SCU_PARMS_SIZE];
};
```

End



26 Storage System Events Detected by Monitor Service (IAStorDataMgrSvc)

Event Type	Event Level	String	Event Displayed		
			NAI[1] (Notification Area Icon)	Event Log	E-Mail Notify[2]
Disk Triggered Events					
Failed	Error	Disk on port {n}: Failed. Open the application for details.	Yes	Yes	Yes
S.M.A.R.T.	Warning	Disk on port {n}: At risk. Open the application for details.	Yes	Yes	Yes
Unlocked	Info	Disk on port {n}: Unlocked.	Yes	Yes	Yes
Added	Info	Disk on port {n}: Detected.	Yes	Yes	Yes
Removed	Info	Disk on port {n}: Removed.	Yes	Yes	Yes
Volume Triggered Events					
Failed	Error	Volume {0}: Failed. Open the application for details.	Yes	Yes	Yes
Degraded	Warning	Volume {0}: Degraded. Open the application for details.	Yes	Yes	Yes
Detected	Info	A new volume was found.	Yes	Yes	Yes
RebuildComplete	Info	Volume {0}: Rebuilding complete.	Yes	Yes	Yes
VerifyStop	Info	Volume {0}: Verification complete.	Yes	Yes	Yes
VerifyAndRepairStop	Info	Volume {0}: Verification and repair complete.	Yes	Yes	Yes
MigrationComplete	Info	Volume {0}: Data migration complete.	Yes	Yes	Yes
InitializeComplete	Info	Volume {0}: Initialization complete.	Yes	Yes	Yes
Unlocked	Info	Volume {0}: Unlocked.	Yes	Yes	Yes
NotPresent	Info	Volume {0}: No longer present on system.	Yes	Yes	Yes
RebuildStarted	Info	Volume {0}: Rebuilding in progress.	Yes	No	No
VerifyStarted	Info	Volume {0}: Verification in progress.	Yes	No	No
VerifyAndRepairStarted	Info	Volume {0}: Verification and repair in progress.	Yes	No	No



Event Type	Event Level	String	Event Displayed		
			NAI[1] (Notification Area Icon)	Event Log	E-Mail Notify[2]
MigrationStarted	Info	Volume {0}: Data migration in progress.	Yes	No	No
InitializeStarted	Info	Volume {0}: Initialization in progress.	Yes	No	No
General Events					
Server start failed	Error	Server failed to start. Additional information:	No	Yes	Yes
Event manager started	Info	Started the event manager.	No	Yes	Yes

NOTES:

1. NAI true only if the user selected to receive notification under Preferences in the UI
2. Refer to Email Section above for Email feature support