

Intel® Rapid Storage Technology enterprise (RSTe)

Technical Product Specification (TPS) Document

August, 2016

Revision 2.2

Intel Confidential

Document Number: 545322



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.

UNLESS OTHERWISE AGREED IN WRITING BY INTEL, THE INTEL PRODUCTS ARE NOT DESIGNED NOR INTENDED FOR ANY APPLICATION IN WHICH THE FAILURE OF THE INTEL PRODUCT COULD CREATE A SITUATION WHERE PERSONAL INJURY OR DEATH MAY OCCUR.

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

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

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

Copies of documents which have an order number and are referenced in this document, or other Intel literature, may be obtained at: http://www.intel.com/design/literature.htm

All products, platforms, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice. All dates specified are target dates, are provided for planning purposes only and are subject to change.

This document contains information on products in the design phase of development. Do not finalize a design with this information. Revised information will be published when the product is available. Verify with your local sales office that you have the latest datasheet before finalizing a design.

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

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

Copyright © 2016, Intel Corporation. All rights reserved.



Contents

1 Introduction 8

1.1 Intended Use 10

1.2 Intended Audience 10

2 Product Overview 11

2.1 Supported Platforms/Chipsets/SKUs 12

2.1.1 Intel® RSTe PCH/SATA Support 12

2.1.2 Intel® RSTe PCH/SATA and NVMe Support 12

2.1.3 Intel® C230 series chipset (Greenlow) Platforms 12

2.2 Supported Chipset SKU Overview 13

2.2.1 Intel® C600 series chipset SKU's 13

2.2.2 Intel® C610 series chipset SKU's 14

2.2.3 Intel® C220/230 series chipset SKU's 14

2.3 Intel® RSTe PCH/SATA Product Major Components 15

2.3.1 Intel® RSTe PCH/SATA Pre-OS Images 15

2.3.2 Intel® RSTe OS Runtime Drivers 16

2.3.3 Intel® RSTe RAID Tools and Utility 17

2.4 Intel® RSTe NVME Product Major Components 18

2.4.1 Intel® RSTe NVMe OS Runtime Drivers 18

2.4.2 Intel® RSTe NVMe OS Tools and Utilities 18

2.5 Unsupported Features 19

2.6 Intel® RSTe PCH/SATA Key Product Features 20

2.6.1 Operating System Supported 20

2.6.2 Intel® RSTe PCH/SATA Pre-OS Package 21

2.6.3 Intel® RSTe PCH/SATA Configuration Tools 25

2.6.4 Intel® RSTe System Configurations supported 27

2.6.5 Software RAID Functional Support 30

2.6.6 Utilities 39

2.7 Intel® RSTe NVMe Key Product Features 40

2.7.1 Operating System Support 40

2.7.2 Intel® RSTe NVMe Configuration Tools 41

2.7.3 Intel® RSTe System Configurations supported 41

2.7.4 Software RAID Functional Support 42

2.7.5 Utilities 50

3 Product Certifications 51

3.1 WHQL 51

4 Appendix A 52



5 Appendix B 55

6 Appendix C 57

6.1 Intel® C600 series chipset OEM Parameters 57

6.2 Structure of OEM Parameter Block 57

6.3 BCFS Bit Settings 65

6.4 Recommended Location in SPI Flash 67

6.5 OEM Parameter and SPI Flash Tools 69

6.5.1 PHY Tune Tool 69

6.5.2 SAS Address Tool (sasaddresses.efi) 70

6.5.3 FITC Tool 70

6.6 Required BIOS Services 70

6.6.1 Legacy OROM 70

6.6.2 UEFI 73

6.7 Intel® C610 series Chipset Legacy OROM Boot Option 74

6.7.1 INT 15 / AX=F300h / BX=0002h (Get RSTe OROM Boot Info) 74

7 Appendix D 76

7.1 Hardware Compatibly 76

7.1.1 External Hardware Compatibility 76



Figures

Figure 1: Structure of OEM Parameter Block 58
Figure 2 Recommended OEM Param Location in SPI Flash 68

Tables

Table 1-1. Common Acronyms 8
Table 5-1. Relevant Specifications 52
Table 5-2. Relevant Documents 52
Table 6-1. Storage System Events Detected by Monitor Service (IAStorDataMgrSvc) 55



Revision History

Document Number	Revision Number	Description	Revision Date
	0.36	Initial release.	Aug 10, 2010
458224	0.5	Pre-Alpha Release	Sept 21. 2010
458224	0.51	Added descriptions for Dirty Strip Journaling, Partial Parity Logging and Verify and Repair features	Sept 22, 2010
458224	0.65	Minor Corrections and Clarifications	December, 2010
458224	0.70	Alpha1 Release	February, 2011
458224	0.71	Update SKU information and silicon support info	March, 2011
458224	0.72	Alpha Release	April, 2011
458224	0.73	Minor Corrections and Clarifications	April, 2011
458224	0.74	Updated the SKU'ing information as well as added documents to Relevant Document table in Appendix A	April, 2011
458224	0.80	Beta Release	June, 2011
458224	0.81	Pre-PC Release	August, 2011
458224	0.82	Update Pre-OS requirements for OROMs Updated Upgrade ROM #9 Device ID	Sept, 2011
458224	0.84	Updated Supported Operating Systems	Nov, 2011
458224	0.85	Updated Hardware Compatibility List	Dec, 2011
458224	1.0	Updated to correct various mistakes and updates.	March, 2012
458224	1.1	Updated to correct various mistakes and updates.	June, 2012
458224	1.2	Updated information regarding expanders	July, 2012
458224	1.3	Updated information on the Supported OSs	August, 2012
458224	1.4	SCU Legacy OROM support for Boot Behind Expander added	November, 2012
458224	1.5	Added new features introduced in RSTe 3.7	April, 2013
458224	1.6	Adding new requirements for RSTe 3.8	May, 2013
545322	1.8	Adding references to Grantley and Denlow Platforms	March, 2014
545322	1.9	Minor Corrections	June, 2014
545322	1.12	Updated for RSTe 4.3 release	June, 2015
545322	2.0	Updated for RSTe 4.3 NVMe release	July, 2015
545322	2.1	Updated with Hot-Plug support	February, 2016
545322	2.2	Updated to include Windows* Server 2016	August 2016



§



1 Introduction

This is the Technical Product Specifications for the Intel® Rapid Storage Technology enterprise (Intel® RSTe) family of products. The intent of this document is to present the Intel® RSTe functional requirements (or features) that make up these products. The features addressed include those for the Pre Operating System (Pre-OS) components, the Intel® RSTe Windows* based drivers and tools and * Graphical User Interface (GUI).

The first product is a RAID package that supports the Intel® PCH and is referred to as Intel® RSTe PCH (or RSTe SATA) or RSTe PCH/SATA. The second product is the RAID package that supports Intel® NVMe SSD plugged into PCI express slots controlled/managed by the CPU. This product is referred to as Intel® RSTe NVMe. Both products are enterprise class RAID products designed to operate in enterprise class platforms. ACRONYMS

The following acronyms will be used throughout this document.

Table 1-1. Common Acronyms

Term	Definition
AHCI	Advanced Host Controller Interface
ATA	Advanced Technology Attachment
ATAPI	Advanced Technology Attachment Packet Interface
BIOS	Basic Input / Output System
BBE	Boot Behind Expander
Chipset	A term used to define a collection of The PNHCI components required to make a PC function.
CIM	Common Information Model
CLI	Command Line Interface
Cougar Point	Platform Controller Hub
CSMI	Common Storage Management Interface
DMA	Direct Memory Access
DOS	Disk Operating System
DIPM	Device Initiated Power Management
DSJ	Dirty Stripe Journaling
Disk's Write Cache	A memory device within a hard drive, which is allocated for the temporary storage of data before that data is copied to its permanent storage location.
EBDA	Extended BIOS Data Area
EN	Entry Server

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



Term	Definition
EP	Efficient Performance
GB	Giga-byte
GUI	Graphical User Interface
HDD	Hard Disk Drive
HIPM	Host Initiated Power Management
HII	Human Interface Infrastructure
Hot Plug	A term used to describe the removal or insertion of a SATA hard drive when the system is powered on.
ICH	Input / Output Controller Hub
IHV	Independent Hardware Vendor
LPM	Link Power Management
MB	Mega-bytes
NAI	Notification Area Icon
NCQ	Native Command Queuing
NTFS	NT File System
NVMe	PCI Express Non Volatile Memory
ODD	Optical Disk Devices
ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
OROM	Option ROM
OS	Operating System
PCH	Platform Control Hub
Pre-OS	Pre Operating System Environment (Legacy OROM and/or UEFI)
Port	The point at which a SATA drive physically connects to the SATA controller.
PPL	Partial Parity Logging
PRD	Product Requirements Document
RHEL	Redhat Enterprise Linux
RRT	Rapid Recover Technology
RST	Rapid Storage Technology
RSTe	Rapid Storage Technology enterprise
SAS	Serial Attached SCSI
SATA	Serial ATA
sSATA	Secondary Serial ATA
SCU	Storage Controller Unit
SES	SCSI Enclosure Service

Term	Definition
SGPIO	Serial General Purpose I/O
SMART	Self-Monitoring, Analysis and Reporting Technology: an open standard for developing hard drives and software systems that automatically monitors a hard drive's health and reports potential problems.
SLES	SUSE Linux Enterpise Server
SMIS	Storage Management Initiative Specification
SSD	Solid State Device – non volatile memory
UI	User Interface
UEFI	Unified Extensible Firmware Interface

1.1 Intended Use

This document is intended to provide high level information on the technical features of the Intel® RSTe family of products (both RSTe PCH/SATA and RSTe NVMe)

1.2 Intended Audience

The intended audience of this document is OEMs and ODMs requiring detailed information of the new features and technical specifications of the Intel® RSTe family of products. It contains only information pertaining to the Intel® RSTe family of products.



2 Product Overview

Intel® Rapid Storage Technology enterprise (RSTe) is a set of products that provides RAID support for the enterprise environment. Some of the RAID features include:

- RAID level 0 (striping)
- RAID level 1 (mirroring)
- RAID level 5 (striping with parity)
- RAID level 10 (striping and mirroring)

Features supported by the RSTe PCH/SATA product include:

- RAID boot support
- LED Management through SGPIO

Features supported by the RSTe NVMe product include:

- Intel NVMe SSD support
- Spanning across NVMe SSD Controllers

The Intel® RSTe PCH/SATA product with support the following Intel® PCHs:

- Intel® C600 series chipset
  - 1 Advanced Host Controller Interface (AHCI)/SATA
  - 1 Storage Controller Unit (SCU)
- Intel® C610 series chipset
  - 1 Advanced Host Controller Interface (AHCI)/SATA
  - 1 secondary Advanced Host Controller Interface (AHCI)/sSATA
- Intel® C220 series chipset
  - 1 Advanced Host Controller Interface (AHCI)/SATA
- Intel® C230 series chipset
  - 1 Advanced Host Controller Interface (AHCI)/SATA

For these platforms, the Intel® RSTe PCH/SATA product package is designed to support Hard Disk Drives (HDDs) and SATA Solid State Drives (SSDs) to enable an optimal enterprise storage solution.

The Intel® RSTe NVMe product will support only Intel® NVMe SSDs directly connected to the PCI express slots controlled by the CPU. Please consult you platform guides to identify which PCI express slots are managed by the CPU.

Please review the entire document for the complete list of features in the Intel® RSTe family of products.

2.1 Supported Platforms/Chipsets/SKUs

The Intel® RSTe product package was designed to work on and had been tested on the Intel Customer Reference Boards (CRBs) outlined in this section.

2.1.1 Intel® RSTe PCH/SATA Support

The following platforms will only support the RSTe PCH/SATA products.

2.1.1.1 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)

2.1.1.2 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

2.1.2 Intel® RSTe PCH/SATA and NVMe Support

The following platforms will support both the RSTe PCH/SATA and RSTe NVMe products.

2.1.2.1 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

2.1.3 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

2.2 Supported Chipset SKU Overview

2.2.1 Intel® C600 series chipset SKU's

2.2.1.1 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.

2.2.1.2 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)

2.2.2 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

2.2.3 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**

2.3 Intel® RSTe PCH/SATA Product Major Components

Intel® RSTe PCH/SATA 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

2.3.1 Intel® RSTe PCH/SATA Pre-OS Images

Intel® RSTe PCH/SATA 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).

2.3.1.1 Intel® RSTe PCH/SATA RAID Legacy Option ROMs

Intel® RSTe PCH/SATA 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 and are include following:

1. Individual images for each of the PCH controllers on the Intel® C600/C610/C220/ and C230 series chipset platforms: There is one RAID Legacy OROM image that supports each of the SATA and sSATA controllers for these PCHs.

2. Individual image for the SCU controller on the Intel® C600 series chipset platforms.

2.3.1.2 Intel® RSTe PCH/SATA RAID Unified Extensible Firmware Interface (UEFI) Drivers

Intel® RSTe PCH/SATA 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 are:

1. Intel® C600/C610/C220/C320 series chipset: There is one RAID UEFI driver image that supports the SATA controller for these PCHs.

2. Intel® C600 series chipset: There is one RAID UEFI driver image that supports the SCU controller for this PCH.

2.3.2 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:

2.3.2.1 Drivers Supporting the Intel® C600 series chipset

1. 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 utility is used to select either AHCI or RAID modes for the SATA controller

2. 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 AHCISATA controllers

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

   NOTE: RAID spanning across controllers is not supported.

2.3.2.2 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.

2.3.2.3 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.

2.3.3 Intel® RSTe RAID Tools and Utility

2.3.3.1 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.

2.3.3.2 Intel® RSTe PCH/SATA 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.

2.3.3.2.1 Intel® RSTe SATA/sSATA RAID Configuration

There are two Intel® RSTe SATA RAID configuration tools. One is a DOS based utility that runs from a DOS bootable USB memory stick. The other is a UEFI tool that runs from a USB memory stick from the platforms UEFI Shell environment. HII tool that is controlled by the system UEFI BIOS. The purpose of these tools is to help identify and manage drives attached to the SATA (AHCI) controller on both the Intel® C600 and C610 series chipsets. There is also a corresponding set of tools that supports the sSATA controller (Intel® C610 series chipset).

2.3.3.2.2 Intel® RSTe SATA/sSATA RAID Compatibility

There are two types of Intel® RSTe SATA RAID compatibility tools. One is a DOS based utility that runs from a DOS bootable USB memory stick. The other is a UEFI tool that runs from a UEFI Shell environment. Both tools need to be loaded/copied to a USB key in order to use in the system being tested. The purpose of these tools is to help verify that the RSTe SATA/sSATA legacy RAID Legacy OROMs and the RSTe SATA RAID UEFI driver have been properly implemented/incorporated into the system BIOS. These tools support both the Intel® C600/C610/C220 and C230 series chipsets.

2.3.3.2.3 Intel® RSTe SCU RAID Configuration

There are two Intel® RSTe SCU RAID configuration tools. One is a DOS based utility that runs from a DOS bootable USB memory stick. The other is a UEFI tool that runs from a UEFI Shell environment. Both tools need to be loaded/copied to a USB key in order to use in the system being tested. The purpose of these tools is to help identify and manage drives attached to the SCU controller on the Intel® C600 series chipsets.

2.3.3.2.4 Intel® RSTe SCU RAID Compatibility

There are two Intel® RSTe SCU RAID compatibility tools. One is a DOS based utility that runs from a DOS bootable USB memory stick. The other is a UEFI tool that runs from the UEFI Shell environment. Both tools need to be loaded/copied to a USB key in order to use in the system

being tested. The purpose of these tools is to help verify that the RSTe SCU legacy RAID OROM and the RSTe SCU RAID UEFI driver have been properly implemented/incorporated into the system BIOS. This also verifies that the OEM Parameters have been properly programmed into the SPI Flash memory. These tools manage drives attached to the SCU controller on the Intel® C600 series chipsets.

2.4 Intel® RSTe NVME Product Major Components

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

- RSTe OS Runtime Drivers
- RSTe Tools and Utilities

2.4.1 Intel® RSTe NVMe OS Runtime Drivers

The RSTe NVMe OS runtime NVMe drivers are based on the Intel® NVMe device drivers, which is Windows* Storport Miniport device driver. This provides for improved performance along with simplified functional maintainability. There a single Windows* runtime RSTe NVMe driver included that installs for each NVMe device attached to the platform. Each instance of the driver installed is reported to the Intel® RSTe environment at a single drive.

2.4.2 Intel® RSTe NVMe OS Tools and Utilities

The RSTe NVMe OS tools and utilities provided support for the installation of the Intel® RSTe NVMe product and the management of the RAID enabled environment.

2.4.2.1 Intel® RSTe NVMe RAID Installation Tool

The Intel® RSTe NVMe installation tool will support the installation of the Intel® RSTe NVMe product onto supported platforms. This includes the installation of:

1. The Intel® RSTe NVMe device drivers
2. The Intel® RSTe NVMe RAID management layer
3. The Intel® RSTe NVMe Graphical User Interface (GUI), which is similar to the Intel® RSTe PCH/SATA GUI.
4. The Intel® RSTe NVMe GUI will also incorporate with the Intel® RSTe PCH/SATA GUI if already present/installed on the platform

2.4.2.1.1 Intel® RSTe NVMe Command Line Interface (CLI) Utility

The Intel® RSTe NVMe 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® C610 and C230 series chipsets.



2.5 Unsupported Features

The following is a non-comprehensive list of features that are not be supported in Intel® RSTe:

- Web Browser RAID configuration utility
- RAID array and volume(s) Spanning across the SATA (AHCI) and SCU ports
- RAID array and volume(s) spanning across the SATA and sSATA controllers (Intel® C610 series chipset)
- Port-Multipliers
- SATA RAID Legacy OROM support for Optical Devices (Optical Devices are managed by Platform BIOS)
- SCU RAID Legacy OROM support for Optical Devices
- SCU UEFI Driver support for Optical Devices
- Intel® Rapid Recovery Technology (Intel® RRT)
- Hard Drive Password
- Zero Power ODD
- NVSRAM
- Link Power Management (LPM)
- Time Limited Error Recovery (TLER)
- Installing Windows 2003 to a RAID Volume
- Combining all 8 SCU ports into a single X8 Wide Port connection – this configuration will behave as a single X4 Wide Port connection.
- SCU Controller operating at 6 Gb/s – Please refer to the latest C600 series chipset Sightings Report and/or Spec. update
- Windows* 8/Server 2012 (or greater) Operating Systems (OS) “inbox” SATA (RAID Mode) drivers to support installing the OS to a RAID Volume.
- Booting from an NVMe RAID Volume
- NVMe SSD Hot Plug not officially supported with RSTe 4.3 release
- Third party NVMe SSD supports
- NVMe RAID Volumes on NVMe devices attached via the PCH PCI express slots
- Intel RSTe 4.5 supports up to 12 Intel NVMe SSDs

Note: Intel® RSTe NVMe only supports NVMe devices connected PCI express slots controlled/managed by the CPU. If an NVMe device(s) is plugged into PCI express slot(s) controlled/managed by the PCH prior to the installation of the Intel® RSTe NVMe product, the RSTe GUI may report those devices. The Intel® NVMe RSTe GUI does not prevent NVMe RAID volumes from including PCH controlled NVMe SSDs. RAID volumes that include PCH controlled NVMe SSDs will encounter a substantial performance degradation.

2.6 Intel® RSTe PCH/SATA Key Product Features

The product features described in this section, along with additional functionality. The following is a summary of the key features of this product.

Name	Key Features	
RAID	• Support for Matrix RAID 0/1/5/10 • Pass-through drives • Hot Plug with I/O[1] • Hot Spare Disk • Auto Rebuild on Hot Insert[1] • Rebuild & Migration Check Pointing • NCQ (SATA) and CQ (SAS) support • UEFI using common metadata • SAS Expanders (Intel® C600 series chipset only) • SMART Support • Bad Block Management • Controller configuration rules • Drive roaming • RAID Volume roaming between Linux* and Windows* • On Line Capacity Expansion • Large Stripe Size Support • RAID-Ready	• Disk Coercion • Manual & Auto Rebuild • Instant Initialization • Read Patrol • SGPIO • Volume creation/verify • Selectable Boot Volume[1] • Email Alerting • CIM • RAID Level Migration (RAID 0, 1, or 10 to RAID 5) • Dirty Stripe Journaling • Partial Parity Logging (PPL) • Verify and Repair • Auto Rebuild on Hot Insert
Utilities	• Install/Uninstall Utility • Configuration and Management Utilities	
Intel® RSTe (Window only)	Intel® RSTe provides software for high-performance for SAS and SATA RAID + • Matrix RAID for two RAID volumes on single array • Improved user interface for enhanced usability	
1	Limited support with Intel® RSTe NVMe 4.5	

2.6.1 Operating System Supported

2.6.1.1 Microsoft OS Support

Intel® RSTe will support both 32 and 64 bit versions of the following Microsoft OSs:

- Intel® C600/C220 Series Chipset platforms only
  - Windows* Vista (Supported until RSTe 4.0)
  - Windows* Server 2003 (Supported until RSTe 4.0)
    - The “In Box” installation image for Windows 2003 does support the Microsoft Storport Miniport Driver Model. Microsoft made enhancements to the Storport Miniport architecture that was

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



implemented following the release of Windows 2003. Microsoft has made a Hot-Fix available for download. The issue will only be encountered when accessing a RAID volume. Detailed instruction on applying the Hot-Fix to the installation images can be found:

- http://technet.microsoft.com/en-us/library/cc766320(WS.10).aspx

▪ To install Windows 2003:

- Install to a single Pass-thru disk, and then apply the appropriate Hot-Fix to bring the Storport driver up to current standards. The hotfix can be downloaded at: http://support.microsoft.com/kb/932755. Reboot the system and use RST GUI to create the desired RAID Volume from the existing boot drive.

- Windows* 7
- Windows* Server 2008 SP2 (Supported until RSTe 4.0)
- Windows* Server 2008 R2 (64 bit only)
- Windows* PE 3.0
- Windows* 8 (Supported until RSTe 4.0)
- Windows* Server 2012
- Windows* 8.1
- Windows* Server 2012 R2
- Windows* 10
- Windows* Server 2016

2.6.2 Intel® RSTe PCH/SATA Pre-OS Package

The Intel® RSTe PCH/SATA product package will provide two options for the Pre-OS implementation. The first option is the two image option and the second is a single image. The two image option contains separate sets of images for the SATA and sSATA Controller. The single image option has one image that will control both SATA and sSATA controller.

NOTE: For the Intel® C600 series chipset, the only option is the two images. One for the SATA controller and the second image for the SCU controller.

2.6.2.1 Intel® RSTe PCH/SATA Pre-OS Package Two Images Option

For the two image options, the Intel® RSTe PCH/SATA product package will provide 2 sets of binary images that are compiled into the system BIOS to provide the RAID mode Pre-OS driver for the Intel® C600/C610/C220 and C230 series chipsets based platforms. The BIOS vendor (or the OEM/ODMs developing their own BIOS) will take these images and compile them into their system BIOS for these platforms. How they are included and how the information is presented is up to the BIOS vendor (or the OEM/ODMs developing their own BIOS).

---

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

2.6.2.1.1 Intel® RSTe PCH/SATA Basic Requirement for the System BIOS

For Intel® RSTe PCH/SATA product to properly operate, the system BIOS for the Intel® C600/C610/C220 and C230 series chipsets based platforms must properly incorporate the Intel® RSTe Pre-OS images. An example of how a BIOS vendor (or the OEM/ODMs developing their own BIOS) may implement the Intel® RSTe Pre-OS package is by adding the following configuration options to the BIOS configuration:

- The option for the user to select which controller they would like to boot from (SATA, sSATA or SCU). Intel® RSTe will provide support for the ability to perform a minimal installation of the Legacy RAID Option ROM for the controllers not selected as the boot controller. Due to limited shadow RAM and EBDA space, and the fact that many platforms may require multiple Option ROMs be loaded, there might not be enough runtime space for multiple RSTe Legacy RAID Option ROMs to load. To support they condition and INT15 interface has been defined. When the user selects the Boot Controller, the INT 15 interface allows the system BIOS to inform the RSTe Legacy OROM image being loaded if it is the boot controller. When the controller is not the boot controller, it will execute the minimal amount of code needed to pass the required information to the controller and to the OS driver. Please reference the platform's BIOS Writer's Guide for more information (or the Release notes included with this package).
- The option to select which pre-boot environment to use (Legacy Option ROM or UEFI). The Intel® RSTe Pre-boot package will support only one boot method of either Legacy RAID OROM or RAID UEFI, but not both.
- The option to select which mode to configure the Intel® C600/C610/C220 and C230 series chipset SATA and/or sSATA controller (AHCI or RAID mode)
    - When AHCI mode is selected the BIOS's own AHCI int13 driver or (UEFI driver) will be used instead of the Intel® RSTe SATA RAID Legacy Option ROM (or Intel® RSTe UEFI driver). In this mode, the Intel® RSTe SATA OS driver will function as a pass-through driver. No RAID functionality will be allowed on the SATA controllers.
    - When RAID mode is selected, the Intel® RSTe SATA and/or sSATA RAID Legacy Option ROM (or Intel® RSTe RAID SATA and/or sSATA UEFI driver) will be loaded and used by the system BIOS to access the devices connected to the SATA and/or sSATA controllers.
    - In both cases, Optical Disk Drives will be managed by the System BIOS.

2.6.2.2 Intel® RSTe PCH/SCU OEM Parameters

For those platforms that include the Intel® C600 series chipset, to properly configure and initialize the SCU Controller there are OEM defined parameters that must be established and properly programed into SPI Flash. The information contained this field must be supplied to the SCU Pre-OS images (Intel® RSTe SCU Legacy Option ROM or Intel® RSTe SCU UEFI driver). The information is then used to properly initialize the SCU controller. These values are very dependent on the design of the platform. The mechanism that is used to pass the OEM Parameter information from SPI Flash to the SCU Pre-OS the INT15 interface outlined in the platform's BIOS Writer's Guide (and the release notes included in this package).



NOTE: The SCU controller will not work properly if the OEM Parameters have not been properly programmed into SPI Flash and properly supplied to the SCU Pre-OS.

2.6.2.2.1 Intel® RSTe PCH/SATA Legacy RAID Option ROMs

The Intel® RSTe PCH/SATA product Pre-OS two image option package will contain the Intel® RSTe SATA Legacy Option ROM image that will be loaded with the SATA controller is configured for RAID Mode. This image provides a menu driven interface that allows for the RAID management of RAID data volumes and RAID boot volumes on drives directly attached to the SATA controller. When there are two or more drives attached to the SATA Controller (in RAID mode) the RSTe menu driven GUI Splash screen will appear. To access the menu driven GUI, just enter "CTRL-I" when prompted. Through this menu driven GUI, RAID volumes can be created before installing the operating system. This will allow the operating system to be installed onto the RAID volume created.

2.6.2.2.2 Intel® RSTe PCH/sSATA Legacy RAID Option ROMs

The Intel® RSTe PCH/SATA product Pre-OS two image option package will contain the Intel® RSTe sSATA Legacy Option ROM image that will be loaded with the sSATA controller is configured for RAID Mode. This image provides a menu driven interface that allows for the RAID management of RAID data volumes and RAID boot volumes on drives directly attached to the sSATA controller. When there are two or more drives attached to the sSATA Controller (in RAID mode) the RSTe menu driven GUI Splash screen will appear. To access the menu driven GUI, just enter "CTRL-I" when prompted. Through this menu driven GUI, RAID volumes can be created before installing the operating system. This will allow the operating system to be installed onto the RAID volume created.

The Intel® SATA and/or sSATA Legacy RAID Option ROMs and Intel® RSTe SATA/sSATA UEFI drivers only supports drives that are direct attached to the SATA and/or sSATA controllers.

2.6.2.2.3 Intel® RSTe PCH/SCU Legacy RAID Option ROM

Intel® RSTe PCH/SCU Legacy RAID Option ROM is part of the two image option and is included to provide support for Legacy mode booting. The BIOS configuration utility may provide an option to select the SCU controller as the boot controller. When the system is configured to boot from the SCU controller, the complete Intel® RSTe SCU Legacy RAID Option ROM will be loaded and will provide the interface to the drives attached to the SCU controller. The Intel® RSTe SCU Legacy RAID Option ROM will manage and boot from drives that are directly attached to the SCU controller from within a single level SAS expander. For the single level SAS expander the Intel® RSTe SCU Legacy Option ROM only supports:

- The SAS expander attached to the first 4 ports to the SCU controller.
- Drive that reside within the first 8 ports (port0 thru port7) of the expander. Those will be the only ports scanned for drives and reported to the BIOS. No other ports will be scanned
- In a mixed configuration (some but not all of the first 4 ports [of the SCU controller] contain direct attached drives and the other [of the first 4] ports are connected to an expander), only the drives in the first 8 ports, of the expander, will be exposed to the BIOS as bootable
    - **NOTE: Intel does not recommend this type of configuration**

- If an expander is attached to any of the first 4 ports of the SCU controller, the second 4 ports will not be scanned for drives.
- An expander connected to second 4 ports of the SCU controller will not be scanned

While booting, a BIOS Splash Screen will appear on the display (provided that there are a least two drives attached) that will show what is attached to the SCU controller. There is also an option to stop the booting process and enter into the Intel® RSTe SCU Legacy RAID Option ROM user interface. This is done by pressing the [CTRL]-I key combination. Through this menu driven GUI, RAID volumes can be created before installing the operating system. This is mainly used to create a RAID volume that can be used as the system OS boot device.

The RSTe Legacy SCU RAID Option ROM Splash Screen can be controlled based off of the BCFS settings outlined in BCFS Bit Settings section.

It is required for the SCU RAID Legacy Option ROM to be used in PCI 3.0 compliant BIOS's only.

2.6.2.2.4 Intel® RSTe PCH/SATA RAID UEFI Drivers

The Intel® RSTe PCH/SATA product Pre-OS two image option package will contain the Intel® RSTe SATA UEFI driver image that will be loaded with the SATA controller is configured for RAID Mode. This image provides an HII menu driven user interface that allows for the RAID management of RAID data volumes and RAID boot volumes on drives directly attached to the SATA controller. The UEFI HII menu driven interface is accessed via the Platform system BIOS setup. Through this menu driven user interface, RAID volumes can be created before installing the operating system. This will allow the operating system to be installed onto the RAID volume created. The Intel® RSTe SATA UEFI RAID Driver will support only drives directly attached to a SATA controller.

2.6.2.2.5 Intel® RSTe PCH/sSATA RAID UEFI Drivers

The Intel® RSTe PCH/sSATA product Pre-OS two image option package will contain the Intel® RSTe sSATA UEFI driver image that will be loaded with the sSATA controller is configured for RAID Mode. This image provides an HII menu driven user interface that allows for the RAID management of RAID data volumes and RAID boot volumes on drives directly attached to the sSATA controller. The UEFI HII menu driven interface is accessed via the Platform system BIOS setup. Through this menu driven user interface, RAID volumes can be created before installing the operating system. This will allow the operating system to be installed onto the RAID volume created. The Intel® RSTe sSATA UEFI RAID Driver will support only drives directly attached to the sSATA controller.

2.6.2.2.6 Intel® RSTe PCH/SCU RAID UEFI Driver

The Intel® RSTe PCH/SCU product Pre-OS two image option package will contain the Intel® RSTe SCU UEFI driver image that will be loaded with the SCU controller. This image provides an HII menu driven user interface that allows for the RAID management of RAID data volumes and RAID boot volumes on drives directly attached to the SCU controller. The UEFI HII menu driven interface is accessed via the Platform system BIOS setup. Through this menu driven user interface, RAID volumes can be created before installing the operating system. This will allow the operating system to be installed onto the RAID volume created. The Intel® RSTe SCU UEFI RAID Driver will support drives directly attached to the SCU controller and will provide at least one level of SAS expander support.



2.6.3 Intel® RSTe PCH/SATA Configuration Tools

The Intel® RSTe PCH/SATA product will support multiple ways for OEMs/ODMs and users to manage RAID arrays and volumes. There is a Pre-boot package, factory installation utilities and an optional end user GUI tool.

2.6.3.1 Intel® RSTe PCH/SATA UEFI Human Interface Infrastructure

Intel® RSTe PCH/SATA product will provide support for a UEFI HII RAID management interface for Intel® C600/C610/C220 and C230 series chipset based platforms. This functionality will be included in the RSTe UEFI driver and need to be included in the system BIOS. The system BIOS must adhere to the UEFI 2.3.1 Spec in order to take advantage of this feature and must expose the interface to the user.

2.6.3.2 Intel® RSTe PCH/SATA RAID Configuration Utilities

Intel® RSTe PCH/SATA product will provide support for Pre-OS environment command line RSTe RAID configuration utilities that can be used in manufacturing environment. There is a DOS base command line RSTe RAID configuration utility that can be used in conjunction with the Legacy Option ROMs. There is also a UEFI based command line RSTe RAID configuration utility that operated from the UEFI Shell environment.

The Intel® RSTe product package provides a set of these tools for each of the SATA and sSATA controllers and one for the SCU controller. There is also a set of tools to work with the single Pre-OS images for the SATA/sSATA controllers.

To use the DOS based tools, copy the tool(s) to a DOS bootable media (floppy drive or USB drive). Once booted to that device the tool(s) can be used.

To use the UEFI based tools, copy the tool(s) to a USB drive. Once the system is booted to a UEFI shell, the tool(s) can be used from the USB drive.

2.6.3.3 Intel® RSTe PCH/SATA RAID Compliance Utilities

Intel® RSTe PCH/SATA product will provide support for Pre-OS environment command line RSTe RAID compliance utilities that can be used in manufacturing environment. This tool is used to verify that the Pre-OS images have been properly implemented in the system BIOS. There is a DOS base command line RSTe RAID compliance utility that can be used in conjunction with the Legacy Option ROMs. This is used to help verify that the Legacy Options ROMs have been properly incorporated into the system BIOS

There is also a UEFI based command line RSTe RAID compliance utility that operated from the UEFI Shell environment. This tool helps verify that the RSTe UEFI drivers have been properly incorporated into the system BIOS.

The Intel® RSTe product package provides a set of these tools for each of the SATA and sSATA controllers and one for the SCU controller. There is also a set of tools to work with the single Pre-OS images for the SATA/sSATA controllers.

To use the DOS based tools, copy the tool(s) to a DOS bootable media (floppy drive or USB drive). Once booted to that device the tool(s) can be used.

To use the UEFI based tools, copy the tool(s) to a USB drive. Once the system is booted to a UEFI shell, the tool(s) can be used from the USB drive.

2.6.3.4 Intel® RSTe PCH/SATA Command Line Interface (CLI) Application

Intel ® RSTe PCH/SATA product will provide support for a command line application that can run under a Windows* command prompt and/or a Windows* PE environments. This application can be used to perform basic RAID operations on the platforms that have or will have Intel® RSTe installed. Intel® RSTe CLI provides basic support for creating and managing RAID arrays and volumes without a dependency on the system OS being installed. (I.e. a factory environment that builds both Windows* and Linux* systems)

2.6.3.5 Intel® RSTe PCH/SATA Graphical User Interface (GUI)

The Intel® RSTe PCH/SATA product will provide support for a graphical user interface for management of RAID arrays and volumes. The Intel® RSTe GUI is used to manage RAID arrays and volumes on the devices attached to the SATA, sSATA and/or SCU controllers. It will be able to distinguish between direct attached devices and expander attached storage devices (expanders are only supported on the SCU controller).Intel® RSTe Management Tools.

NOTE: The RSTe GUI does not support Volumes spanned across controllers.

2.6.3.6 Common Information Model (CIM)

Intel® RSTe PCH/SATA product will support an industry standard management API based on CIM model and Storage Management Initiative Specification (SMIS) specification. Samples of the CIM Profiles that will be included in the initial Intel® RSTe release are as follows:

- Host hardware raid controller profile
- Block services profile
- Physical asset profile
- Software inventory profile
- Generic initiator ports profile
- Direct attached target ports profile
- Job control profile
- Indication profile

Intel® RSTe will support an industry standard management API based on CIM model and Storage Management Initiative Specification (SMIS) specification (Linux).

This feature will be supported on platforms that have installed Linux, Windows*7 and Windows* 2008R2 (64 and 32 bit).

2.6.3.7 Common Storage Management Interface (CSMI)

The Intel® RSTe PCH/SATA product will support the Common Storage Management Interface (CSMI) for reporting RAID configurations and SMP, SSP, STP pass through.

---

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



2.6.4 Intel® RSTe System Configurations supported

This section addresses to physical components of the system configuration supported by the Intel® RSTe product package.

2.6.4.1 Intel® C600 series chipset SATA Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SATA ports of the Intel® C600 series chipset.

2.6.4.2 Intel® C600 series chipset SCU Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SCU ports of the Intel® C600 series chipset.

2.6.4.3 Intel® C600 series chipset - SCU OEM Parameters

The Intel® RSTe PCH/SATA product will provide support for the OEM to define some of their own operational parameters (i.e. SAS Addresses) into the PDR region of SPI Flash. This information will be passed to RSTe components from the BIOS who will access the OEM parameters. Please review the platform BIOS Writer's Guide for more details.

2.6.4.4 Intel® C600 series chipset - SAS Expander Support

The Intel® RSTe PCH/SCU product will support 1 expander attached to an SCU controller (provide external HW drive and expander compatibility list). Intel® RSTe will not support the use of port multipliers on either the SATA or SCU controller. This expander will also be reported and exposed to the Device Manager.

2.6.4.5 Maximum Hard Drives supported on Intel® C600 series chipset Based Platforms

The Intel® RSTe PCH/SCU product will provide support for the following quantity of physical disks (direct attached or through SAS expanders):

- SKU –A:
    - Physical Disks (Pass Through and RAID):
        - Single SCU configuration (SATA disks only) – 4
        - SATA - 6
            - 3 Gb/s – 4
            - 6 Gb/s - 2
    - Maximum number of RAID volumes on the direct attached physical disks:
        - SCU – 4
        - SATA - 4
- SKU -B:
    - Physical Disks (Pass Through and RAID):
        - Single SCU configuration (SAS and SATA supported) – 32
        - SATA - 6

- 3 Gb/s – 4
- 6 Gb/s - 2

- Maximum disks managed by Intel® RSTe RAID – 32
- Maximum number of RAID volumes on the direct attached physical disks:
  - SCU – 4
  - SATA - 4

NOTE: With both SKUs –A and –B, all of the SCU traffic is moved up to the processor through the DMI uplink, which is shared by the other components managed by Intel® C600 series chipset.

- SKU –D and -T:
  - Physical Disks:
    - Dual SCU configuration (SAS and SATA supported) – 64
    - SATA
      - 3 Gb/s – 4
      - 6 Gb/s - 2
  - Maximum disks managed by Intel® RSTe RAID – 32
  - Maximum number of RAID volumes on the direct attached physical disks:
    - SCU – 8
    - SATA - 4

There is no OS RAID software limitations imposed.

2.6.4.6 Intel® C600 series chipset - SCU Controller RAID Management Limitations

The Intel® RSTe PCH/SATA product will support the RAID management of up to 32 physical drives attached to the SCU controller. Drives added beyond this limitation (up to a total of 64 drives) will be supported as pass-through drives but will not be validated as part of supported RAID array configurations. The Intel® RSTe GUI will allow up to 8 RAID volumes to be created across the 32 drives. For example, a RAID array that encompasses all 32 drives will result RAID volume limitation of up to 2 volumes (Matrix RAID allows 2 RAID volumes per RAID array).

No OS based software RAID (non-Intel® RSTe 3.0) limitations are imposed.

2.6.4.7 Intel® C220 series chipset SATA Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SATA ports of the Intel® C220 series chipset.

2.6.4.8 Intel® C610 series chipset SATA/sSATA Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SATA and sSATA ports of the Intel® C610 series chipset.

NOTE: Spanning RAID volumes across the SATA and sSATA controllers is not supported.



2.6.4.9 Intel® C230 series chipset SATA Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SATA ports of the Intel® C230 series chipset.

2.6.4.10 Pass-through drives

Intel® RSTe PCH/SATA will support the ability to expose non-RAID configured disks (pass-through) to Host OS. Please refer to section 2.6.4.5 for additional information.

2.6.4.11 Hot Plug

The Intel® RSTe PCH/SATA will support the ability to Hot Plug (remove and replace) disk drives on the controllers of the Intel® C600/C610/C220 and C230 series chipsets, whether or not I/O is being processed. To enable this functionality, the system BIOS must have this feature enabled

The Intel® RSTe NVMe 4.5 will introduce Hot Plug support on Intel® C610 and C230 series chipsets on platforms that include the Intel® SAS/SATA/NVMe Combo-Enclosures (attached to the PCI express slot managed by the CPU). The latest BIOS updates to the Intel® C610 series chipset is required to support this feature. Please contact your Intel FAE to obtain these latest BIOS updates.

2.6.4.12 Drive roaming

Intel® RSTe PCH/SATA will support the ability to move and recognize RAID arrays and volumes between the controllers of the Intel® C600/C610/C220 series chipsets. RAID Boot volumes may have limitations based on the drivers loaded and the controller configurations.

2.6.4.13 Volume Roaming between Linux* and Windows*

Intel® RSTe PCH/SATA will support the ability to move RAID data volumes (configured appropriately) between Linux* and Windows* environments and the RAID data volumes will be recognized and available for use.

2.6.4.14 SGPIO on SATA Controller (RAID Mode Only)

Intel® RSTe PCH/SATA will support enclosure management, compliant to SFF-8485 as well as SFF-8489, to identify drive location or unit failures on the SATA/sSATA controllers.

2.6.4.15 SGPIO on SCU

Intel® RSTe PCH/SCU will support enclosure management, compliant to SFF-8485 as well as SFF-8489, to identify drive location or unit failures on the SCU controller.

---

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



2.6.4.16 NCQ (SATA/sSATA) and CQ (SCU) support

The Intel® RSTe PCH/SATA will support Native Command Queuing (SATA/sSATA) and Command Queuing (SAS).

2.6.4.17 SCSI Enclosure Service (SES) v2

The Intel® RSTe PCH/SCU product will provide support management of enclosures that are compliant with SES (SCSI Enclosure Services) v2 attached to the SCU controller. Intel® RSTe also supports in-band management to SES compliant expanders attached to the SCU.

2.6.5 Software RAID Functional Support

This section will focus on RAID specific features unless the particular requirement specifies differently.

2.6.5.1 Matrix RAID

The Intel® RSTe PCH/SATA will support up to two logical RAID volumes on the same array. A RAID array simply refers to the set of disk drives that can be formed into a RAID volume.

2.6.5.2 RAID 0/1/5/10 Volumes

The I Intel® RSTe PCH/SATA will support base level RAID volumes on both drives connected to the SATA, sSATA and/or SCU controllers. RAID volume spanning across the SATA, sSATA and SCU controllers is not supported.

2.6.5.3 Software RAID 5 on SCU

The Intel® RSTe PCH/SCU product will support software RAID 5 on Intel® C600 series chipset–T SKUs as an optional feature. This feature is enabled by means of the “upgrade ROM” capabilities or by purchasing the –T SKU of Intel® C600 series chipset.

2.6.5.4 Simultaneous RAID Arrays

The Intel® RSTe PCH/SATA will provide support for RAID volume management on disks attached to the SCU controller separate from disks attached to the SATA controllers. However, Intel® RSTe will provide support for simultaneous RAID management on both.

2.6.5.5 Disk Coercion

The Intel® RSTe PCH/SATA will 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.



2.6.5.6 Hot Spare Disk

The Intel® RSTe PCH/SATA will support the ability to set a drive as a hot spare that would automatically be used to rebuild a failed or degraded RAID volume without any user interaction. This applies to both the SATA, sSATA and/or SCU controllers. This support is provided in the Intel® RSTe GUI as well as in the Pre-OS images.

2.6.5.7 Auto Rebuild on Hot Insert

The Intel® RSTe PCH/SATA product will support the ability to initiate an automatic RAID rebuild when a physical disk of the appropriate size is hot inserted into the same directly attached port that the failed drive was removed from. When configured appropriately, if a RAID volume issue occurs (failure, degradation, or SMART event) and the questionable drive is hot removed, if a drive of the appropriate size (new or and from an off-line RAID volume) is hot inserted into that same port, the volume will be rebuilt on the inserted drive.

2.6.5.8 Manually Invoked Rebuild

The Intel® RSTe PCH/SATA will provide a manual method to initiate a RAID volume rebuild if a hot spare has not been configured or is not available.

2.6.5.9 RAID SMART Support

The Intel® RSTe PCH/SATA will provide support for SMART Alerts for SAS and SATA disks. A SMART drive event response alert on failure will initiate rebuild to hot spare disk.

2.6.5.10 RAID-Ready Mode

A RAID-Ready system refers to a system that has been configured to support Intel® RSTe PCH/SATA product. The system BIOS has the appropriate pre-boot drivers and has been configured for RAID mode. RAID mode can be either:

- The system is configured to boot off the SATA/sSATA controllers and it is in RAID mode
- The system is configured to boot off the SCU controller

Intel® RSTe will support an Intel® C600/C610/C220 and C230 series chipset based platform configured in RAID-Ready mode.

2.6.5.11 RAID Volume Creation with Data Preservation

The Intel® RSTe PCH/SATA will support the ability to preserve the data from one of the disks used for the volume creation. A non-RAID disk can be migrated to a RAID volume while retaining the existing data on that disk.

Note: When creating a system boot volume (the Intel® RSTe family of products only), the maximum stripe size supported is 128K.

In a RAID-Ready configuration, the user can take their single system drive and turn it into a supported RAID volume by using the Intel® RSTe GUI application. This process does not require the reinstallation of the operating system. All applications and data will remain intact.

The following are examples of RAID level creations that will be supported by Intel® RSTe (depending on the chipset being used):

- Individual pass-through to 2 -16 drives for RAID 0 (Intel® RSTe PCH/SATA product)
- Individual pass-through to 2 - 8 drives for RAID 0 (Intel® RSTe NVMe product)
- Individual pass-through to 2 drive RAID 1
- Individual pass-through to 4 drive RAID 10
- Individual pass-through to 3 to 6 drive RAID 5

2.6.5.12 Instant Initialization

The Intel® RSTe PCH/SATA will support a newly created volume to be used immediately (no reboot required), protecting newly written data and creating parity data concurrently.

For a RAID 5 configuration that consists of 3 or 4 drives, the RAID volume will be shown as normal as soon as the volume is created. Parity will be computed and written with every RAID 5 write activity.

For a RAID 5 configuration that consists of 5 or more drives, the parity initialization will begin as soon as the volume is created. This is done to improve the operational performance of RAID 5 volumes.

2.6.5.13 RAID Level Migrations

The RAID level migration feature in the Intel® RSTe PCH/SATA will enable the ability to convert the contents of a drive (attached to the SATA, sSATA and/.or SCU controller) into a RAID volume (RAID 0, RAID 1, RAID 5, or RAID 10). The RAID level migration feature also supports the ability to migrate from a one RAID volume to another.

The size of the hard drives determines how much time is required to complete the migration but the system will remain fully functional during the migration process. The only limitation is that some disk-intensive tasks may have slower performance during a RAID migration.

NOTE: Single volume per array only. This is dependent on required capacity and implicit array expansion.

NOTE: When using a GPT partition, make sure there is at least 5 Megabytes of disk space available **for RAID Metadata when the OS is being installed.**

The following are some examples of RAID level migrations supported by Intel® RSTe (depending on the chipset being used):

- N-drive RAID 0 to N+1 – 32 drive RAID 5 (where N can be 2 to 31)
- 2-drive RAID 1 to 3 - 32 drive RAID 5 (8 drives max for Intel® RSTe NVMe product)
- 4-drive RAID 10 to 4 - 32 drive RAID 5 (8 drives max for Intel® RSTe NVMe product)

NOTE: Additional information on RAID level migration can be found in the RSTe GUI help

2.6.5.14 RAID Reconfiguration (Stripe size)

The Intel® RSTe PCH/SATA will provide the ability to change stripe size on existing volumes (migration required). Intel® RSTe will support a stripe size migration in conjunction with a RAID level migration.



Note: Migration supports stripe sizes for the respective RAID levels supported. Stripe Size Support for (values are in Kilobytes):

- RAID 0 volumes – 4, 8, 16, 32, 64, 128, 256, 512, 1024
- RAID 10 volumes - 4, 8, 16, 32, 64, 128, 256, 512, 1024
- RAID 5 volumes - 4, 8, 16, 32, 64, 128, 256, 512, 1024

2.6.5.15 Expanded Stripe Size

The Intel® RSTe PCH/SATA will support the ability to expand the RAID volume stripe size for the following RAID volumes (values are in Kilobytes):

- RAID 0 volumes – 256, 512, 1024
- RAID 10 volumes – 256, 512, 1024
- RAID 5 volumes – 256, 512, 1024

2.6.5.16 Online Array / Volume Capacity Expansion

The Intel® RSTe PCH/SATA will provide the ability to add new drives to an existing array and expand existing volumes accordingly. This is supported only under RAID 0 and RAID 5.

2.6.5.17 Read Patrol

The Intel® RSTe PCH/SATA will provide support for Read Patrol, which checks the RAID volumes for errors that could result in a failure. The checks are done periodically in background and will verify all sectors of all RAID volumes that are connected to the SATA, sSATA and/or SCU controllers and NVMe SSDs. If an issue is discovered an attempt at corrective action is taken. Read Patrol can be enabled or disabled manually.

The background process begins when there is no I/O to the RAID volume, though it can continue to run while I/O's are being processed.

2.6.5.18 Verify and Repair

The Intel® RSTe PCH/SATA will provide support for Verify and Repair.

The RAID volume data verification process identifies any inconsistencies or bad data on a RAID 0, RAID 1, RAID 5, or RAID 10 volume.

The RAID volume data verification and repair process identifies and repairs any inconsistencies or bad data on a RAID 1, RAID 5, or RAID 10 volume.

The following describes what occurs for each RAID level:

RAID Level	Verify	Verify & Repair
RAID 0	Bad blocks are identified.	N/A

RAID 1	Bad blocks are identified	Bad blocks are reassigned.
	Data on the mirror drive is compared to data on the source drive.	If the data on the mirror drive does not match the data on the source drive, the data on the mirror is overwritten with the data on the source.
RAID 5	Bad blocks are identified.	Bad blocks are reassigned.
	Parity is recalculated and compared to the stored parity for that stripe.	If the newly calculated parity does not match the stored parity, the stored parity is overwritten with the newly calculated parity.
RAID 10	Bad blocks are identified.	Bad blocks are reassigned.
	Data on the mirror is compared to data on the source.	If the data on the mirror does not match the data on the source, the data on the mirror is overwritten with the data on the source.

2.6.5.19 Check Pointing

The Intel® RSTe PCH/SATA will provide the ability to perform Check Pointing to be able to track forward progress on read patrol, array rebuilds and volume migration if interrupts occur. After resuming, the operation will restart from the last valid stage reached.

2.6.5.20 Bad Block Management

The Intel® RSTe PCH/SATA cts will provide support for Bad Block Management.

In the course of rebuilding a degraded RAID volume, where one of the member disks has failed or been removed, and is being replaced by a 'spare' drive, the redundant contents of the other drive(s) are read and then used to reconstruct data to be written to the spare drive. In case a read failure occurs sometime during this rebuild process, the data to be written to the spare will not be available and therefore lost. In this scenario, rather than mark the entire RAID volume as failed, we can mark only those sectors on the spare that are known to have indeterminate data, in a log of such bad sectors. This bad block management log can be used to reflect error status whenever any attempts are made to access those sectors of the spare.

2.6.5.21 Dirty Stripe Journaling

The Intel® RSTe PCH/SATA will provide support for Dirty Stripe Journaling (DSJ). DSJ is used to help speed up RAID 5 write power loss recovery by storing the write stripes that were in progress



at the time of the failure. The DSJ allows rapid recovery without having to rebuild the entire volume. The DSJ is only utilized when disk write cache is DISABLED

2.6.5.22 Partial Parity Logging (PPL)

The Intel® RSTe PCH/SATA will provide support for Partial Parity Logging (PPL). PPL is used to record the results of XORing old data with old parity. PPL is currently saved as part of the RAID member information and is only utilized when writing RAID 5 parity. It helps protect against data loss when a power failure or a system crash occurs by allowing data to be rebuilt by utilizing the PPL information.

2.6.5.23 OS Installation

The Intel® RSTe PCH/SATA product will provide the OS appropriate RSTe driver files required for installation during the OS setup onto a drive or RAID volume attached to either the SATA, sSATA or SCU controllers.

2.6.5.24 Selectable Boot Volume

The Intel® RSTe PCH/SATA product will support the ability to select any volume as the OS boot volume. The OS installer will be able to install the operating system onto RAID volume. There will be no need for RAID management (e.g. volume creation/deletion) support from within OS installer.

2.6.5.25 Auto Rebuild

The Intel® RSTe PCH/SATA will provide support for the ability to automatically rebuild a failed or degraded RAID volume. This feature will begin when a member disk of the array has failed and a suitable replacement disk with sufficient capacity is available. As soon as the failure occurs the rebuild process will begin automatically, using the marked Hot Spare disk, without user intervention.

If a marked Hot Spare disk is not present, the automatic rebuild process will begin under the following conditions:

- Another free disk is plugged into the same directly attached physical location as the failed drive
- The newly inserted disk size is at least as large as the amount of space used per disk in the current array
- The newly inserted disk must be the same type (SAS/SATA) as the disk being replaced or the rebuild will not start.
- If the newly inserted disk contains Intel® RSTe (or Intel® RST) metadata with current status of member being offline or contains no Intel® RSTe (or Intel® RST) metadata.
- The newly inserted disk has not reported a SMART event.

The following table summarizes the functionality:

Controller	Auto Rebuild Support	Action
SATA & SCU	Previously marked Hot Spare available.	Rebuild starts when spare found. This takes precedence over auto-spare disk.
SATA	No Hot Spare previously marked	No auto rebuild: Manual steps required to rebuild array using new disk
SCU	Auto rebuild conditions described above are met.	Auto rebuild starts without any user intervention
SCU	One or more of the above conditions was not met.	No auto rebuild: Manual steps required to rebuild array using new disk
sSATA	Auto rebuild conditions described above are met.	Auto rebuild starts without any user intervention
sSATA	One or more of the above conditions was not met.	No auto rebuild: Manual steps required to rebuild array using new disk

Automatic rebuild support will default to OFF for Intel® RSTe and can be enabled through the Intel® RSTe GUI.

2.6.5.26 Error Threshold Monitoring/Handling

The Intel® RSTe PCH/SATA will support the ability to initiate an automatic RAID rebuild to a marked hot spare drive in the event of a drive SMART event alert that indicates a failure. (Windows*Only)

2.6.5.27 Unified Extensible Firmware Interface (UEFI)

The Intel® RSTe PCH/SATA product will support UEFI for the SCU, SATA and sSATA controllers using common metadata.

2.6.5.28 Extended SCSI Passthru

The Intel® RSTe PCH/SATA product will support the following Extended SCSI Passthru commands as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D:

- EFI_EXT_SCSI_PASS_THRU_PROTOCOL protocol for physical disk devices
- Support for EFI_DEVICE_PATH_PROTOCOL for physical disk devices
- Support for EFI_SCSI_IO_PROTOCOL for physical disk devices

---

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



2.6.5.29 Disk Write Cache

The Intel® RSTe PCH/SATA will support the ability to enable/disable Disk Write Cache through the Intel® RSTe GUI. Disk Data Cache will be disabled by default.

2.6.5.30 RAID Volume Read Cache

The Intel® RSTe PCH/SATA will support the ability to enable/disable RAID Volume Read Cache through the Intel® RSTe GUI. RAID Volume Read Cache will be enabled by default.

2.6.5.31 Write Back Cache

The Intel® RSTe PCH/SATA will support the ability to enable/disable Write Back Cache through the Intel® RSTe GUI. Write Back Cache will be disabled by default.

2.6.5.32 Volume Cache Increase

The Intel® RSTe PCH/SATA will increase the volume cache size to 16MB for SCU volumes and 16MB for SATA and sSATA volumes.

2.6.5.33 RAID Data Volume Size

Intel® RSTe will provide support for RAID volumes that are larger than 2 Terabytes.

2.6.5.34 RAID Boot Volume Size

The Intel® RSTe PCH/SATA product will provide support for RAID Boot volumes that are larger than 2 Terabytes.

2.6.5.35 Disk Monitor Service

The Intel® RSTe PCH/SATA will support the ability to provide a disk monitoring service. The service will be active by default and executed as a system service. The service will monitor the system for SMART and RAID volume state changes events. The changes will be logged in the system log.

2.6.5.36 Failed Drive Reinsertion

The Intel® RSTe PCH/SATA will support the ability to recognize a failed drive re-inserted into the array. If able, Intel® RSTe will attempt to rebuild the volume to that drive. If not able, Intel® RSTe will mark the drive accordingly in the GUI.

2.6.5.37 Drives Supported

The Intel® RSTe PCH/SATA will provide support for current production SATA drives from “any” manufacturer on the SATA/sSATA controllers; current production SAS and SATA drives supported on the SCU controller and Intel NVMe SSD (on supported platforms).



2.6.5.38 Safe Mode Support

The Intel® RSTe PCH/SATA will provide support for booting into Safe Mode for all supported OSs.

2.6.5.39 Non-Intel Controller Support

The Intel® RSTe PCH/SATA will not hinder, break or prevent operation of non-Intel® Controllers (SATA, PATA, SATA or RAID) or the existence of non-Intel NVMe SSDs

2.6.5.40 Device Configuration

The Intel® RSTe PCH/SATA will support the ability, at initialization, to read the system registry to get configuration setting in order to set the appropriate operational parameters.

2.6.5.41 Power Management

The Intel® RSTe PCH/SATA will support all the necessary power management functions required by the OSs.

2.6.5.42 Staggered Spin-up

The Intel® RSTe PCH/SATA product will provide support for staggered spin-up on the SCU controller for those hard drives that support this feature.

2.6.5.43 Exporting SATA Drives on SATA/sSATA Controller

The Intel® RSTe PCH/SATA RAID Legacy Option ROMs will export those drives directly attached on a port order basis. This will be for both the SATA and sSATA controllers.

2.6.5.44 ATAPI on SATA/sSATA Controller

The Intel® RSTe PCH/SATA product will provide support for ATAPI devices connected to the SATA/sSATA controller. Intel® RSTe Legacy RAID Option ROM will only support HDD devices (not ATAPI).

Pre-OS ATAPI support is handled by the platform system BIOS.

2.6.5.45 Solid State Drives (SSD)

The Intel® RSTe PCH/SATA product will support SSDs as if they are Hard Disk Drives.

2.6.5.46 Hybrid Drives

The Intel® RSTe PCH/SATA product will support Hybrid Drives as if they are Hard Disk Drives.

2.6.5.47 TRIM Command

Note: This feature is not an end-user visible feature. There is no Intel® RSTe application or user interface control to configure the feature. Registry settings are provided for OEM use.



Support for the TRIM command allows the OS to pass information to the Solid State Disk (SSD) that identifies sectors that can be deleted. The SSD will then go through and clear out that information in the background thereby minimizing the chances of an “Overwriting” process happening at crucial times. The SSD is also free to do some additional optimizations with those sectors. E.g. an SSD can pre-erase any sector that has been TRIM'ed. The TRIM command improves the long term Write performance and the life-span of SSDs.

RAID 0/1/10

2.6.5.47.1 SATA/sSATA Controllers

The Intel® RSTe PCH/SATA product will support TRIM on both SATA controllers in a non-RAID configurations as well as on RAID0, RAID1 and RAID10 volumes.

2.6.5.47.2 SCU Controller

The Intel® RSTe PCH/SCU product will support TRIM on the SCU controller in a non-RAID configurations as well as on RAID0, RAID1 and RAID10 volumes.

2.6.5.48 Email Alerting and Notification[1]

The Intel® RSTe PCH/SATA will support email notification of certain storage events.

2.6.6 Utilities

2.6.6.1 Install/Uninstall Utility

The Intel® RSTe PCH/SATA will be available through the use of an install package. The install package will automatically install the proper RSTe driver and GUI that corresponds to the OS being installed on. There will also be a mechanism available to uninstall the driver and GUI.

NOTE: Great care must be taken when trying to perform the uninstall process. Removal of the driver could result in a system crash that could require a complete reinstallation of the operating system.

[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.

2.7 Intel® RSTe NVMe Key Product Features

Name	Key Features	
RAID	• Support for Matrix RAID 0/1/5/10 • Pass-through drives • Hot Spare Disk • Auto Rebuild on Hot Insert[1] • Rebuild & Migration Check Pointing • NCQ (SATA) and CQ (SAS) support • SMART Support • Bad Block Management • Controller configuration rules • Drive roaming • RAID Volume roaming between Linux* and Windows* • On Line Capacity Expansion • Large Stripe Size Support • RAID-Ready	• Disk Coercion • Manual & Auto Rebuild • Instant Initialization • Read Patrol • Volume creation/verify • Email Alerting • RAID Level Migration (RAID 0, 1, or 10 to RAID 5) • Dirty Stripe Journaling • Partial Parity Logging (PPL) • Verify and Repair • NVMe Pass-Through • SCSI Pass-Through
Utilities	• Install/Uninstall Utility • Configuration and Management Utilities	
Intel® RSTe NVMe	Intel® RSTe NVMe provides software for high-performance Intel® NVMe SSDs • Matrix RAID for two RAID volumes on single array • Improved user interface for enhanced usability	
1	Feature supported with Intel® RSTe NVMe 4.5	

2.7.1 Operating System Support

Intel® RSTe NVMe will support both 32 and 64 bit versions of the following Microsoft OSs:

- Windows* 7
- Windows* Server 2008 SP2
- Windows* Server 2008 R2 (64 bit only)
- Windows* PE 3.0
- Windows* Server 2012
- Windows* 8.1
- Windows* Server 2012 R2
- Windows* 10
- Windows* Server 2016

---

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

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



2.7.2 Intel® RSTe NVMe Configuration Tools

The Intel® RSTe NVMe product will support two OS level configuration.

2.7.2.1 Intel® RSTe NVMe Command Line Interface (CLI) Application

Intel ® RSTe NVMe product will provide support for a command line application that can run under a Windows* command prompt and/or a Windows* PE environments. This application can be used to perform basic RAID operations on the platforms that have or will have Intel® RSTe installed. Intel® RSTe CLI provides basic support for creating and managing RAID arrays and volumes without a dependency on the system OS being installed. (I.e. a factory environment that builds both Windows* and Linux* systems)

2.7.2.2 Intel® RSTe NVMe Graphical User Interface (GUI)

The Intel® RSTe NVMe product will provide support for a graphical user interface for management of RAID arrays and volumes on Intel® NVMe SSDs directly connected via PCI express slots controlled/managed by the CPU.

This feature will be supported on platforms that have installed Linux, Windows*7 and Windows* 2008R2 (64 and 32 bit).

2.7.3 Intel® RSTe System Configurations supported

This section addresses to physical components of the system configuration supported by the Intel® RSTe product package.

2.7.3.1 Intel® C610 series chipset

The Intel® RSTe NVMe product will provide support for managing RAID volumes on Intel NVMe SSDs attached to PCI express slots managed by the PCU of the Intel® C610 series chipset based platform.

2.7.3.2 Intel® C230 series chipset SATA Controller Support

The Intel® RSTe PCH/SATA product will provide support for managing RAID volumes on drives attached to the SATA ports of the Intel® C230 series chipset.

2.7.3.3 Intel® NVMe SSDs

The Intel® RSTe NVMe product will provide support Intel® NVMe SSDs directly connected to the CPU managed PCI express slots.

2.7.3.4 Pass-through drives

Intel® RSTe NVMe will support the ability to expose non-RAID configured Intel NVMe SSD (pass-through) to Host OS.

2.7.3.5 Drive roaming

Intel® RSTe NVMe will support the ability to move and recognize RAID arrays and volumes between the controllers of the Intel® C600/C610/C220 series chipsets. RAID Boot volumes may have limitations based on the drivers loaded and the controller configurations.

2.7.3.6 Volume Roaming between Linux* and Windows*

Intel® RSTe NVMe will support the ability to move RAID data volumes (configured appropriately) between Linux* and Windows* environments and the RAID data volumes will be recognized and available for use.

2.7.3.7 NCQ

The Intel® RSTe NVMe will support Native Command).

2.7.3.8 NVMe Status LEDs

The Intel® RSTe NVMe product will provide support NVMe Status LED management. Support for this requires the OEM/ODMs to development customer DLL to support the enclosure selected for the platform being developed. Please reference OEM Technical Guild for additional information.

2.7.4 Software RAID Functional Support

This section will focus on RAID specific features unless the particular requirement specifies differently.

2.7.4.1 Matrix RAID

The Intel® RSTe NVMe will support up to two logical RAID volumes on the same array. A RAID array simply refers to the set of disk drives that can be formed into a RAID volume.

2.7.4.2 RAID 0/1/5/10 Volumes

The Intel® RSTe NVMe will support base level RAID volumes on Intel® NVMe SSDs attached to PCI express slots managed by the CPU.

---

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



2.7.4.3 Simultaneous RAID Arrays

The Intel® RSTe NVMe will provide support for RAID volume management on Intel® NVMe SSDs separate from disks attached to the SATA/sSATA controllers. However, Intel® RSTe will provide support for simultaneous RAID management on both areas.

2.7.4.4 Disk Coercion

The Intel® RSTe NVMe will 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.

2.7.4.5 Hot Spare Disk

The Intel® RSTe NVMe will support the ability to set a drive as a hot spare that would automatically be used to rebuild a failed or degraded RAID volume without any user interaction. This applies to both the SATA, sSATA and/or Intel NVMe SSDs. This support is provided in the Intel® RSTe GUI as well as in the Pre-OS images.

2.7.4.6 Auto Rebuild on Hot Insert

The Intel® RSTe NVMe will support the ability to initiate an automatic RAID rebuild when a physical disk of the appropriate size is hot inserted into the same directly attached port that the failed drive was removed from. When configured appropriately, if a RAID volume issue occurs (failure, degradation, or SMART event) and the questionable drive is hot removed, if a drive of the appropriate size (new or and from an off-line RAID volume) is hot inserted into that same port, the volume will be rebuilt on the inserted drive.

2.7.4.7 Manually Invoked Rebuild

The Intel® RSTe NVMe will provide a manual method to initiate a RAID volume rebuild if a hot spare has not been configured or is not available.

2.7.4.8 RAID SMART Support

The Intel® RSTe NVMe will provide support for SMART Alerts for SAS and SATA disks. A SMART drive event response alert on failure will initiate rebuild to hot spare disk.

2.7.4.9 RAID-Ready Mode

A RAID-Ready NVMe to a system that has been configured to support Intel® RSTe NVMe.

Intel® RSTe will support an Intel® C610 and C230 series chipset based platform.

2.7.4.10 RAID Volume Creation with Data Preservation

The Intel® RSTe NVMe will support the ability to preserve the data from one of the disks used for the volume creation. A non-RAID disk can be migrated to a RAID volume while retaining the existing data on that disk.

Note: When creating a system boot volume (the Intel® RSTe family of products only), the maximum stripe size supported is 128K.

In a RAID-Ready configuration, the user can take their single system drive and turn it into a supported RAID volume by using the Intel® RSTe GUI application. This process does not require the reinstallation of the operating system. All applications and data will remain intact.

The following are examples of RAID level creations that will be supported by Intel® RSTe (depending on the chipset being used):

- Individual pass-through to 2 - 8 drives for RAID 0 (Intel® RSTe NVMe product)
- Individual pass-through to 2 drive RAID 1
- Individual pass-through to 4 drive RAID 10
- Individual pass-through to 3 to 6 drive RAID 5

2.7.4.11 Instant Initialization

The Intel® RSTe NVMe will support a newly created volume to be used immediately (no reboot required), protecting newly written data and creating parity data concurrently.

For a RAID 5 configuration that consists of 3 or 4 drives, the RAID volume will be shown as normal as soon as the volume is created. Parity will be computed and written with every RAID 5 write activity.

For a RAID 5 configuration that consists of 5 or more drives, the parity initialization will begin as soon as the volume is created. This is done to improve the operational performance of RAID 5 volumes.

2.7.4.12 RAID Level Migrations

The RAID level migration feature in the Intel® RSTe NVMe will enable the ability to convert the contents of a drive (Intel NVMe SSDs) into a RAID volume (RAID 0, RAID 1, RAID 5, or RAID 10). The RAID level migration feature also supports the ability to migrate from a one RAID volume to another.

The size of the hard drives determines how much time is required to complete the migration but the system will remain fully functional during the migration process. The only limitation is that some disk-intensive tasks may have slower performance during a RAID migration.

NOTE: Single volume per array only. This is dependent on required capacity and implicit array expansion.

NOTE: When using a GPT partition, make sure there is at least 5 Megabytes of disk space available **for RAID Metadata when the OS is being installed.**

The following are some examples of RAID level migrations supported by Intel® RSTe (depending on the chipset being used):

- 2-drive RAID 1 to 3 - 32 drive RAID 5 (8 drives max for Intel® RSTe NVMe product)
- 4-drive RAID 10 to 4 - 32 drive RAID 5 (8 drives max for Intel® RSTe NVMe product)

NOTE: Additional information on RAID level migration can be found in the RSTe GUI help



2.7.4.13 RAID Reconfiguration (Stripe size)

The Intel® RSTe NVMe will provide the ability to change stripe size on existing volumes (migration required). Intel® RSTe will support a stripe size migration in conjunction with a RAID level migration.

Note: Migration supports stripe sizes for the respective RAID levels supported. Stripe Size Support for (values are in Kilobytes):

- RAID 0 volumes – 4, 8, 16, 32, 64, 128, 256, 512, 1024
- RAID 10 volumes - 4, 8, 16, 32, 64, 128, 256, 512, 1024
- RAID 5 volumes - 4, 8, 16, 32, 64, 128, 256, 512, 1024

2.7.4.14 Expanded Stripe Size

The Intel® RSTe NVMe will support the ability to expand the RAID volume stripe size for the following RAID volumes (values are in Kilobytes):

- RAID 0 volumes – 256, 512, 1024
- RAID 10 volumes – 256, 512, 1024
- RAID 5 volumes – 256, 512, 1024

2.7.4.15 Online Array / Volume Capacity Expansion

The Intel® RSTe NVMe will provide the ability to add new drives to an existing array and expand existing volumes accordingly. This is supported only under RAID 0 and RAID 5.

2.7.4.16 Read Patrol

The Intel® RSTe NVMe will provide support for Read Patrol, which checks the RAID volumes for errors that could result in a failure. The checks are done periodically in background and will verify all sectors of all RAID volumes that are connected to the Intel NVMe SSDs. If an issue is discovered an attempt at corrective action is taken. Read Patrol can be enabled or disabled manually.

The background process begins when there is no I/O to the RAID volume, though it can continue to run while I/O's are being processed.

2.7.4.17 Verify and Repair

The Intel® RSTe NVMe will provide support for Verify and Repair.

The RAID volume data verification process identifies any inconsistencies or bad data on a RAID 0, RAID 1, RAID 5, or RAID 10 volume.

The RAID volume data verification and repair process identifies and repairs any inconsistencies or bad data on a RAID 1, RAID 5, or RAID 10 volume.

The following describes what occurs for each RAID level:

RAID Level	Verify	Verify & Repair
RAID 0	Bad blocks are identified.	N/A
RAID 1	Bad blocks are identified	Bad blocks are reassigned.
	Data on the mirror drive is compared to data on the source drive.	If the data on the mirror drive does not match the data on the source drive, the data on the mirror is overwritten with the data on the source.
RAID 5	Bad blocks are identified.	Bad blocks are reassigned.
	Parity is recalculated and compared to the stored parity for that stripe.	If the newly calculated parity does not match the stored parity, the stored parity is overwritten with the newly calculated parity.
RAID 10	Bad blocks are identified.	Bad blocks are reassigned.
	Data on the mirror is compared to data on the source.	If the data on the mirror does not match the data on the source, the data on the mirror is overwritten with the data on the source.

2.7.4.18 Check Pointing

The Intel® RSTe NVMe will provide the ability to perform Check Pointing to be able to track forward progress on read patrol, array rebuilds and volume migration if interrupts occur. After resuming, the operation will restart from the last valid stage reached.

2.7.4.19 Bad Block Management

The Intel® RSTe NVMe will provide support for Bad Block Management.

In the course of rebuilding a degraded RAID volume, where one of the member disks has failed or been removed, and is being replaced by a 'spare' drive, the redundant contents of the other drive(s) are read and then used to reconstruct data to be written to the spare drive. In case a read failure occurs sometime during this rebuild process, the data to be written to the spare will not be available and therefore lost. In this scenario, rather than mark the entire RAID volume as failed, we can mark only those sectors on the spare that are known to have indeterminate data, in a log of such bad sectors. This bad block management log can be used to reflect error status whenever any attempts are made to access those sectors of the spare.



2.7.4.20 Dirty Stripe Journaling

The Intel® RSTe NVMe will provide support for Dirty Stripe Journaling (DSJ). DSJ is used to help speed up RAID 5 write power loss recovery by storing the write stripes that were in progress at the time of the failure. The DSJ allows rapid recovery without having to rebuild the entire volume. The DSJ is only utilized when disk write cache is DISABLED

2.7.4.21 Partial Parity Logging (PPL)

The Intel® RSTe NVMe will provide support for Partial Parity Logging (PPL). PPL is used to record the results of XORing old data with old parity. PPL is currently saved as part of the RAID member information and is only utilized when writing RAID 5 parity. It helps protect against data loss when a power failure or a system crash occurs by allowing data to be rebuilt by utilizing the PPL information.

2.7.4.22 Auto Rebuild

The Intel® RSTe NVMe will provide support for the ability to automatically rebuild a failed or degraded RAID volume. This feature will begin when a member disk of the array has failed and a suitable replacement disk with sufficient capacity is available. As soon as the failure occurs the rebuild process will begin automatically, using the marked Hot Spare disk, without user intervention.

If a marked Hot Spare disk is not present, the automatic rebuild process will begin under the following conditions:

- Another free disk is plugged into the same directly attached physical location as the failed drive
- The newly inserted disk size is at least as large as the amount of space used per disk in the current array
- The newly inserted disk must be the same type (Intel® NVMe SSD) as the disk being replaced or the rebuild will not start.
- If the newly inserted disk contains Intel® RSTe (or Intel® RST) metadata with current status of member being offline or contains no Intel® RSTe (or Intel® RST) metadata.
- The newly inserted disk has not reported a SMART event.

The following table summarizes the functionality:

Controller	**Auto Rebuild Support**	**Action**
NVMe SSDs Connected to the CPU	Auto rebuild conditions described above are met.	Auto rebuild starts without any user intervention
NVMe SSDs Connected to the CPU	One or more of the above conditions was not met.	No auto rebuild: Manual steps required to rebuild array using new disk

Automatic rebuild support will default to OFF for Intel® RSTe and can be enabled through the Intel® RSTe GUI.

2.7.4.23 Error Threshold Monitoring/Handling

The Intel® RSTe NVMe will support the ability to initiate an automatic RAID rebuild to a marked hot spare drive in the event of a drive SMART event alert that indicates a failure. (Windows*Only)

2.7.4.24 Extended SCSI Passthru

The Intel® RSTe NVMe will support the following Extended SCSI Passthru commands as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D:

- EFI_EXT_SCSI_PASS_THRU_PROTOCOL protocol for physical disk devices
- Support for EFI_DEVICE_PATH_PROTOCOL for physical disk devices
- Support for EFI_SCSI_IO_PROTOCOL for physical disk devices

2.7.4.25 Extended NVMe Passthru

The Intel® RSTe NVMe product will support the following Extended NVMe Passthru commands as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D:

- EFI_EXT_SCSI_PASS_THRU_PROTOCOL protocol for physical disk devices
- Support for EFI_DEVICE_PATH_PROTOCOL for physical disk devices
- Support for EFI_SCSI_IO_PROTOCOL for physical disk devices

2.7.4.25.1 Windows driver NVMe_PASSTHRU_IOCTL_PROTOCOL

NVMe pass-through IOCTLs are supported for pass-through drives. RSTe 4.5 supports the ability to send NVMe pass-through IOCTLs directly to RAID member drives. The following is supported:

1) NVME Pass Through Admin commands
    a. Identify
    b. Read Log
2) NVME Pass Through Namespace commands
    a. Read (data commands are not supported as pass-through commands)
    b. Flush cache

2.7.4.25.2 EFI_NVM_EXPRESS_PASS_THRU_PROTOCOL

This is currently supported in NVMe UEFI. It can also be used for RAID member drives with RSTe 4.5.

2.7.4.25.3 NVMe Deallocate (Similar to Trim)

The Intel® RSTe NVMe will support for pass-through and for RAID 0/1/10.

---

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



2.7.4.26 Disk Write Cache

The Intel® RSTe NVMe will support the ability to enable/disable Disk Write Cache through the Intel® RSTe GUI. Disk Data Cache will be disabled by default.

2.7.4.27 RAID Volume Read Cache

The Intel® RSTe NVMe will support the ability to enable/disable RAID Volume Read Cache through the Intel® RSTe GUI. RAID Volume Read Cache will be enabled by default.

2.7.4.28 Write Back Cache

The Intel® RSTe NVMe will support the ability to enable/disable Write Back Cache through the Intel® RSTe GUI. Write Back Cache will be disabled by default.

2.7.4.29 RAID Data Volume Size

Intel® RSTe NVMe will provide support for RAID volumes that are larger than 2 Terabytes.

2.7.4.30 RAID Boot Volume Size

The Intel® RSTe NVMe product will provide support for RAID Boot volumes that are larger than 2 Terabytes.

2.7.4.31 Disk Monitor Service

The Intel® RSTe NVMe will support the ability to provide a disk monitoring service. The service will be active by default and executed as a system service. The service will monitor the system for SMART and RAID volume state changes events. The changes will be logged in the system log.

2.7.4.32 Failed Drive Reinsertion

The Intel® RSTe NVMe will support the ability to recognize a failed drive re-inserted into the array. If able, Intel® RSTe will attempt to rebuild the volume to that drive. If not able, Intel® RSTe will mark the drive accordingly in the GUI.

2.7.4.33 Drives Supported

The Intel® RSTe NVMe will provide support for current production Intel® NVMe.

2.7.4.34 Device Configuration

The Intel® RSTe NVMe will support the ability, at initialization, to read the system registry to get configuration setting in order to set the appropriate operational parameters.

2.7.4.35 Power Management

The Intel® RSTe NVMe will support all the necessary power management functions required by the OSs.

2.7.4.36 TRIM Command

Note: This feature is not an end-user visible feature. There is no Intel® RSTe application or user interface control to configure the feature. Registry settings are provided for OEM use.

Support for the TRIM command allows the OS to pass information to the Solid State Disk (SSD) that identifies sectors that can be deleted. The SSD will then go through and clear out that information in the background thereby minimizing the chances of an "Overwriting" process happening at crucial times. The SSD is also free to do some additional optimizations with those sectors. E.g. an SSD can pre-erase any sector that has been TRIM'ed. The TRIM command improves the long term Write performance and the life-span of SSDs.

RAID 0/1/5/10

2.7.4.37 Email Alerting and Notification[1]

The Intel® RSTe NVMe will support email notification of certain storage events. Please see OEM Technical Guide for more information.

2.7.5 Utilities

2.7.5.1 Install/Uninstall Utility

The Intel® RSTe NVMe will be available through the use of an install package. The install package will automatically install the proper RSTe driver and GUI that corresponds to the OS being installed on. There will also be a mechanism available to uninstall the driver and GUI.

NOTE: Great care must be taken when trying to perform the uninstall process. Removal of the driver could result in a system crash that could require a complete reinstallation of the operating system.

[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.*



3 Product Certifications

3.1 WHQL

The Intel® RSTe drivers are required to be logo certified on the Microsoft Windows* WLK 1.5 test suite for storage device drivers running on legacy OSs.

The Intel® RSTe drivers are required to be logo certified on the Microsoft Windows* WLK 2.0 test suite for storage device drivers.

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

4 Appendix A

Table 4-1. Relevant Specifications

ATA/ATAPI-7 (http://www.t13.org/Documents/UploadedDocuments/project/d1532v3r4a-ATA-ATAPI-7.pdf)
ATA Command Set 2 (http://www.t13.org/Documents/UploadedDocuments/docs2009/d2015r2-ATAATAPI_Command_Set_-_2_ACS-2.pdf)
ATA8-ACS-8 (http://www.t13.org/Documents/UploadedDocuments/docs2007/D1699r4c-ATA8-ACS.pdf)
SATA 1.0 Specification (http://www.serialata.org)
SATA II Specification (http://www.serialata.org)
SATA 3 (http://www.sata-io.org/documents/SATA-Revision-3.0-Press-Release-FINAL-052609.pdf)
Serial Attached SCSI - 2 (SAS-2) (http://www.t10.org)

Table 4-2. Relevant Documents

CDI / IBL #	Title/Location
Reference Documents	
441979	**Intel® 6 Series Chipset/ Intel® C200 Series Chipset/ Intel® C600 series chipset Platform Controller Hub (PCH) BIOS Specification Update – NDA**
473093	Intel Patsburg PhyTune Tool -RC Ver 2.0 **Note: This package contains the PhyTune tool along with the SASAddress efi utility. Please refer to the documentation included in the package for additional information.**
453321	**Intel® Server Platform Services Manageability Engine Firmware for Patsburg chipset Product Line Firmware Startup Guide**
454672	**Patsburg Chipset SPI Programming Guide**
450911	**Patsburg Chipset External Design Specification (EDS)**
445721	**Patsburg Chipset External Design Specification (EDS) Specification Update - NDA**
458143	**Sandy Bridge-E Processor External Design Specification (EDS) - Volume One of Two**
459924,	**Sandy Bridge-E Processor External Design Specification - Volume Two of Two**



CDI / IBL #	Title/Location
30051	RS – Intel® 6 Series Chipset/ Intel® C200 Series Chipset/ Intel® C600 series chipset Platform Controller Hub (PCH) BIOS Spec ***Contact your Intel FAE to get access to this document through Anacapa***
Kit 33272	Intel® Server Platform Services Alpha SPS_02.01.01.009.0 Note: This package is the Intel® Server Platform Services Manageability Engine Firmware for Intel® C600 series chipset Product Line - Alpha Full Release and contains key tools such as FITc and fpt for the Intel® C600 series chipset **This document can be downloaded from ARMS/VIP**
511555	Intel® C610 Series Chipset and Intel® X99 Chipset Platform Controller Hub (PCH)
558771	Skylake Server Processor External Design Specification (EDS)
546955	Skylake Platform Controller Hub (SKL PCH-H) External Design Specification (EDS)
509027	Haswell-EN/EP/EP 4S/EX Processor BIOS Writer's Guide (BWG)
541497	Broadwell-EP/EP 4S/EX Processor External Design Specification (EDS)
558749/ 486708	Intel® 8 Series/C220 Series Chipset Family Platform Controller Hub (PCH) External Design Specification (EDS)
547817	Lewisburg Platform Controller Hub External Design Specification
RST Legacy Requirements Documents	
443658	***Intel® Rapid Storage Technology (Intel® RST) 9.6 Product Requirements Document (PRD) Differences Document*:** http://www.intel.com/cd/edesign/library/asmo-na/eng/443658.htm
409399	*Intel® Rapid Storage Technology 9.5 Product Requirements Document (PRD) Différences Document :* **http://www.intel.com/cd/edesign/library/asmo-na/eng/409399.htm**
375707	*Intel® Matrix Storage Manager 8.5 Product Requirements Document (PRD) Differences Document :*
373388	*Intel® Matrix Storage Manager 8.0 Product Requirements Document (PRD) :* **http://www.intel.com/cd/edesign/library/asmo-na/eng/373388.htm**
558079	*Intel® Xeon® Processor E3-1200 v4 Product Family External Design Specification (EDS)*

CDI / IBL #	Title/Location
Technical Guides / White papers	
N/A	*Intel® Rapid Storage Technology OEM Technical Guide:* (location) Intel® Validation Internet Portal (VIP **https://platformsw.intel.com/**) in each major or maintenance release version of Intel® Rapid Storage Technology release.
445153	*Reference Manual for Safe Removal with Link Power Management (LPM) on Hot Plug Capable Port (HPCP)* : **http://www.intel.com/cd/edesign/library/asmo-na/eng/445153.htm**



5 Appendix B

Table 5-1. 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
MigrationStarted	Info	Volume {0}: Data migration in progress.	Yes	No	No
InitializeStarted	Info	Volume {0}: Initialization in progress.	Yes	No	No

Event Type	Event Level	String	Event Displayed		
			NAI[1] (Notification Area Icon)	Event Log	E-Mail Notify[2]
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



6 Appendix C

6.1 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.

6.2 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 1 shows this structure.

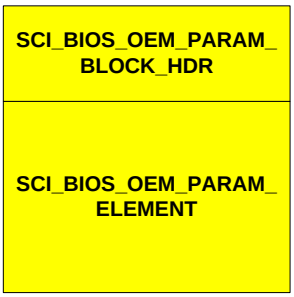


Figure 1: Structure of OEM Parameter Block

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 6.3 BCFS Bit Settings 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;
```

6.3 BCFS Bit Settings

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 delay will be set to 2 seconds
		Raid type enabled	0x1	
1	Enable/disable Raid1	Raid type disabled	0x0	
		Raid type enabled	0x1	
2	Enable/disable Raid10	Raid type disabled	0x0	
		Raid type enabled	0x1	

BCFS bit number	Bit meaning	Values		Additional info
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	
6-8	RESERVED			
9	Dnable/disable RSTe caching	Feature disabled	0x0	
		Feature enabled	0x1	
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	



BCFS bit number	Bit meaning	Values		Additional info
		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			

6.4 Recommended Location in SPI Flash

Figure 2 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 *6.5 OEM Parameter and SPI Flash Tools*.

Note:	The recommended SPI flash format is descriptor mode.

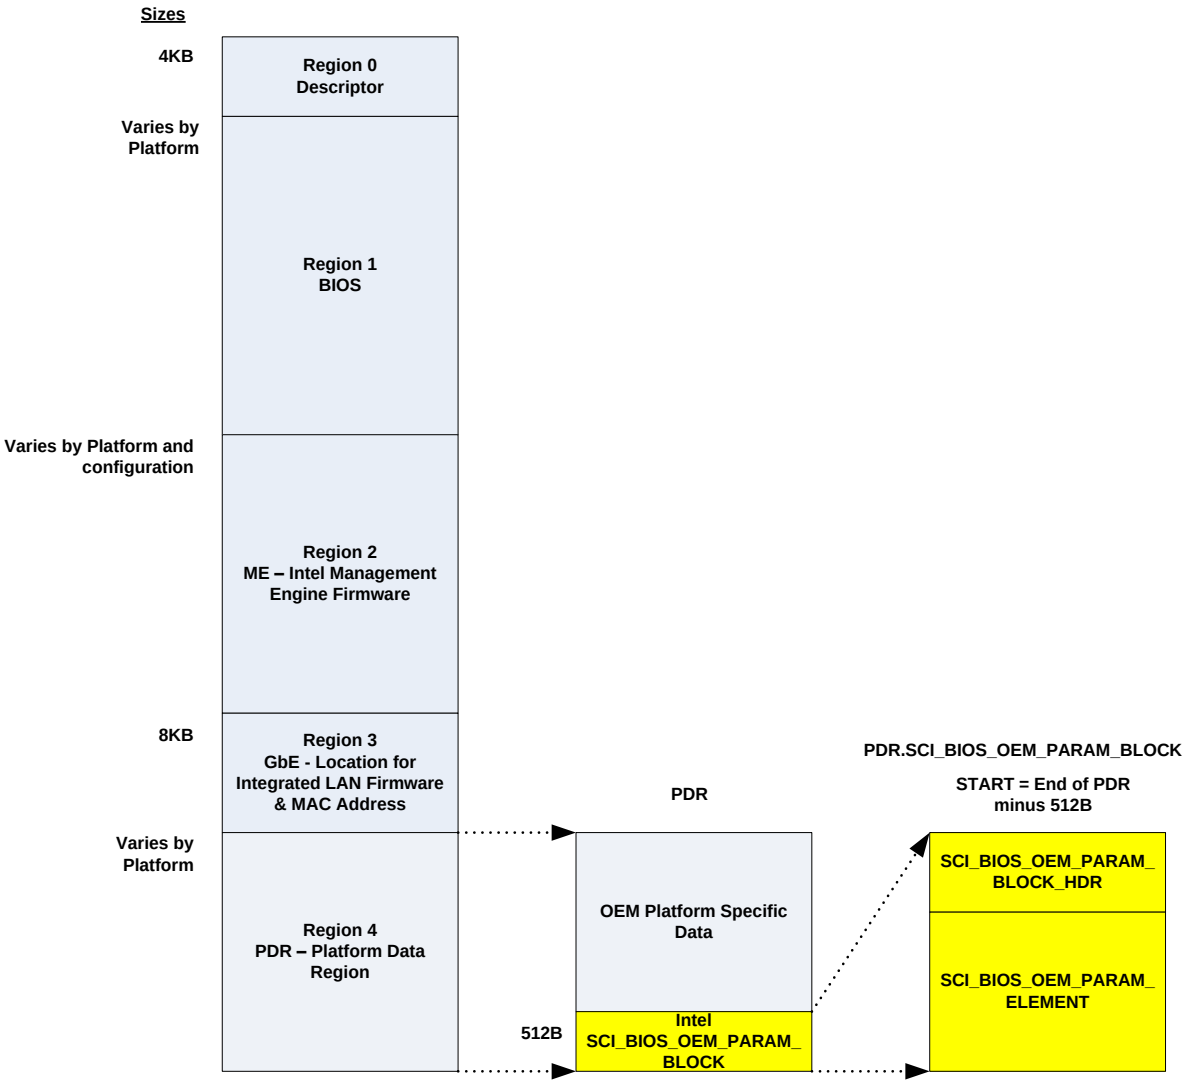


Figure 2 Recommended OEM Param Location in SPI Flash

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 *6.6 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.



6.5 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 at 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 *6.2 Structure of OEM Parameter Block* for a detailed description of the PHY parameters.

6.5.1 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 6.4 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.



6.5.2 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.

6.5.3 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 6.4 Recommended Location in SPI Flash.

6.6 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 6.4 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.

6.6.1 Legacy OROM

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

6.6.1.1 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

6.6.1.2 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 RAID configuration 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. There will 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)

6.6.2 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];
};
```

6.7 Intel® C610 series Chipset Legacy OROM Boot Option

6.7.1 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 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 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 another device. 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. There will 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)

7 Appendix D

7.1 Hardware Compatibly

7.1.1 External Hardware Compatibility

The embedded file indicates the current list of external hardware used in validation and is subject to change without notice. Please contact your factory representative for questions on any specific hardware item.

Enterprise SAS Drives

Vendor	Family	Model Name/Number
Fujitsu	AL9Se Series (2.5")	MAY2036RC
Fujitsu	AL9LX Series (3.5")	MAX3036RC,
Fujitsu	AL10Se Series (2.5")	MBB2 Series
Seagate	SAS	Barracuda ES.2 7.2k rpm
Seagate	SAS	Cheetah 15k.6 (3.5")
Seagate	SAS	Cheetah 15K.4 (3.5")
Seagate	SAS	Cheetah 15K.5 (3.5")
Seagate	SAS	Cheetah 15K.7
Seagate	SAS	Savvio 10K.1 (2.5")
Seagate	SAS	Savvio 10K.2 (2.5")
Seagate	SAS	Savvio 15K.1 (2.5")
Seagate	SAS	Cheetah NS
Hitachi	Ultrastar 15K147 3.5" (Viper A')	HUC101473CSS300,
Hitachi	Ultrastar 15K147 3.5" (Viper B)	HUS153014VLS300, HUS153073VLS300
Hitachi	Ultrastar C10K147 2.5" (Cobra B)	HUC101473CSS300,



Enterprise SATA Drives

Vendor	Family	Model Name/Number
Fujistu	A160 (2.5") 7200 RPM FDE Option Extended Duty	MHZ2080BK
Hitachi	Ultrastar A7k1000 (3.5") 7.2rpm	
Seagate	Barracuda 7200.10 Serial ATA	
Seagate	Barracuda 7200.11 Serial ATA	
Seagate	Barracuda ES	
Western Digital		WD1002FAEX
Western Digital		WD6000HLHX

Expanders and Enclosures

Vendor	Model Number
AIC	XJ1100
Xyratex	RS1603X
Supermicro	CSE-M28x
Promise	Vtrak E-Class E310
	Vtrak J-class
Supermicro	SC836E1-R800V