

Intel® Rapid Storage Technology enterprise (RSTe)

Technical Product Specification (TPS) Document

September, 2013

Revision 1.7

Intel Confidential



INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

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 © 2013, Intel Corporation. All rights reserved.



Contents

1	Introduction	7
1.1	ACRONYMS	7
1.2	Intended Use	9
1.3	Intended Audience	9
2	Product Overview	10
2.1	Supported Platforms/Chipsets/SKUs	11
2.1.1	Grantley Platforms	11
2.1.2	Romley Platforms	11
2.2	Intel® C600 series chipset SKU's	11
2.2.1	Upgrade ROM Options	12
2.3	Intel® RSTe Major Components	12
2.3.1	Intel® RSTe Pre-OS Images	12
2.3.2	Intel® RSTe OS Runtime Drivers	13
2.3.3	Intel® RSTe RAID Tools and Utility	14
2.4	Unsupported Features	15
2.5	Key Product Features	16
2.5.1	Operating Systems Supported	16
2.5.2	Open Source Operating Systems Support	17
2.5.3	Intel® RSTe Pre-OS Package	17
2.5.4	Intel® RSTe OEM Parameters	20
2.5.5	Intel® RSTe Configuration Tools	20
2.5.6	Intel® RSTe Management Tools	21
2.5.7	Intel® RSTe System Configurations supported	22
2.5.8	Software RAID Functional Support	24
2.6	Intel® C600 series chipset OEM Parameters	31
2.7	Structure of OEM Parameter Block	32
2.8	BCFS Bit Settings	39
2.9	Recommended Location in SPI Flash	41
2.10	OEM Parameter and SPI Flash Tools	43
2.10.1	PHY Tune Tool	43
2.10.2	SAS Address Tool (sasaddresses.efi)	44
2.10.3	FITC Tool	44
2.11	Required BIOS Services	44
2.11.1	Legacy OROM	44
2.11.2	UEFI	49
2.11.3	Utilities	53
3	Features, Specifications, and Specification Updates	54
3.1	New Features and Supported Specifications Introduced in Intel® RSTe 3.8	54
3.1.1	Beta Support for Grantley Platforms with Intel® C610 series chipset sSATA Controller	54
3.1.2	Windows® 8.1 and Server 2012 R2 Support	54
3.1.3	RSTe INF File Split	54
3.1.4	Device Information Display in UEFI	55
3.2	New Features and Supported Specifications Introduced in Intel® RSTe 3.7	55



	3.2.1	Human Interface Infrastructure (HII) support	55
	3.2.2	Expander Exposed to Device Manager	56
	3.2.3	Trim on RAID0, RAID1 and RAID 10.....	56
	3.2.4	RSTe UEFI Command Line Tool to support Supports Testing SGPIO	56
	3.2.5	RSTe GUI Supports Showing SAS Link Width.....	56
	3.2.6	SCU Controller Legacy Option ROM Splash Screen Timing Adjustment.....	56
	3.3	Updated Features/Specifications Implemented in Intel® RSTe 3.8	56
4		Product Certifications.....	57
	4.1	WHQL	57
5		Appendix A.....	58
6		Appendix B.....	60
7		Appendix C	62
	7.1	Intel® C600 series chipset OEM Parameters	62
	7.2	Structure of OEM Parameter Block	62
	7.3	BCFS Bit Settings.....	70
	7.4	Recommended Location in SPI Flash	72
	7.5	OEM Parameter and SPI Flash Tools.....	74
		7.5.1 PHY Tune Tool.....	74
		7.5.2 SAS Address Tool (sasaddresses.efi)	75
		7.5.3 FITC Tool.....	75
	7.6	Required BIOS Services	75
		7.6.1 Legacy OROM.....	75
		7.6.2 UEFI	80
8		Appendix D.....	81
	8.1	Hardware Compatibly	81
		8.1.1 External Hardware Compatibility.....	81



Figures

Figure 1 Email Format	52
Figure 2: Structure of OEM Parameter Block	63
Figure 3 Recommended OEM Param Location in SPI Flash	73

Tables

Table 1-1. Common Acronyms	7
Table 5-1. Relevant Specifications	58
Table 5-2. Relevant Documents	58
Table 6-1. Storage System Events Detected by Monitor Service (IAStorDataMgrSvc)	60



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

§



1 Introduction

This is the Technical Product Specifications for the Intel® Rapid Storage Technology enterprise (Intel® RSTe) product. The intent of this document is to present the functional requirements (or features) that make up this product. Although this document has references to Linux, the main intent is to describe the features that will be in the Pre Operating System (Pre-OS) and Windows® based Intel® RSTe drivers.

1.1 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.
EN	Entry Server
EP	Efficient Performance
GB	Giga-byte
HDD	Hard Disk Drive

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



Term	Definition
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
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
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 Enterprise Server
SMIS	Storage Management Initiative Specification
SSD	Solid State Device - non volatile memory



Term	Definition
UI	User Interface
UEFI	Unified Extensible Firmware Interface

1.2 Intended Use

This document is intended to give detailed information on the technical features of the Intel® Rapid Storage Technology enterprise (Intel® RSTe) product.

1.3 Intended Audience

The intended audience of this document is OEMs, ODMs, IHV, System Integrators, and end users requiring detailed information of the new features and technical specifications of the Intel® RSTe product. It contains only information pertaining to the initial features introduced with the Intel® RSTe product.



2 *Product Overview*

The Intel® RSTe product supports enterprise storage platforms that may contain one Advanced Host Controller Interfaces (AHCI)/SATA and the Storage Controller Unit (SCU) (Romley) or two SATA Controllers (Grantley). Intel® RSTe is the enterprise version of Intel® Rapid Storage Technology (RST) for supported systems equipped with Serial ATA (SATA) devices, Serial Attached SCSI (SAS) devices, and/or solid state drives (SSD) to enable an optimal enterprise storage solution. It offers value-add features such as RAID and advanced SAS* and/or SATA* capabilities for various Operating System (for detailed OS support, review the Release Notes for each software release).

Some of the RAID features supported by Intel® RSTe include RAID level 0 (striping), RAID level 1 (mirroring), RAID level 5 (striping with parity) and RAID level 10 (striping and mirroring).

The new features introduced with Intel® RSTe include but are not limited to:

- RAID support for SAS devices
- System Management (i.e. CIM provider and CSMI)
- Read Patrol

This document outlines and explains the key features available with the Intel® RSTe product.



2.1 Supported Platforms/Chipsets/SKUs

2.1.1 Grantley Platforms

- Intel® Server Series (Haswell) 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
- [Arandas](#) 4S EV/SV VVP platform

2.1.2 Romley Platforms

- Intel® Server Series (Sandy Bridge) Chipset-based platform CRBs with the Intel® C600 series chipset (PCH)
 - Waimea Bay (High End Desktop, 1 CPU)
 - River City (Workstation, 2 CPUs/Sockets)
 - Harbor City (EN Server, 2 CPUs/Sockets)
 - Rose City (EP, 2 CPU)
 - Potter City (EX, 4 CPUs/Sockets)

2.2 Intel® C600 series chipset SKU's

SKUs	-A	-B	-D	-T
Features	(SATA Only)			
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 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 -A 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)

2.3 Intel® RSTe Major Components

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

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

2.3.1 Intel® RSTe Pre-OS Images

Intel® RSTe will provide a Pre-OS package for BIOS vendors and OEM/ODMs developing their own BIOS. The Pre-OS package will contain the binary images necessary to be compiled into the system BIOS to support Intel® C600/610 series chipset based platforms.



2.3.1.1 Intel® RSTe RAID Legacy Option ROMs

There are three Legacy RAID OROM images that support Int13 BIOS environments. These images also include the RSTe RAID Pre-OS boot (menu-driven) configuration utility. The three Legacy RAID OROM images are for:

1. Intel® C600 series chipset: There is one image (each) for the SATA and SCU controllers.
2. Intel® C610 series chipset: There is one additional image for the sSATA controllers. This is used in conjunction with the SATA image for Intel® C600 series chipset.

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

There are three sets pre-OS SATA UEFI drivers to support a UEFI BIOS environment. The three RAID UEFI images are:

1. Intel C600 series chipset: There is one image (each) for the SATA and SCU controllers.
2. Intel C610 series chipset: Adds an image for the sSATA controllers. This is used in conjunction with the SATA image for Intel® C600 series chipset.

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

Note: The SCU controller does not currently support a non-RAID mode of operation analogous to "AHCI Mode" on the AHCI SATA controllers

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, sSATA and SCU controllers.

NOTE: RAID spanning across controllers is not supported.

2.3.2.2 Drivers Supporting the Intel® C610 series chipset driver

Along with the SATA (AHCI) Controller drivers described above, 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.

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 series chipsets.

2.3.3.2 Intel® RSTe Pre-OS Tools

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

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 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 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 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 legacy RAID OROM and the RSTe SATA RAID UEFI driver have been properly implemented/incorporated into the system BIOS. These tools support 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).

2.3.3.2.3 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 Unsupported Features

The following is a list is a non-comprehensive list of features that will 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 individual SATA controllers (Intel® C610 series chipset)
- Port-Multipliers
- SATA RAID Legacy OROM support for Optical Devices
- SCU RAID Legacy OROM support for Optical Devices
- SCU UEFI Driver support for Optical Devices
- Intel® Rapid Recovery Technology (Intel® RRT)
- Expander support on SATA controller
- 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.



2.5 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 0/1/5/10 on SAS	<ul style="list-style-type: none">• SCU support for Matrix RAID 0/1/5/10• Pass-through drives• Hot Plug with I/O• Hot Spare Disk• Auto Rebuild on Hot Insert• Rebuild & Migration Check Pointing• NCQ (SATA) and CQ (SAS) support• UEFI using common metadata• SAS Expanders• SMART Support• Bad Block Management• SAS & SATA controller configuration rules• SAS & SATA drive roaming• RAID Volume roaming between Linux* and Windows*• On Line Capacity Expansion• Large Stripe Size Support• RAID-Ready	<ul style="list-style-type: none">• Disk Coercion• Manual & Auto Rebuild• Instant Initialization• Read Patrol• SGPIO for SAS & SATA• volume creation/verify• Selectable Boot Volume• 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	<ul style="list-style-type: none">• Install/Uninstall Utility• Configuration and Management Utilities	
Intel® RSTe (Window only)	Intel® RSTe provides software for high-performance for SAS and SATA RAID + <ul style="list-style-type: none">▪ Matrix RAID for two RAID volumes on single array▪ Improved user interface for enhanced usability	

2.5.1 Operating Systems Supported

2.5.1.1 Microsoft OS Support

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

- Windows* Vista
- Windows* 7
- Windows* Server 2008 SP2
- Windows* Server 2008 R2 (64 bit only)
- Windows* Server 2003
 - 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 implemented following the release of Windows 2003. Microsoft has made a Hot-Fix available for download. The issue will only be

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



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](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* PE 3.0
- Windows* 8
- Windows* Server 2012
- Windows* 8.1
- Windows* Server 2012 R2

2.5.2 Open Source Operating Systems Support

Intel® RSTe will support the latest minor revision (dot release) of the latest shipping major revision of the Open Source OS (such as Linux).

- Red Hat Enterprise Linux Server* (x32 and x64) (Contact Red Hat for details)
- Novell SuSE Linux Enterprise Server* (x32 and x64) (Contact Novell for details)

Core functionality support will also be provided for the N-1 minor revision via a driver back port, Zstream and Driver Update Disk processes beginning with the following:

- Red Hat Enterprise Linux Server* (x32 and x64) (Contact Red Hat for details)
- Novell SuSE Linux Enterprise Server* (x32 and x64) (Contact Novell for details)

Support for Open Source version greater than the above specified N-1 version will not be supported.

2.5.3 Intel® RSTe Pre-OS Package

The Intel® RSTe will provide binary images that will be compiled into the system BIOS to provide a Pre-Boot driver to the SATA, sSATA, and SCU controllers. The BIOS vendor (or the OEM/ODMs developing their own BIOS) will take these images and compile them into their BIOS package for Intel® C600 series chipset based 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). The Intel® RSTe Pre-OS package will support the following:

- An SATA Controller configured for RAID Mode
- Using SATA Legacy OROM to creating a bootable RAID volume on drives connected to a SATA controller
- Booting from a bootable drive or RAID volume on drives connected to a SATA controller
- SCU Controller

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



- Using SCU Legacy OROM to create a bootable RAID volume on drives either directly attached, or behind a SAS expander connected, to the SCU Controller
- Using a UEFI driver to support UEFI boot environment to provide access the devices attached to either a SATA or SCU controllers.
- Using a UEFI Command Line Interface (CLI) utility to create bootable RAID volumes on drives attached to either a SATA or SCU controller.

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 some configuration options to the BIOS configuration utility (not addressed in this document). The BIOS configuration utility may provide the following configuration options:

- An option for the user to select which controller they would like to boot from (SATA or SCU).
- An option to select which pre-boot environment to use (Legacy Option ROM or UEFI)
- An option to select which mode to configure the Intel® C600 series chipset SATA 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 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 RAID Legacy Option ROM (or Intel® RSTe RAID UEFI driver) will be loaded and used by the system BIOS to access the devices connected to the SATA controllers.

When the SCU controller is selected as the boot controller, the system BIOS will need to load and use either the Intel® RSTe SCU RAID Legacy Option ROM or the Intel® RSTe UEFI driver to access the drives connected to the SCU controller. The Intel® RSTe Pre-boot package will support only one mode of operation (RAID Mode) on the SCU controller.

The Intel® RSTe Pre-boot package will support only one boot method, either Legacy OROM or UEFI, but not both.

The Intel® SATA RAID Legacy Option ROMs only supports drives that are direct attached to the SATA controllers.

The Intel® SCU RAID Legacy Option ROM supports drives that are either directly attached, or located in a SAS expander connected, to the SCU Controller.

2.5.3.1 Intel® RSTe SATA & sSATA RAID Legacy Option ROMs

The Intel® RSTe will support a SATA and an sSATA RAID Legacy Option ROM. The BIOS configuration utility may provide an option to select the one of the SATA controllers as the boot controller. When the system is configured to boot from a SATA controller in RAID mode, an Intel® RSTe SATA RAID Legacy Option ROM will be loaded and will provide the interface to the drives attached to the SATA controllers. The Intel® RSTe SATA RAID Legacy Option ROM will only support drives directly attached to the SATA controller.

While booting, a BIOS Splash Screen will appear on the display when there are a two or more drives attached to a controller. If there two or more drives attached to each SATA and sSATA controller, two separate BIOS Splash Screens will appear. The Splash Screen will show what is attached to the



controller. There is also an option to stop the booting process and enter into the Intel® RSTe SATA RAID Legacy Option ROM user interface. This is done by pressing the [CTRL]-I key combination. Once entered, user interface will allow the user to create/manage/delete RAID volumes on drives attached to the SATA controllers. This is mainly used to create a RAID volume that can be used as the system OS boot device.

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

2.5.3.2 Intel® RSTe SCU RAID Legacy Option ROM

Intel® RSTe will provide support for an SCU RAID Legacy Option ROM. 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 Intel® RSTe SCU RAID Legacy Option ROM will be loaded and will provide the interface to the drives attached to the SCU controller. The Intel® RSTe SCU RAID Legacy Option ROM will manage and boot from drives that are:

- Directly attached to the SCU controller (all 8 ports)
- Located in an expander (single level) attached to the first 4 ports to the SCU controller.
 - The first 8 ports (port0 thru port7) of the expander will be 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

When using the RSTe SCU RAID Legacy Option ROM, the system BIOS must support the INT15 function call to obtain the OEM Parameter information programmed into SPI Flash. Please reference the platform's BIOS Writer's Guide for more information.

NOTE: The RSTe SCU RAID Legacy Option ROM may not work properly if the OEM Parameters have not been properly programmed into SPI Flash.

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 RAID Legacy Option ROM user interface. This is done by pressing the [CTRL]-I key combination. Once entered, the user interface will allow the user to create/manage/delete RAID volumes on drives attached to the SCU controller. This is mainly used to create a RAID volume that can be used as the system OS boot device.

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

2.5.3.3 Intel® RSTe SATA and sSATA RAID UEFI Drivers

Intel® RSTe will provide a SATA RAID UEFI driver for each of the SATA and sSATA controllers. The drivers will provide the interface to a controller and connected devices. The Intel® RSTe SATA UEFI RAID Driver will support only drives directly attached to a SATA controller



2.5.3.4 Intel® RSTe SCU RAID UEFI Driver

Intel® RSTe will provide support for an SCU RAID UEFI driver. This driver will provide the interface driver to the devices connected to the SCU controller. The Intel® RSTe SCU RAID UEFI Driver will support directly attached drives and will provide at least one level of SAS expander support.

2.5.4 Intel® RSTe OEM Parameters

Intel® RSTe will provide support for an SCU OEM Parameter field to reside in the SPI Flash. This field must be properly configured in order for the SCU Controller to initialize properly. Please reference the platform's BIOS Writer's Guide for more information. This driver will provide the interface driver to the devices connected to the SCU controller. The Intel® RSTe SCU RAID UEFI Driver will support directly attached drives and will provide at least one level of SAS expander support.

Note: The SCU Controller may not work properly if the SCU OEM Parameter field is not programmed properly.

2.5.5 Intel® RSTe Configuration Tools

The Intel® RSTe 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.

2.5.5.1 Intel® RSTe UEFI Command Line Interface (CLI) Utility

Intel® RSTe will provide support for a UEFI command line interface utility. An Intel® RSTe UEFI Command Line Interface (CLI) utility will be made available to manage RAID volumes when booted into the UEFI environment. The Intel® RSTe UEFI CLI utility will need to be **launched from USB drive**.

Note: This Intel® RSTe UEFI CLI utility will provide a command line interface to the user to allow to create/manage/delete RAID volumes on drives attached to either the SATA or SCU controllers. The utility will access the appropriate controller and is only available when they system boots into the UEFI shell. This is mainly used to create a RAID volume that can be used as the system OS boot device. A Microsoft* signed version of the Intel® RSTe UEFI RAID Configuration utility will also be made available to support secure boot environments. When the system is configured to boot using UEFI, the user must boot into the UEFI Shell to use this utility to manage the RAID volumes (check the status, initiate rebuilds, expand, etc.).

2.5.5.2 Intel® RSTe UEFI Human Interface Infrastructure

Intel® RSTe will provide support for a UEFI HII RAID management interface. 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.

2.5.5.3 Intel® RSTe RAID Configuration Utilities

Intel® RSTe will provide support for a DOS base command line utility that can be used in conjunction with the Legacy Option ROMs. There will be an Intel® RSTe RAIDCFG utility each of the SATA controllers and one for the SCU controller. The utility is accessed through DOS bootable media (floppy drive or USB drive) and 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)

2.5.5.4 Intel® RSTe Command Line Interface (CLI) Application

Intel® RSTe 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 (similar to the RAIDCfg utility) 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)

2.5.5.5 Intel® RSTe Graphical User Interface (Intel® RSTe GUI)

Intel® RSTe 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 ACHI 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).

Note: Intel® RSTe GUI will provide RAID management functionality for up to 32 drives.

2.5.6 Intel® RSTe Management Tools

2.5.6.1 Common Information Model (CIM)

Intel® RSTe 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).

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

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



2.5.6.2 Common Storage Management Interface (CSMI)

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

2.5.7 Intel® RSTe System Configurations supported

This section addresses to physical components of the system configuration supported by Intel® RSTe 3.0.

2.5.7.1 SATA Controller Support

Intel® RSTe will provide support for managing RAID volumes on drives attached to the SATA ports.

2.5.7.2 Intel® C610 series chipset sSATA Controller Support

Intel® RSTe 3.8 will provide support for managing RAID volumes on drives attached to the sSATA ports of the Intel® C610 series chipset.

2.5.7.3 SCU Controller Support

Intel® RSTe will provide support for managing RAID volumes on drives attached to the SCU ports.

2.5.7.4 SCU OEM Parameters

Intel® RSTe 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.

2.5.7.5 SAS Expander Support

Intel® RSTe 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.

2.5.7.6 Maximum Hard Drives supported on Intel C600 series chipset Based Platforms

Intel® RSTe 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.

2.5.7.7 Pass-through drives

Intel® RSTe will support the ability to expose non-RAID configured disks (pass-through) to Host OS. Please refer to section 2.5.7.6 for additional information.

2.5.7.8 SCU Controller RAID Management Limitations

Intel® RSTe 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.



2.5.7.9 Hot Plug

Intel® RSTe will support the ability to Hot Plug (remove and replace) disk drives on the SATA controllers whether or not I/O is being processed, provided that the capabilities are enabled in the BIOS.

Intel® RSTe will support the ability to Hot Plug (remove and replace) disk drives attached to the SCU controller whether or not I/O is being processed.

2.5.7.10 SCU & SATA drive roaming

Intel® RSTe will support the ability to move RAID volumes on SATA drives between the SATA and SCU controllers and have RAID arrays and volumes recognized, available and bootable via common metadata.

2.5.7.11 Volume Roaming between Linux* and Windows*

Intel® RSTe 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.

2.5.7.12 SGPIO on SATA Controller

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

2.5.7.13 SGPIO on SCU

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

2.5.7.14 NCQ (SATA) and CQ (SCU) support

Intel® RSTe will support Native Command Queuing (SATA) and Command Queuing (SAS).

2.5.7.15 SCSI Enclosure Service (SES) v2

Intel® RSTe 3.7 will provide support management of enclosures that are compliant with SES (SCSI Enclosure Services) v2 attached to the SCU controller. Intel® RSTe 3.7 also supports in-band management to SES compliant expanders attached to the SCU.

2.5.8 Software RAID Functional Support

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

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



2.5.8.1 Matrix RAID

Intel® RSTe 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.

2.5.8.2 RAID 0/1/5/10 Volumes

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

2.5.8.3 Software RAID 5 on SCU

Intel® RSTe 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.

2.5.8.4 Simultaneous RAID Arrays

Intel® RSTe 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.

2.5.8.5 Disk Coercion

Intel® RSTe 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.

2.5.8.6 Hot Spare Disk

Intel® RSTe 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 and SCU controllers.

2.5.8.7 Auto Rebuild on Hot Insert

Intel® RSTe 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.

2.5.8.8 Manually Invoked Rebuild

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



2.5.8.9 RAID SMART Support

Intel® RSTe 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.

2.5.8.10 RAID-Ready Mode

A RAID-Ready system refers to a system that has been configured to support Intel® RSTe 3.0. 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 controllers and it is in RAID mode
- The system is configured to boot off the SCU controller

Intel® RSTe will support an Intel® C600 series chipset based platform configured in RAID-Ready mode.

2.5.8.11 RAID Volume Creation with Data Preservation

Intel® RSTe 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 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 3.0:

- Individual pass-through to 2 16 drives for RAID 0
- 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

2.5.8.12 Instant Initialization

Intel® RSTe 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.

2.5.8.13 RAID Level Migrations

The RAID level migration feature in Intel® RSTe product will enable the ability to convert the contents of a drive (attached to the SATA 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 3.0:

- 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
- 4-drive RAID 10 to 4 - 32 drive RAID 5

2.5.8.14 RAID Reconfiguration (Stripe size)

Intel® RSTe 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

2.5.8.15 Expanded Stripe Size

Intel® RSTe 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

2.5.8.16 Online Array / Volume Capacity Expansion

Intel® RSTe 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.

2.5.8.17 Read Patrol

Intel® RSTe 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 either the SATA or SCU controllers. 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.



2.5.8.18 Verify and Repair

Intel® RSTe 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 Data on the mirror drive is compared to data on the source drive.	Bad blocks are reassigned. 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. Parity is recalculated and compared to the stored parity for that stripe.	Bad blocks are reassigned. 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. Data on the mirror is compared to data on the source.	Bad blocks are reassigned. 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.

2.5.8.19 Check Pointing

Intel® RSTe 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.

2.5.8.20 Bad Block Management

Intel® RSTe 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.

2.5.8.21 Dirty Stripe Journaling

Intel® RSTe 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

2.5.8.22 Partial Parity Logging (PPL)

Intel® RSTe 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.

2.5.8.23 OS Installation

Intel® RSTe 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 or SCU controllers.

2.5.8.24 Selectable Boot Volume

Intel® RST 3.0 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.

2.5.8.25 Auto Rebuild

Intel® RSTe 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

Automatic rebuild support will default to OFF for Intel® RSTe3.0 and can be enabled through the Intel® RSTe GUI.

2.5.8.26 Error Threshold Monitoring/Handling

Intel® RSTe 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)

2.5.8.27 Unified Extensible Firmware Interface (UEFI)

Intel® RSTe will support UEFI for the SCU and SATA controllers using common metadata.

2.5.8.28 Disk Write Cache

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

2.5.8.29 RAID Volume Read Cache

Intel® RSTe 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.

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



2.5.8.30 Write Back Cache

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

2.5.8.31 Volume Cache Increase

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

2.5.8.32 RAID Volume Size

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

2.5.8.33 RAID Boot Volume Size

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

2.5.8.34 Disk Monitor Service

Intel® RSTe 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.

2.5.8.35 Failed Drive Reinsertion

Intel® RSTe 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.

2.5.8.36 Drives Supported

Intel® RSTe will provide support for current production SATA drives from “any” manufacturer on the SATA controllers. SAS and SATA drives supported on the SCU controller are outlined in Appenix C

2.6 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.

2.7 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 2 shows this structure.




Figure 2: 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 SCL_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 (SCL_PREBOOT_SOURCE_OROM and SCL_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 7.3 BCFS Bit Settings below for details on the allowed settings.

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

} SCL_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[SCL_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[SCL_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;
```

2.8 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	



BCFS bit number	Bit meaning	Values		Additional info
2	Enable/disable Raid10	Raid type disabled	0x0	
		Raid type enabled	0x1	
3	Enable/disable Raid5	Raid type disabled	0x0	
		Raid type enabled	0x1	
4	RESERVED			
5	Enable/disable UI	Feature disabled	0x0	If you set it to 0 (disable) bits 10-14 are ignored
		Feature enabled	0x1	
6-8	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 occurred
		Show only if error	0x01	
		Never show UI	0x10	
		Show always	0x11	
15	RESERVED			

2.9 Recommended Location in SPI Flash

Figure 3 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 7.5 *OEM Parameter and SPI Flash Tools*.

Note: The recommended SPI flash format is descriptor mode.

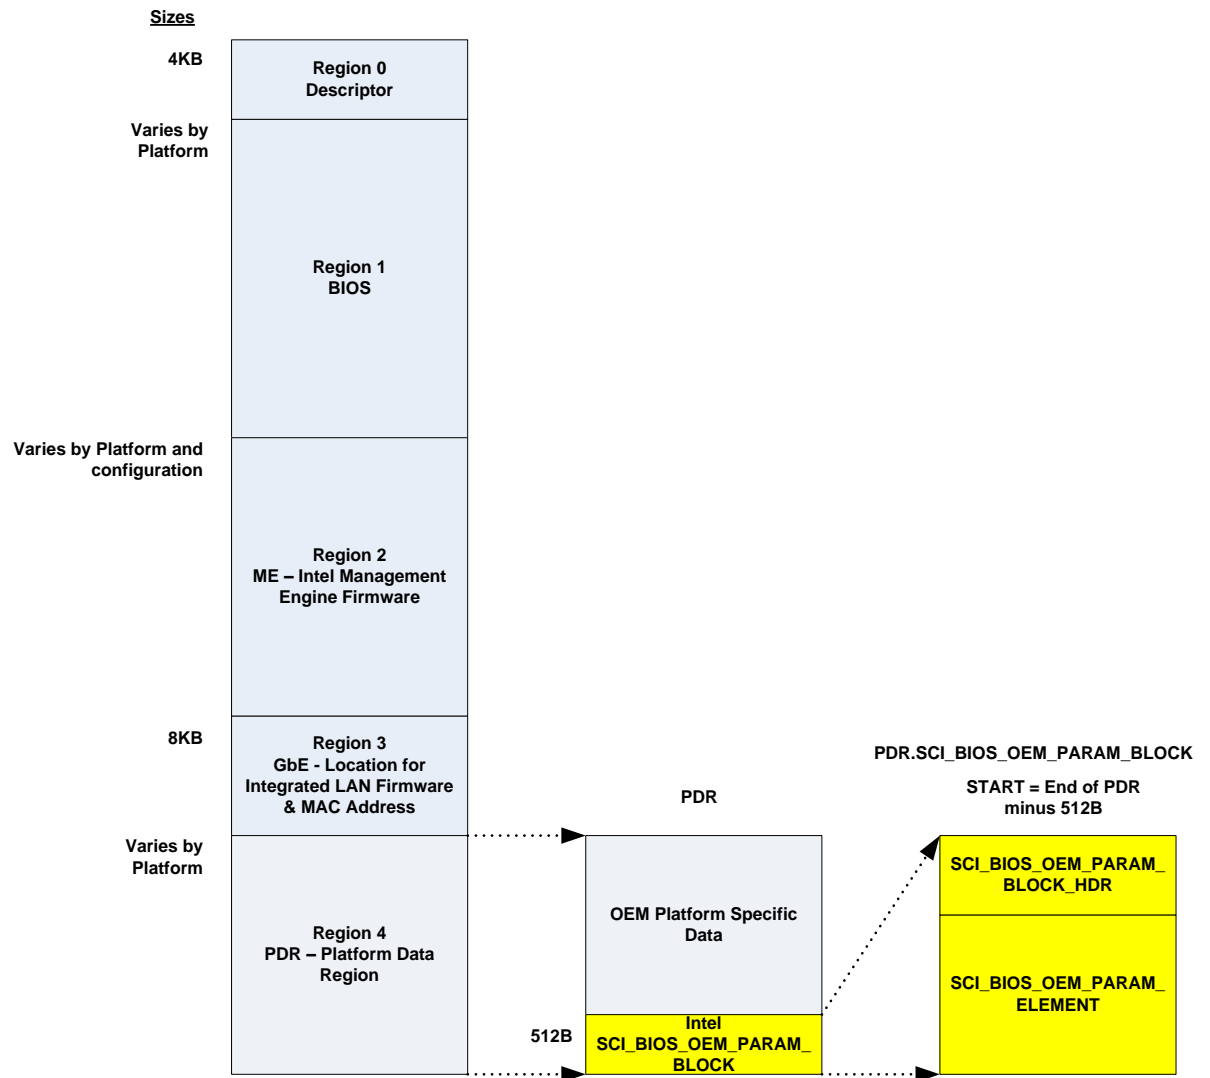


Figure 3 Recommended OEM Param Location in SPI Flash

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 7.6 *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.



2.10 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 7.2 *Structure of OEM Parameter Block* for a detailed description of the PHY parameters.

2.10.1 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. Each individual PHY can have its own unique settings that best fit its situation. As indicated in section 7.2 *Structure of OEM Parameter Block*, the main areas available for OEM adjustment are Tx Equalization to compensate for chip-to-connector trace lengths and spread spectrum clocking parameters.

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 7.4 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.



2.10.2 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.

2.10.3 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 7.4 Recommended Location in SPI Flash.

2.11 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 7.4 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.

2.11.1 Legacy OROM

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

2.11.1.1 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

2.11.1.2 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_om_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 SCU OROM 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_om_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)



2.11.2 UEFI

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

```
//
// Define SCU Parameters protocol GUID
//
// EDK and EDKII have different GUID formats
//
#ifdef(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];
};
```

Appendix (Hardware Compatibility List). There will also be support for drives that are larger than 2 Terabytes as well as drives that support 4KB physical (512B logical) sectors.

2.11.2.1 Safe Mode Support

Intel® RSTe will provide support for booting into Safe Mode for all supported OSs.

2.11.2.2 Non-Intel Controller Support

Intel® RSTe will not hinder, break or prevent operation of non-Intel® Controllers (SATA, PATA, SATA or RAID).

2.11.2.3 Device Configuration

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



2.11.2.4 Power Management

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

2.11.2.5 Staggered Spin-up

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

2.11.2.6 Exporting SATA Drives on SATA Controller

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

2.11.2.7 ATAPI

Intel® RSTe will provide support for ATAPI devices. Intel® RSTe RAID Legacy Option ROM will only support HDD devices (not ATAPI).

2.11.2.8 Solid State Drives (SSD)

Intel® RSTe will support SSDs as if they are Hard Disk Drives.

2.11.2.9 Hybrid Drives

Intel® RSTe will support Hybrid Drives as if they are Hard Disk Drives.

2.11.2.10 TRIM Command : Windows*7¹

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.

Win7 and ATA8 introduce OS support for the TRIM command. Support for the TRIM command allows Windows*7 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.

2.11.2.10.1 SATA Controllers

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

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

¹ Legacy feature; introduced in the Intel® RST 9.6 Release.



2.11.2.10.2 SCU Controller

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

2.11.2.11 Email Alerting and Notification¹

Intel® RSTe will support email notification of certain storage events (see Appendix B for supported events). The Intel® RSTe UI will provide the interface for enabling/disabling and configuring the email notification feature. **The default setting in the UI is 'disabled'.**

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

2.11.2.11.1 Configuration

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

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

2.11.2.11.2 Email Message Format

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

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.

Figure 1 Email Format shows detail of how an email alert is formatted.

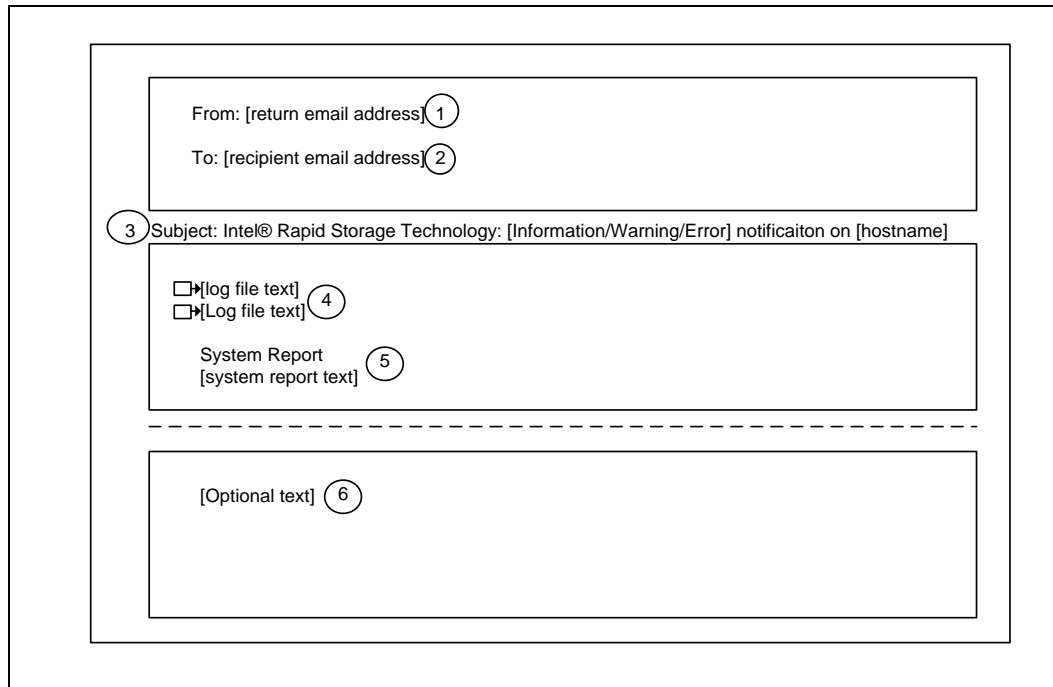


Figure 1 Email Format

2.11.2.11.3 Protocol Support

Email alert shall support SMTP host & SMTP port.

2.11.2.11.4 Error Conditions

See Appendix B for list of supported storage events and their notification mechanism:

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



2.11.3 Utilities

2.11.3.1 Install/Uninstall Utility

Intel® RSTe 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.



3 *Features, Specifications, and Specification Updates*

3.1 **New Features and Supported Specifications Introduced in Intel® RSTe 3.8**

3.1.1 **Beta Support for Grantley Platforms with Intel® C610 series chipset sSATA Controller**

With the release of 3.8.0.1107, introduces Beta level support for the Intel® Grantley platform with the Intel® C610 series chipset. These platforms contain two SATA controllers that can be configured for either AHCI mode or RAID mode. The first SATA controller is a six- 6 Gigabits per second port controller. The second (sSATA Controller) will provide an additional four - 6 Gigabits per second ports. These two controllers are treated as two separate and independent controllers. As such, the RSTe driver included in the release does not support spanning RAID volumes across the two controllers.

When running on these platforms, the customer will see two instances of the RSTe driver installed (one for each controller). To support these two controllers, the platform BIOS will need to contain two RAID Legacy OROM and two RAID EFI drivers (one image for each of the SATA controllers).

3.1.2 **Windows* 8.1 and Server 2012 R2 Support**

The release of 3.8.0.1107 introduces support for Windows* 8.1 and Windows* Server 2012 R2 operating systems. These new operating systems will contain an "inbox" driver that will support the SATA Controllers for the Intel® C600 and C610 series chipset Platform Controller Hub (PCH) when configured for RAID mode. It is strongly recommended that the RSTe 3.8.0.1107 F6 drivers be used instead of the available "inbox" driver. The provided "inbox" driver is intended only for those customers who may not have the RSTe 3.8.0.1107 F6 drivers readily available and ONLY for installing to a single drive (NOT to a RAID volume). Once the OS is installed, it is strongly recommended that the RSTe 3.8.0.1107 package be installed immediately. At that point, it will be safe to migrate the system disk into a RAID Volume (using the RSTe GUI).

When upgrading the existing operating system to Windows* 8.1 or Windows* Server 2012 R2, it is strongly recommended that the RSTe driver be updated to the RSTe 3.8.0.1107 package prior to the OS upgrade.

3.1.3 **RSTe INF File Split**

The release of 3.8.0.1107 introduces support for those customers who need to dynamically install either the 32-bit or 64-bit versions of the RSTe driver in a manufacturing environment. Now a single manufacturing environment utility can be used to install either the 32-bit or the 64-bit version of the RSTe 3.8.0.1107 driver.



3.1.4 Device Information Display in UEFI

The release of 3.8.0.1107 introduces the support of the UEFI driver reporting the physical port a device is connected to. In previous releases of the RSTe UEFI driver, the device information provided (i.e. to the UEFI shell environment) was based off of the enumeration values created during the discovery of the devices attached. To support a manufacturing environment that relies on this information to identify the physical port the device is connected too, the UEFI driver now reports out the physical port value instead of the enumerated value. The data is displayed as a set of 3 values, in the following order:

X- 0: passthru disk, 1: volume

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

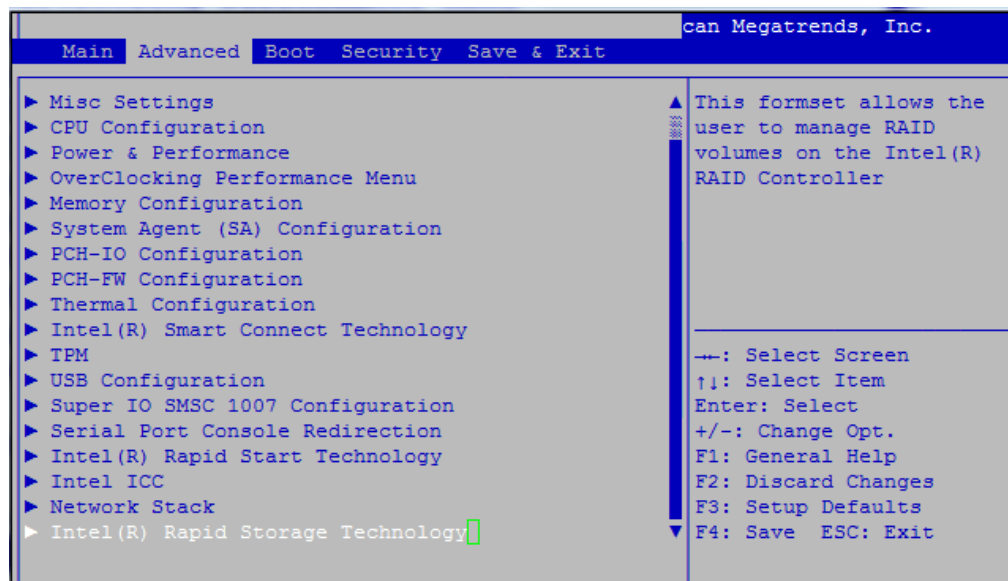
Z - Disk Number on PHY (in case of an expander)

3.2 New Features and Supported Specifications Introduced in Intel® RSTe 3.7

3.2.1 Human Interface Infrastructure (HII) support

With the release 3.7.0.1093, customers will be able to configure RAID volumes using Intel® RSTe UEFI Human Interface Infrastructure (HII). HII is a simple “Walk Up and Use” interface that allows users to create, delete, and manage RSTe RAID Volumes within the BIOS Setup Menu.

The UEFI driver will provide an entry point which can be called to invoke the HII interface. The entry point can be called from any place in the BIOS UI that an IBV chooses.





BIOS must support the 2.3.1 UEFI Specification. In addition, the following protocols must be supported for the RSTe UEFI HII to be used in the BIOS Setup Screen:

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

3.2.2 Expander Exposed to Device Manager

With the release of RSTe 3.7 driver, the RSTe driver now supports exposing expanders that are connected to the Storage Controller Unit (SCU). So, when the Device Manager is opened, the expander will be seen.

3.2.3 Trim on RAID0, RAID1 and RAID 10

With the release of RSTe 3.7 driver, the RSTe driver now supports TRIM on RAID1 and RAID10 volumes. This feature is automatic and nothing is required to enable this functionality.

3.2.4 RSTe UEFI Command Line Tool to support Supports Testing SGPIO

With the release of RSTe 3.7 driver, a new RSTe UEFI CLI utility (LedTool.efi) is being introduced to support the issuing of SGPIO commands to an SGPIO enabled backplane. The purpose is to ensure that the platform SGPIO functionality is working properly in a factory environment. As before, the platform has to be booted into a UEFI Shell and LedTool.efi needs to be copied to a USB key in order to utilize this feature. Please reference the help option available in the utility for detailed instructions.

3.2.5 RSTe GUI Supports Showing SAS Link Width

With the release of RSTe 3.7 driver, the RSTe GUI now provides support for showing the SAS link width. This information can be obtained by viewing the System Report.

3.2.6 SCU Controller Legacy Option ROM Splash Screen Timing Adjustment

With the release of RSTe 3.7 driver, the RSTe SCU Legacy OROM splash screen display duration is now adjustable. To utilize this feature, a change must be made to the OEM Parameter field in SPI Flash. Please reference the section 7.3 BCFS Bit Settings for detailed information.

NOTE: If no changes are made to the OEM Parameter field, there will be no change in the behavior of the SCU Legacy OROM splash screen behavior.

3.3 Updated Features/Specifications Implemented in Intel® RSTe 3.8

Please refer to the product release notes.



4 *Product Certifications*

4.1 **WHQL**

The Intel® RSTe driver is required to be logo certified on the Microsoft Windows* WLK 1.5 test suite for storage device drivers.

The Intel® RSTe 3.7 drivers is 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.

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



5 Appendix A

Table 5-1. 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 5-2. 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 Chipset Product Line Firmware Startup Guide
454672	PatsburgChipset 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 <i>Contact you Intel FAE to get access to this document through Anacapa</i>
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
RST Legacy Requirements Documents	
443658	<i>Intel® Rapid Storage Technology (Intel® RST) 9.6 Product Requirements Document (PRD) Differences Document:</i> http://www.intel.com/cd/edesign/library/asm-na/eng/443658.htm
409399	<i>Intel® Rapid Storage Technology 9.5 Product Requirements Document (PRD) Differences Document :</i> http://www.intel.com/cd/edesign/library/asm-na/eng/409399.htm
375707	<i>Intel® Matrix Storage Manager 8.5 Product Requirements Document (PRD) Differences Document :</i>
373388	<i>Intel® Matrix Storage Manager 8.0 Product Requirements Document (PRD) :</i> http://www.intel.com/cd/edesign/library/asm-na/eng/373388.htm
Technical Guides / White papers	
N/A	<i>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.</i>
445153	<i>Reference Manual for Safe Removal with Link Power Management (LPM) on Hot Plug Capable Port (HPCP) :</i> http://www.intel.com/cd/edesign/library/asm-na/eng/445153.htm



6 Appendix B

Table 6-1. Storage System Events Detected by Monitor Service (IAStorDataMgrSvc)

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



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

NOTES:

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

7 Appendix C

7.1 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.

7.2 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 2 shows this structure.




Figure 2: 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 7.3 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[SCL_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;

7.3 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	
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	



BCFS bit number	Bit meaning	Values		Additional info
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	
		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	



BCFS bit number	Bit meaning	Values	Additional info
15	RESERVED		

7.4 Recommended Location in SPI Flash

Figure 3 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 7.5 *OEM Parameter and SPI Flash Tools*.

Note: The recommended SPI flash format is descriptor mode.

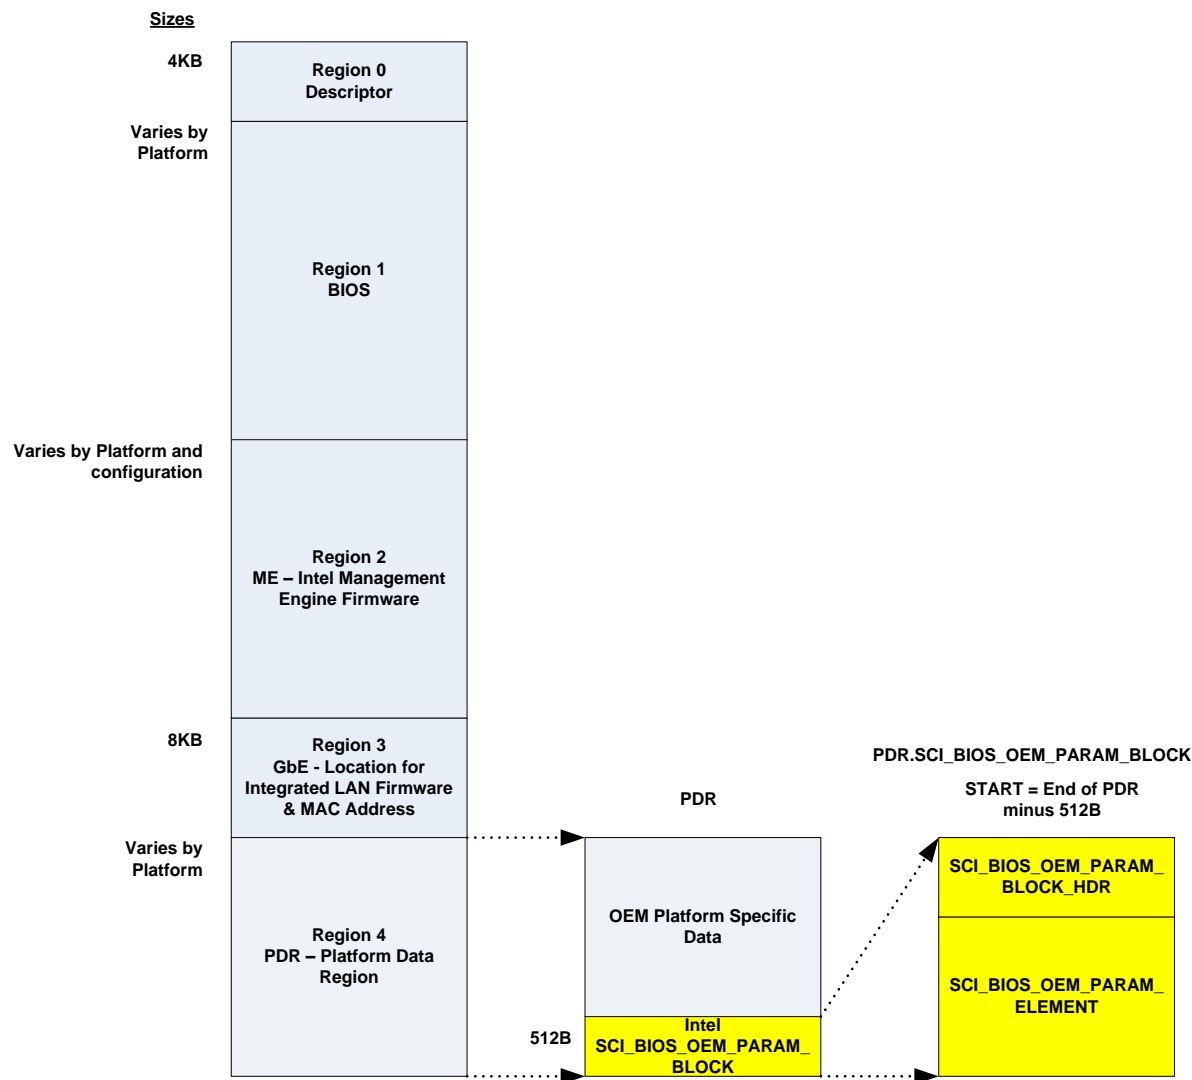


Figure 3 Recommended OEM Param Location in SPI Flash

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 7.6 *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.



7.5 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 7.2 *Structure of OEM Parameter Block* for a detailed description of the PHY parameters.

7.5.1 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. Each individual PHY can have its own unique settings that best fit its situation. As indicated in section 7.2 *Structure of OEM Parameter Block*, the main areas available for OEM adjustment are Tx Equalization to compensate for chip-to-connector trace lengths and spread spectrum clocking parameters.

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 7.4 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.



7.5.2 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.

7.5.3 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 7.4 Recommended Location in SPI Flash.

7.6 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 7.4 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.

7.6.1 Legacy OROM

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

7.6.1.1 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

7.6.1.2 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_om_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 SCU OROM 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_om_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)



7.6.2 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];  
};
```




8 Appendix D

8.1 Hardware Compatibility

8.1.1 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	HUC101473CSS300,



	2.5" (Cobra B)	
--	----------------	--

Enterprise SATA Drives

Vendor	Family	Model Name/Number
Fujitsu	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