

Intel® Rapid Storage Technology

OEM Technical Guide

For the Intel® Rapid Storage Technology 11.2 Release Version

Revision 0.8

Intel Confidential



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

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

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

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

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

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

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

Copyright © 2012, Intel Corporation. All rights reserved.



Contents

1 About This Document 9
1.1 Purpose and Scope of this Document 9
2 Intel® Rapid Storage Technology 10
2.1 Overview of RAID Levels 10
2.2 Typical Usage Model for RAID Levels 12
3 Intel® Rapid Storage Technology Suite 13
3.1 Intel® Rapid Storage Technology Software 13
3.2 Intel® Rapid Storage Technology Option ROM 14
3.3 Intel RAID Configuration Utility 14
3.3.1 RCfgSata Utility for MS-DOS* 14
3.4 RAIDCfg32/64 Utilities 15
4 Creating a RAID Volume 18
4.1 Using the Intel Rapid Storage Technology UI 18
4.2 Using the Intel Rapid Storage Technology Option ROM User Interface 19
4.3 Using the RAID Configuration Utility 19
5 Deleting a RAID Volume 21
5.1 Using the Windows User Interface Utility 21
5.2 Using the Option ROM User Interface 21
5.3 Using the RAID Configuration Utility 22
6 Common RAID Setup Procedures 23
6.1 Build a SATA RAID 0, 1, 5 or 10 System 23
6.2 Build a SATA "RAID Ready" System 24
6.3 Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System 25
6.4 Migrate an Existing Data Hard Drive to a RAID 0 or RAID 1 Volume 26
6.5 Migrating From one RAID Level to Another 26
6.6 Create a RAID Volume on SATA While Booting to PATA 27
6.7 Build a RAID 0 or RAID 1 System in an Automated Factory Environment 28
6.7.1 Part 1: Create the Master Image 28
6.7.2 Part 2: Apply the Master Image 28
7 RAID Volume Data Verification and Repair Feature 30
7.1 Verify and Repair Volume Feature 30
8 Intel® Rapid Recover Technology 31
8.1 Creating a Recovery Volume Through the RAID Option ROM 31
8.2 Creating a Recovery Volume Through the Intel® RST UI 32
8.3 Changing Recovery Volume Modes 33
8.4 Update Recovery Volume in On Request Update Policy 33
8.5 Access Recovery Drive Files 33



8.6 Hide Recovery Drive Files 34
8.7 Scenarios of Recovering Data 34

9 System Running off Recovery Drive 37
9.1 Drive Offline or Missing 37

10 Pre-OS Installation of the Intel® Rapid Storage Technology Driver 38
10.1 Pre-OS Installation Using the "F6" Method 38
10.2 Unattended Installations of Microsoft Windows XP 39
10.3 Creating a Floppy Disk Containing the Intel® RST Driver 40
10.3.1 Method 1 – Using Install Applications 40
10.3.2 Method 2: Using F6 Application 41

11 Determining the Version of the RAID Driver 42
11.1 Using Intel® Rapid Storage Technology User Interface (UI) 42
11.2 Using Intel® RST File Properties (Alternate) 42
11.3 Determining the Version of the Option ROM 43
11.3.1 Using the Intel Rapid Storage Technology UI 43
11.3.2 Using the Intel® RST Option ROM User Interface 43

12 Un-installation 44
12.1 Uninstalling the Intel® RST Software (except the RAID Driver) 44
12.2 Disabling the RAID Driver by Disabling the RAID Controller 44

13 Registry Customizations 46
13.1 Native Command Queuing Settings 47
13.2 Zero Power ODD Settings 47
13.3 E-mail Notification UI Visible Enable/Disable 48
13.4 Disabling the Intel® SRT Option to Accelerate to Maximized Mode 48

14 Link Power Management with Intel® Rapid Storage Technology 50
14.1 Instructions to disable/enable LPM 50
14.2 Safe Removal with LPM on hot plug capable ports 51

15 RAID Option ROM and Utilities 52

16 HDD Password Support With RAID Volumes 53
16.1 HDD Password Use Cases 53
16.2 Unlocking Password Protected Disks in RST UI 54

17 Intel® Smart Response Technology 55
17.1 Overview 55
17.1.1 Warnings! 55
17.1.2 Requirements and Limitations 56
17.1.3 Acceleration Modes 58
17.2 Build a New System with Disk/Volume Acceleration Enabled 58
17.2.1 Prepare New Computer 59
17.2.2 Installing the OS to a New System Prepared for Disk/Volume Acceleration 68
17.3 OEM System Manufacturing and Intel® Smart Response Technology 69
17.3.1 Imaging an OS onto a Pre-Configured Acceleration-enabled HDD 69
17.3.2 Enabling Acceleration post end user OOBE 71

17.4 OEM System Manufacturing and Cache Pre-load for Intel® Smart Response Technology 73
17.4.1 Requirements 73
17.4.2 Process 73
17.4.3 Replication for Mass Production of Accelerated HDD and SSD 75
17.5 HDD ATA Power-Up in Standby (PUIS) Supporting Intel® Smart Connect Technology‡ 75
17.5.1 Overview 76
17.5.2 Theory of Operation 77

18 Intel® Rapid Storage Technology UI 80

18.1 Introduction 80
18.1.1 Getting Started 81
18.1.2 System Requirements 83
18.1.3 Understanding the Application 84
18.1.4 Notification Area 86
18.2 Storage System Status 87
18.2.1 Understanding the Status 87
18.2.2 Storage System View 88
18.3 Creating a Volume 90
18.3.1 Volume Requirements 90
18.3.2 Creation Process 92
18.3.3 Creating Additional Volumes 96
18.4 Managing the Storage System 97
18.4.1 Managing Arrays 98
18.4.2 Managing Volumes 102
18.4.3 Managing Disks 116
18.4.4 Managing Ports 122
18.4.5 Managing ATAPI Devices 122
18.5 Accelerating the Storage System 123
18.5.1 Cache Device Properties 123
18.5.2 Enabling Acceleration 126
18.5.3 Disabling Acceleration 127
18.5.4 Changing Acceleration Mode 127
18.5.5 Accelerating a Disk or Volume 128
18.5.6 Resetting a Cache Device to Available 129
18.5.7 Disassociating the Cache Memory 130
18.6 Preferences 130

19 Using the BCFS to Differentiate Platform SKUs 132

19.1 Configuring the Platform's RAID Related Features 132
19.1.1 Configuring the Standard Supported RAID Levels 132
19.1.2 Configuring Intel® RRT Related RAID Features 133
19.1.3 Configuring the Behavior of the OROM UI and Banner 134
19.1.4 Configuring Intel® RST UI Capabilities 135
19.1.5 Configuring the Platform to Support Intel® Smart Response Technology 135
19.1.6 BIOS Control Feature Set 136

20 Testing, Certification Notes 138

20.1 Correcting Microsoft* Windows 7 (Win7) WHQL test failure 138

21 Glossary 139



22 Troubleshooting 146
22.1 Failed Volumes 146
22.2 Degraded Volumes 148
22.3 Other Volume States 151
22.4 Disk Events 154
22.5 Caching Issues 155
22.6 Software Errors 158
23 Appendix A: Common Storage Management Interface Support (CSMI) 160



Revision History

Revision Number	Description	Intel® Rapid Storage Release Version	Revision Date
Intel® RST 11.2 Release			
0.82	• Updated section17.4 • Added Appendix A: CSMI Support	11.2	Q2 '2012
Intel® RST 11.1 Release			
0.82	• Corrected some general typos • Corrected installation procedures in chapter 6 • Added section 17.4: Cache Preload	11.1	Q1 '2012
Intel® RST 11.0 Release			
1.0	• Edited section 10.1; removed support for CD/DVD driver installation during OS installation • Edited section 17.2.1.4 to include SRT and FFS configuration info • Added section 17.4 to cover the HDD ATA PUIS support for Intel® Smart Connect Technology	11.0	Q4 '2011
0.9	• Minor edits to section 17.3	11.0	Q4 '2011
0.82	• Added section 17.3 ***OEM System Manufacturing and Intel® Smart Response Technology***	11.0	Q4 '2011
0.81	• Added section 3.4.1.5: included registry setting information to disable/enable UI functionality	11.0	Q3 '2011
0.8	• Changed RAIDCfg to RCfgSata • Changed RAIDCply to RCmpSata • Updated sections: 3.3.1, 15, 17.1.2, 17.2.1.3.3 step 4, • Added sections: 17.2.1.3.3 step 4a, 17.3, 17.4	11.0	Q3 '2011
0.5	• Added sections: 17.1.1, 17.2.3 • Updated sections:	11.0	Q2 '2011
Intel® RST 10.5 Release			
1.01	• Added section 19: Using the BCFS to Differentiate Platform SKUs	10.5.0.1027	Q2 '2011



1.0	• Added section 16.2: HDD Password Unlock via RST UI is configurable • Section 18.4.3.2 updated	10.5.0.1027	Q2 '2011
0.8	Initial Draft Version	10.5.0.1027	Q1 '2011



1 About This Document

1.1 Purpose and Scope of this Document

This document will assist customers in evaluating, testing, configuring, and enabling RAID and AHCI functionality on platforms using the *Intel® Rapid Storage Technology* software for the chipset components as listed in the product's Readme.txt file.

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



2 Intel® Rapid Storage Technology

Intel® Rapid Storage Technology (Intel® RST) provides added performance and reliability for systems equipped with serial ATA (SATA) hard drives and solid state disk (SSD) drives to enable an optimal PC storage solution. It offers value-add features such as RAID and advanced Serial ATA* capabilities for the Microsoft* client and server operating systems (for detailed OS support, review the Release Notes for each software release). The driver also offers Non-volatile (NV) caching for performance and application acceleration with optional SSD used as the cache memory device.

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

A configuration supporting two RAID levels can also be achieved by having two volumes in a single RAID array that use Intel® RST. These are called matrix arrays. Typical for desktops, workstations, and entry level servers, Intel® RST RAID solution addresses the demand for high-performance or data-redundant platforms. OEMs are also finding it beneficial to implement this RAID capability into mobile platforms as well.

2.1 Overview of RAID Levels

RAID 0 (striping)

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

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

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

**Note: the number of drives supported in a RAID 0 array is dependent upon the chipset model and desktop vs. mobile. Please consult the specification for your chipset to determine the maximum number of drives supported in a RAID array.



RAID 1 (mirroring)

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

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

RAID 5 (striping with parity)

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

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

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

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

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

RAID 10 (striping and mirroring)

RAID level 10 uses four hard drives to create a combination of RAID levels 0 and 1. The data is striped across a two-disk array forming a RAID 0 component. Each of the drives in the RAID 0 array is mirrored to form a RAID 1 component. This provides the performance benefits of RAID 0 and the redundancy of RAID 1.

The RAID 10 volume appears as a single physical hard drive with a capacity equal to two drives of the four drive configuration (the minimum RAID 10 configuration). The space on the remaining two drives will be used for mirroring.



2.2 Typical Usage Model for RAID Levels

RAID 0	This provides end-users the performance necessary for any disk-intensive applications; these include video production and editing, image editing, and gaming applications.
RAID 1	This provides end-users with data redundancy by mirroring data between the hard drives.
RAID 5	This provides end-users with good performance and data redundancy by striping data and parity across all the hard drives. The write performance is enhanced with volume write-back cache.
RAID 10	This provides end-users with the benefits of RAID 0 (performance) and RAID 1 (data mirroring).



3 Intel® Rapid Storage Technology Suite

The Intel® Rapid Storage Technology Suite contains three core component groups:

1. Intel® Rapid Storage Technology OS runtime software package
    a. AHCI/RAID driver
    b. Graphical User Interface (Intel® RST UI), optional
    c. Event Monitor service (IAStorDataMgrSvc), optional; communicates with:
        i. Intel® RST UI (Windows graphical user interface)
        ii. Event Notification Tray Icon (IAStorIcon.exe)
        iii. Windows system NT Event log
2. Intel® Rapid Storage Technology RAID Option ROM (OROM)
3. Intel® Rapid Storage Technology Utilities (***ONLY for OEM manufacturing use! Must not be distributed to end-user systems!***)
    a. RCfgSata DOS utility (formerly RAIDCfg: not an end-user tool)
    b. RCmpSata DOS utility (formerly RAIDCply: not an end-user tool)
    c. RAIDCfg32/64 Windows utilities (not an end-user tool)

3.1 Intel® Rapid Storage Technology Software

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



3.2 Intel® Rapid Storage Technology Option ROM

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

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

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

3.3 Intel RAID Configuration Utility

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

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

3.3.1 RCfgSata Utility for MS-DOS*

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

```
=============================================================================

RCfgSata.exe [/?] [/Y] [/Q] [/C:vol_name] [/SS:strip_size]
     [/L:raid_level] [/S:vol_size] [/DS:disk_ports] [/D:vol_name] [/X]
     [/I] [/P] [/U] [/ST] [/V]

/?   Displays Help Screen.  Other options ignored.
/Y   Suppress any user input. Used with options /C, /D & /X.
/Q   Quiet mode / No output. Should not be used with status commands.

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

```
/C   Create a volume with the specified name. /S, /DS, /SS, & /L can be
     specified along with /C.
/SS  Specify strip size in KB.  Only valid with /C.
/L   Specify RAID Level (0, 1, 10, or 5). Only valid with /C.
/S   Specify volume size in GB or percentage if a '%' is appended.
     Percentage must be between 1-100.  Only valid with /C.
/DS  Selects the disks to be used in the creation of volume. List should
     be delimited by spaces.
/D   Delete Volume with specified name.
/X   Remove all metadata from all disks. Use with /DS to delete metadata
     from selected disks.
/I   Display All Drive/Volume/Array Information. /P can be specified.
/P   Pause display between sections. Only valid with /I or /ST.
/U   Do not delete the partition table. Only valid with /C on RAID 1
     volumes.
/ST  Display Volume/RAID/Disk Status. /P can be specified.
/V   Display version information
=============================================================================
```

3.4 RAIDCfg32/64 Utilities

The Intel RAIDCfg32/64 utility is an executable. It provides OEMs with the ability to create, delete, and manage RAID volumes on a system within a windows environment using command line parameters that make it possible to perform these functions by using scripts or shell commands.

The command syntax for the Intel RAIDCfg32/64 utilities is shown below:

```
=============================================================================

RaidCfg32.exe [/?] [/Y] [/Q] [/RD] [/SPARE:drive] [/ACCEL:vol_name]
              [/ACCELD:port] [/MODE:<Max|Enh|Off>] [/NVCACHE]
              [/C:vol_name] [/S:vol_size] [/SS:strip_size]
              [/L:raid_level] [/DS:port] [/D:vol_name] [/X] [/I] [/P]
              [/ST] [/STD] [/STV] [/F:vol_name] [/IRRT]
              [/SYNC:<Auto|Manual>] [/M:port]

/?        Displays Help Screen.  Other options ignored.
/Y        Suppress any user input. Used with options /C, /D, /X, /F,
          /ACCEL & /ACCELD.
/Q        Quiet mode / No output. Should not be used with status
          commands.
          COMMANDS - Only one at a time.
/RD       Inquire if the raid driver is running.
/SPARE    Set a drive to be a spare.  List should be delimited by
          spaces.  Use /SPARE:port-target for drives behind a port
          multiplier.
/ACCEL    Accelerate a volume using a pre-configured SSD as cache.
          /MODE must be specified with the /ACCEL command.
/ACCELD   Accelerate a single disk using a pre-configured SSD as cache.
          /MODE must be specified with the /ACCELD command.
/MODE     Set the acceleration mode to 'Maximum', 'Enhanced' or 'Off'.
          Only valid with /ACCEL and /ACCELD.
```



```
/NVCACHE  Create a cache volume(18.6GB or 64GB).  /SS & /L not supported
          with the /NVCACHE option.  Only valid with /C. An example input
          for using /NVCACHE with an SSD on port 2 is as follows:
              raidcfg32 /C CacheVolume /DS 2 /S 18.6 /NVCACHE
/NVLOAD   Loads content into the NV cache.  This switch requires a
          directory path and filename.  A wildcard ("*") can be given for
          the filename.
          This operation will not perform sub-directory recursion.
          Usage examples:
              RaidCfg32 /NVLOAD:"C:\Windows\*"
              RaidCfg32
              /NVLOAD:"C:\ProgramFiles(x86)\Adobe\Reader9.0\Reader\AcroRd
              32.exe"
          *NOTE: To ensure content is loaded, make sure the following
          registry setting is set upon boot:
              HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor
              \Parameters\
              NvCachePolicy = 0, (inserts all into cache).
              When cache loading is complete, delete the NvCachePolicy
              key from the registry and reboot so the RST driver will
              resume standard cache policies.
/NVLOADR  Loads content into the NV cache.  Equivalent to /NVLOAD except
          performs sub-directory recursion.
/NVSTAT   Provides NV cache statistics e.g. full capacity indicator.

/C        Create a volume with the specified name.
          /S, /DS, /SS, & /L can be specified along with /C.
/SS       Specify strip size in KB.  Only valid with /C.
/L        Specify RAID Level (0, 1, 10, or 5).  Only valid with /C.
/S        Specify volume size in GB or percentage if a '%' is appended.
          Percentage must be between 1-100. Only valid with /C.
/DS       Selects the disks to be used in the creation of volume.
          List should be delimited by spaces.  Use /DS:port-target for
          drives behind a port multiplier.
/D        Delete Volume with specified name.
/X        Remove metadata from all disks. Selecting a single disk with
          the /DS option will delete the metadata from all volume
          members.
/I        Display All Drive/Volume/Array Information.  /P can be
          specified.
/P        Pause display between sections. Only valid with /I.
/ST       Display Volume/RAID/Disk Status.
/STD      Display delimited Disk Status
          Port,Model,SerialNumber,FirmwareVersion,Array,Status,Size,Free,
          Type.
/STV      Display delimited Volume Status
          Index,Level,StripSize,Size,Status,Bootable,Array,Name.

          The "Size" value represents the number of logical blocks of the
          disk.  Each logical block is 512 bytes, regardless of the
          Physical Sector Size of the disk (which can be 512 bytes or
          1024 bytes).
          Example of calculation:
              Size = 156301488 logical blocks.
              156301488 logical blocks * (512 bytes / logical block) =
                80026361856 bytes
              80026361856 bytes * (1 KB / 1024 bytes) * (1 MB / 1024 KB)
                * ( 1GB / 1024 MB) = 74.53 GB

              Size = 312581808 logical blocks.
              156301488 logical blocks * (512 bytes / logical block) =
                160041885696 bytes
```



```
                160041885696 bytes * (1 KB / 1024 bytes) * (1 MB / 1024
                  KB) * ( 1GB / 1024 MB) = 149.1 GB

/F          Repair failed RAID0 Volume.
/IRRT       Create a recovery volume. Only valid with /C.
/SYNC       Set sync type for 'Recovery' volume. Only valid with /IRRT.
/M          Specify the port number of the Master disk for 'Recovery'
            volume. Only valid with /IRRT.  Use /M:port-target for drives
            behind a port multiplier.

===========================================================================
```



4 Creating a RAID Volume

RAID volumes can be created three different ways. The method most widely used by end-users is to use the Intel Rapid Storage Technology UI in Windows*. The second method to create a RAID volume is to use the Intel Rapid Storage Technology option ROM user interface. The third way, used by OEMs only, is using the RCfgSata or RAIDCfg32/64 utilities.

4.1 Using the Intel Rapid Storage Technology UI

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

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

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

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

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

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

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

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

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



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

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

4.2 Using the Intel Rapid Storage Technology Option ROM User Interface

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

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

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

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

5. Select the strip size (128 KB is the default for RAID 0) by using the arrow keys, then press Enter when done.

6. Enter the size for the RAID volume in gigabytes. The default value will be the maximum size. If you specify a smaller size, you will be able to create a second volume in the remaining space using the same procedure.

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

4.3 Using the RAID Configuration Utility

Note: RAIDCFG32 and RAIDCFG64 can be used interchangeably below.

Run "RCfgSata.exe" (DOS environment) or "raidcfg32.exe (or raidcfg64.exe)" (Windows environment) with the following command line flags to create a RAID volume.

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

```
C:\>rcfgsata.exe /C:OEMRAID0 /DS:0 1 /SS:128 /L:0 /S:120
C:\>raidcfg32.exe /C:OEMRAID0 /DS:0 1 /SS:128 /L:0 /S:120
```



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

```
C:\>RCfgSata.exe /C:OEMRAID0
```

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

```
C:\>rcfgsata.exe /?
C:\>raidcfg32.exe /h
```

Note:

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



5 Deleting a RAID Volume

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

5.1 Using the Windows User Interface Utility

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

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

2. Under 'Status' or 'Manage' Click on the volume you want to delete. The user will be presented with the volume properties on the left.

3. Click on 'Delete volume'

4. Review the warning message, and click 'Yes' to delete the volume.

5. The 'Status' page refreshes and displays the resulting available space in the storage system view. You can now use it to create a new volume.

5.2 Using the Option ROM User Interface

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

2. Within this UI, select option #2 'Delete RAID volume'.

3. You should be presented with another screen listing the existing RAID volume.

4. Select the RAID volume you wish to delete using the up and down arrow keys.

5. Press the Delete key to delete the RAID volume

6. Press Y to confirm.

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



5.3 Using the RAID Configuration Utility

Run “RCfgSata.exe” (DOS environment) or “raidcfg32.exe” (Windows environment) with the following command line flag to delete a RAID volume. The following command line will instruct the utility to delete a RAID 0 volume named “OEMRAID0”

```
C:\>rcfgsata.exe /D:OEMRAID0
C:\>raidcfg32.exe /D:OEMRAID0
```

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

```
C:\>RCfgSata.exe /?

C:\>raidcfg32.exe /h
```



6 Common RAID Setup Procedures

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

This is the most common setup. This configuration will have the operating system striped for RAID 0, or mirrored for RAID 1, or striped with parity for RAID 5, or mirrored and striped across two or up to four SATA hard drives for RAID 10. To prepare for this, you must have the Intel RAID driver on a floppy disk. See the procedure for creating this floppy further down in this document.

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

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

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

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

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

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

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

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

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

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



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

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

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

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

7. Installation procedures as follows:

   a. Windows XP and previous OS installs: At the beginning of Windows Setup, press F6, when prompted select the command to install a third-party SCSI* or RAID driver. Insert the floppy disk with Intel® RST driver (floppy drive required). The floppy will list a set of controllers; select the proper controller for your system from the list. Continue the install.

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

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

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

6.2 Build a SATA "RAID Ready" System

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

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

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

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

4. Installation procedures as follows:



a. Windows XP and previous OS installs: At the beginning of Windows Setup, press F6, when prompted select the command to install a third-party SCSI* or RAID driver. Insert the floppy disk with Intel® RST driver (floppy drive required). The floppy will list a set of controllers; select the proper controller for your system from the list. Continue the install.

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

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

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

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

If you have an existing "RAID Ready" system as defined in section 6.2: Build a SATA "RAID Ready" System, then you can use the following steps to migrate from a single-drive non-RAID configuration to a two drive RAID 0 or RAID 1 configuration. The resulting configuration will be identical to that created by the procedure in section 6.1: Build a SATA AID0, 1, 5 or 10 System. To prepare for this, you will need to install another SATA hard drive with a capacity equal to or greater than the capacity of the hard drive being used as the source hard drive.

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

2. Install the second SATA hard drive on the available SATA port.

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

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

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

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

7. After the migration is complete, reboot the system. If you migrated to a RAID 0 volume, use Disk Management from within Windows in order to partition and format the empty space created when the two hard drive capacities are combined.



You may also use third-party software to extend any existing partitions within the RAID volume.

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

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

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

1. Note the serial number of the SATA hard drive that is already installed. You will use this to select it as the source hard drive when initiating the migration.
2. Physically attach the second SATA hard drive to the available SATA port.
3. Boot to Windows, install the Rapid Storage Technology software, if not already installed, using the setup package obtained from a CD-ROM or from the Internet. This will install the necessary Intel Rapid Storage Technology UI and start menu links.
4. Open the Intel Rapid Storage Technology UI from the Start Menu.
5. Follow steps 4 to 7 in section 6.3

6.5 Migrating From one RAID Level to Another

RAID level migration allows an existing RAID configuration to be migrated to another RAID configuration. The following migrations are possible.
NOTE: Not all migrations are supported on all chipsets. The support varies depending on the chipset and the ports supported on the chipset (For supported migrations for each chipset please Intel Rapid Storage Technology product requirements document):

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



	3, 4, 5 or 6-disk RAID 5
2-disk RAID 0	3, 4, 5 or 6-disk RAID 5
3-disk RAID 0	4, 5 or 6-disk RAID 5
4-disk RAID 0	5 or 6-disk RAID 5
4-disk RAID 10	4, 5 or 6-disk RAID 5

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

Please follow the procedure illustrated below

1. Start the Intel Rapid Storage Technology UI application:

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

2. Under 'Status' or 'Manage', in the storage system view, click the array or volume to which you want to modify. The volume properties now display on the left.

3. Click 'Change type'.

4. In the 'Change Volume Type' dialog, type a new name if you want to change the default name.

5. Select the new volume type, and then click 'OK'.

6. The 'Manage' page refreshes and reports the new volume type.

7. After the migration starts, you can view the migration progress under status.

8. When the Status field indicates volume as ‘Normal’, the migration is complete.

6.6 Create a RAID Volume on SATA While Booting to PATA

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

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

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

3. Boot to Windows; install the Intel Rapid Storage Technology software, if not already installed, use the setup package obtained from a CD-ROM or from the



Internet. This will install the necessary Intel Rapid Storage Technology UI and Start menu links.

4. Use the Intel Rapid Storage Technology UI to create a RAID 0 volume on two SATA drives according to the procedure in section 6.1 of this document.

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

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

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

6.7.1 Part 1: Create the Master Image

1. Build a RAID 0 or RAID 1 System as described in section 6.1 of this document.

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

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

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

6.7.2 Part 2: Apply the Master Image

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

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

3. Within a DOS environment, use the Intel RAID Configuration utility (RCfgSata.exe) to create a RAID volume. The following command line will instruct the utility to create a RAID 0 volume named "OEMRAID0" on the hard drives on Port 0 and 1 with a strip size of 128 KB and a size of 120GB:



C:\>RCfgSata.exe /C:OEMRAID0 /DS:0 1 /SS:128 /L:0 /S:120.
The following command line will display all supported command line parameters and their usage: *C:\>RCfgSata.exe /?*

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

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



7 RAID Volume Data Verification and Repair Feature

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

7.1 Verify and Repair Volume Feature

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

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

Follow the below steps to start RAID volume data verification

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

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

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

4. The verification progress is shown under 'status'

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



8 Intel® Rapid Recover Technology

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

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

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

- Continuous Update Policy
- On Request Update Policy

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

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

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

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

8.1 Creating a Recovery Volume Through the RAID Option ROM

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



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

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

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

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

OROM Recovery menu provides the following options

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

8.2 Creating a Recovery Volume Through the Intel® RST UI

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

Follow the below steps to create a Recovery Volume

1. Under Create select the volume type as 'Recovery' and click 'Next'
2. Under the 'Configure Volume' you can change the default volume name if you want, then select the 'master' disk and then the 'recovery' disk. Now change the 'update' mode if needed to 'On Request'. The default selection is 'continuous'.
3. Once all the above selections are made, click 'Next'
4. Under 'Confirm' review the selected configuration. If you are not ok with the configuration click 'back' or click 'create volume' if you are fine with the configuration.
5. Now you will see a dialog box with warning message and read the warning message before clicking 'ok' to make sure you are erasing data on the right disk.
6. Once you click 'ok' the volume creation starts and progress of the volume creation can be viewed under status. Once the status is set to 'normal' the volume creation is completed.
7. The system will synchronize the Primary with the Recovery disk once after the creation of the Recovery volume.



8.3 Changing Recovery Volume Modes

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

1. Open Intel® Rapid Storage Technology UI.
2. Under 'Manage' or 'Status' click on the recovery volume under the storage system view on right where you need to change the update mode. The volume properties now display on the left view
3. Click 'change mode' and then click 'yes' to confirm.
4. The page refreshes and the volume properties report the new update mode. NOTE: Disabling the continuous update policy requires the end-user to request updates manually. Only changes since the last update process are copied. The recovery volume will remain in On Request Policy until the end-user enables continuous updates.

8.4 Update Recovery Volume in On Request Update Policy

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

1. Open Intel® Rapid Storage Technology UI.
2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.
3. Click 'Update data'.
4. A dialog box is shown stating that the only changes since the last update will be copied. Select the check box if you don't want this confirmation message to display each time you request an update. Click 'Yes' to confirm.
5. The progress of update process can be viewed under 'status' or 'manage'.

8.5 Access Recovery Drive Files

When data recovery to the master disk of a recovery volume is required, you can use 'access the recovery disk files' option. This action is only available if a recovery volume is present, in a normal state, and in on request update mode. Follow the



below instructions to access the recovery drive file when you have a recovery volume in 'on request' mode on your system (If the recovery drive is not in continuous mode, use the instructions in section 8.3 to change the mode)

1. Open the 'Intel Rapid Storage Technology UI'.
2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.
3. Click on 'Access recovery disk files'.
4. Now you can view recovery disk files using Windows Explorer*.

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

8.6 Hide Recovery Drive Files

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

1. Open the 'Intel Rapid Storage Technology UI'.
2. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties now display on the left.
3. Click 'Hide recovery disk files'.
4. Now the recovery driver files are no longer accessible in Windows Explorer.
5. The page refreshes and data updates on the volume are now available.

8.7 Scenarios of Recovering Data

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

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

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



Scenario 2:
What happens if the Master Drive fails and/or the user would like to do a reverse synchronization to a new Master Drive?

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

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

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

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

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

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

Solution:

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

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



Note: If the system is running is on battery, the volume will not synchronize if it is in continuous update policy. If the volume is in Update on Request policy, then the synchronization will be successful.

Additional comments: need to call out that an on update volume should first be updated before the recovery disk is valid.

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

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

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

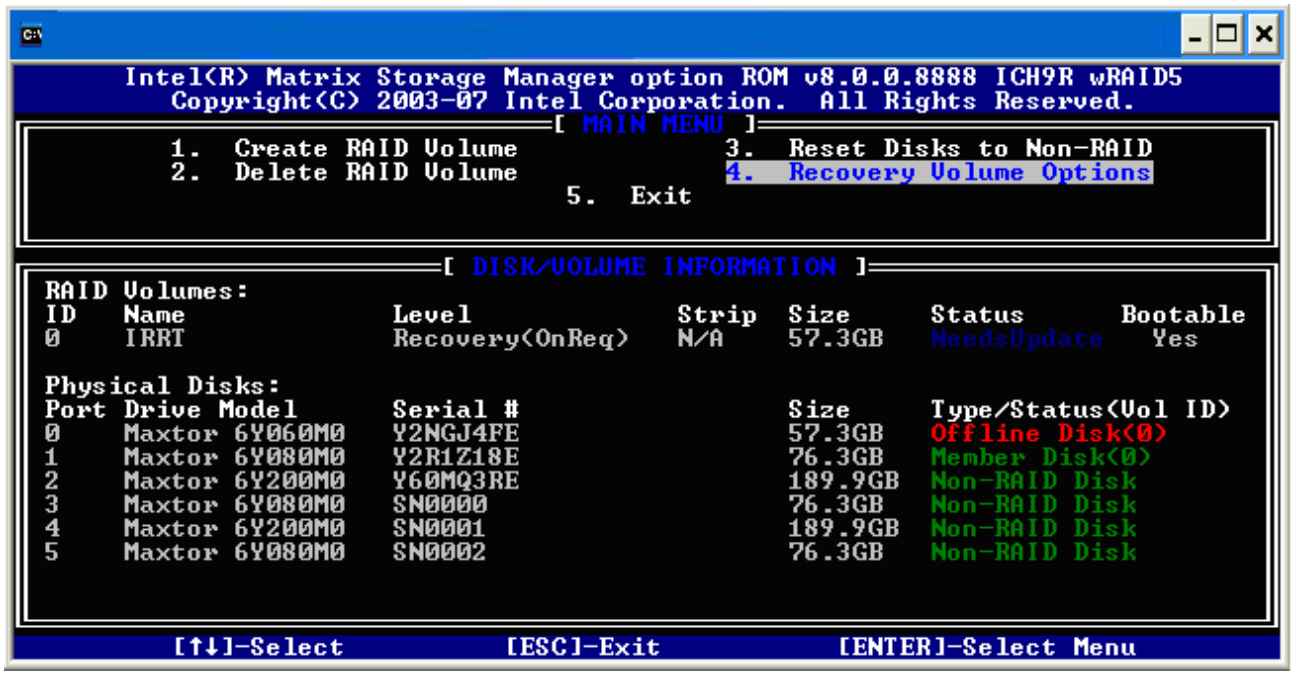
9 System Running off Recovery Drive

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

9.1 Drive Offline or Missing

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

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



Then Select Option 2 Enable Only Recovery Disk.



10 Pre-OS Installation of the Intel® Rapid Storage Technology Driver

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

10.1 Pre-OS Installation Using the "F6" Method

1. When you start installing Windows XP and older operating systems, you may encounter a message stating, "Setup could not determine the type of one or more mass storage devices installed in your system". If this is the case, then you are already in the right place and are ready to supply the driver. If this is not the case, then press F6 when prompted at the beginning of Windows setup.

2. Press the 'S' key to select 'Specify Additional Device".

3. You should be prompted to insert a floppy disk containing the manufacturer-supplied driver into the A: drive. This disk should contain the driver which includes:

Note: For Windows Vista or later you can also use a USB Key.

- Driver binary (iastor.sys),
- INF files (iastor.inf and iaAhci.inf),
- cat files (iaStor.cat and iaAhci.cat) and
- TxtSetup OEM file (txtsetup.oem) required for pre-Vista OS versions.

These should be in the root directory of the floppy. See Section 10.3 : Creating a Floppy Disk Containing the Intel Rapid Storage Technology Driver for instructions on making a driver installation disk.

For post-Windows XP OS versions:

1. During the Operating system installation, after selecting the location to install Vista click on 'Load Driver' button to install a third party SCSI or RAID driver.



2. When prompted, insert the floppy disk or media (Floppy or USB key) with the Intel® RST driver files and press Enter.

3. You should be shown a list of available RAID / SCSI Adapters.

4. Select the appropriate Intel RAID controller and press ENTER.

5. The next screen should confirm that you have selected the Intel® RAID controller. Press ENTER again to continue.

6. You have successfully installed the Intel® Rapid Storage Technology driver, and Windows setup should continue.

7. Leave the disk in the floppy drive until the system reboots itself. Windows setup will need to copy the files from the floppy again after the RAID volume is formatted, and Windows setup starts copying files.

10.2 Unattended Installations of Microsoft Windows XP

To install the driver as outlined in the Microsoft document, "Deployment Guide Automating Windows NT* Setup" use the TXTSETUP.OEM file included in this package and insert the lines below into the UNATTEND.TXT file. This method is also available for Microsoft Windows XP operating system. For Windows XP, extract the iaStor.inf, iaAhci.inf, iaStor.sys, iaStor.cat, iaAhci.cat, and Txtsetup.oem files.

For Microsoft Windows XP, insert the following text into the UNATTEND.TXT file:

For system in RAID mode:

```
[MassStorageDrivers]

"Intel® 82801IR/IO  SATA RAID Controller (ICH9R/IO)" = OEM (replace with
actual name for the controller in your system)

[OEMBootFiles]

iaStor.inf

iaStor.sys

iaStor.cat

Txtsetup.oem
```

For systems in AHCI Mode:

```
[MassStorageDrivers]

"Intel® 82801IR/IO  SATA AHCI Controller (ICH9R/IO)" = OEM
```



```
[OEMBootFiles]

iaAhci.inf

iaStor.sys

iaAhci.cat

Txtsetup.oem
```

10.3 Creating a Floppy Disk Containing the Intel® RST Driver

This procedure should be used to create a floppy disk containing the Intel® Rapid Storage Technology driver for use in installing the driver using the "F6 method".

10.3.1 Method 1 – Using Install Applications

1. On a system running Microsoft Windows, download the Intel® Rapid Storage Technology package or obtain it from your Intel representative. The package contains the following files:
   - iata_enu.exe (English only executable)
   - iata_enu.zip (English only zip package)
   - iata_cd.exe (Multi-language executable)
   - iata_cd.zip (Multi-language zip package)
2. For convenience copy the above files into the C:\ drive.
3. Insert a blank DOS formatted floppy disk into the A: drive.
4. Unpack the iata_enu.zip or iata_cd.zip files. The setup.exe will be located in path "..\iATA_ENU\Disk1" for English version.
5. Choose the language in which you want the driver to be supported and execute the commands described below. From Windows command prompt, type the following commands:

   a. Using Setup.exe – English version:
      - `C:\>setup.exe -a -p c:\IAAdriver`

   b. Using iata_enu.exe - English version (Alternate Method):
      - `C:\>iata_enu.exe -a -a -p c:\IAAdriver`

   c. Using iata_cd - Multi-language version:
      - `C:\>iata_cd.exe -a -a -p c:\IAAdriver`

   **Where `C:\IAAdriver` is an already existing folder on the c: drive



6. After this is done, copy the driver binary files, INF files and the txtsetup.oem file from the local directory `C:\IAAdriver\Driver` to the root directory of the floppy.

10.3.2 Method 2: Using F6 Application

On a system running Windows, download the application f6flpy32.zip for 32-bit operating system or f6flpy64.zip for 64-bit operating system to your local drive.

1. Double click on the zip file and extract the files on to your system and then copy all the files on to the DOS floppy.

2. Now your floppy disk will have the valid Intel® Rapid Storage Technology Driver which can be used for the F6 installation method.



11 Determining the Version of the RAID Driver

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

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

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

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

11.2 Using Intel® RST File Properties (Alternate)

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

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



11.3 Determining the Version of the Option ROM

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

11.3.1 Using the Intel Rapid Storage Technology UI

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

11.3.2 Using the Intel® RST Option ROM User Interface

1. Early in system boot-up, during post, or when you see the "Intel® RAID for Serial ATA" status screen output, type CTRL-I. This will open the Option ROM user interface.
2. The following banner will be displayed:
3. Intel® Rapid Storage Technology option ROM v9.0.0.xxxx ICH10R/DO
4. V9.0.0.xxxx is the version of the Option ROM currently installed on your system. The 9.0.0 portion is the product release number; the xxxx portion is the build number.



12 Un-installation

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

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

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

12.1 Uninstalling the Intel® RST Software (except the RAID Driver)

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

12.2 Disabling the RAID Driver by Disabling the RAID Controller

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

1. Enter System BIOS Setup and disable RAID Mode. This setting may be different for each motherboard manufacturer. Consult your manufacturer's user manual if necessary. When done, exit Setup.
2. Reboot the system (The OS must have been installed on a disk not attached to the Intel® SATA RAID controller). You should no longer see the RAID Option ROM



status screen during boot, and you should no longer see the Intel® SATA RAID Controller in Device Manager.

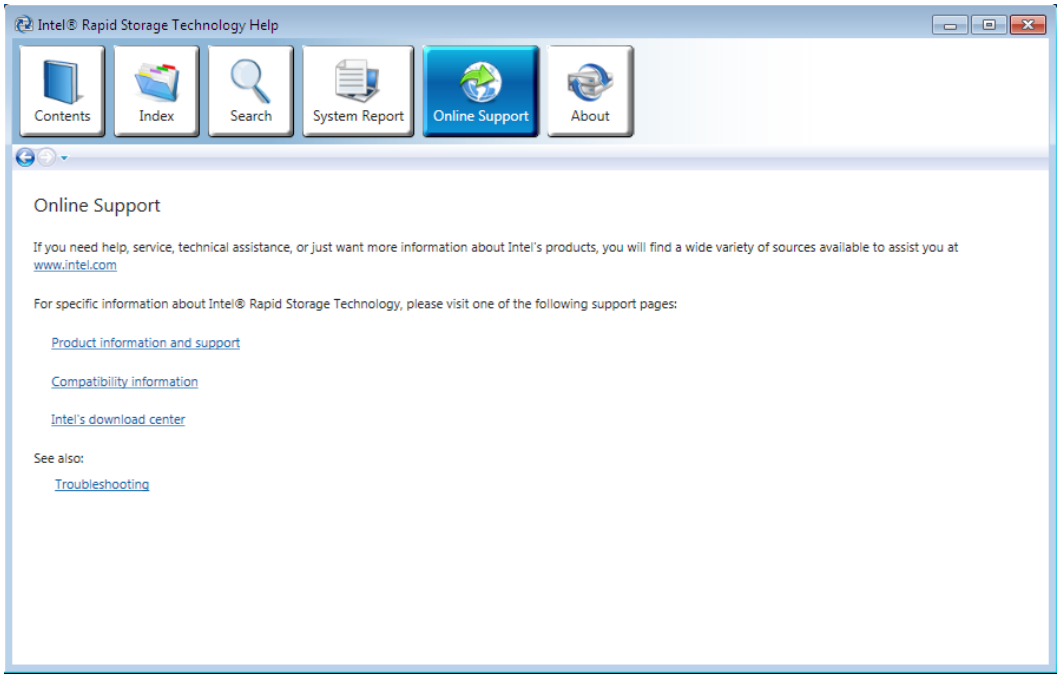
3. At this point, Windows will no longer be using the RAID driver and you will not have Intel RAID functionality. **All data contained in existing RAID volumes will no longer be accessible.** To re-enable Intel RAID functionality, re-enter System BIOS Setup and re-enable RAID mode.

Uninstall Note: End-users can use this same procedure to disable the Intel® SATA RAID Controller if necessary. In fact, the uninstall program used in section 12.1 of this document will display a text file with a similar procedure. Run the Uninstall Program, click 'Cancel' when presented with the first dialog box, then click 'Yes' at the second dialog box to read the text document containing the procedure.

13 Registry Customizations

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

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



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

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

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



13.1 Native Command Queuing Settings

Native Command Queuing (NCQ) is enabled by default on both mobile and desktop platforms.

NCQ can be enabled or disabled using the below registry key.

HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\SATA

"NCQEnable"=dword:00000000; Native Command Queuing (1=Enabled)

13.2 Zero Power ODD Settings

Beginning with the Intel® RST 10.0 release and the Intel® Mobile Express Chipset SATA AHCI and the Intel® Desktop/Workstation/Server Express Chipset SATA AHCI controllers (PCH), the product supports the zero-power ODD feature (also referred to as ZPODD). It is intended to allow an unused ODD to be powered off, and then powered on only upon receipt of demand requests or when the ODD eject button is pressed. This goal is achieved by utilizing ACPI** methods to change the power condition of the drive when several platform (HW) conditions exist.

**Note: This is not supported on Windows XP and older operating systems.

Associated with this feature are two registry keys located at *[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters]*

1. ZPODD enable/disable

```
[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\]

"OddZeroPowerEnable"=dword(0, 1)
```

This key determines a platform's eligibility for the feature. When the value is zero then this feature will be disabled. When the value is non-zero or not present the feature will be enabled. Default value will be enabled (**1**).

2. ODD idle timeout

```
[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\]

"SecondsToOddZeroPower"=dword:(30, 300)
```

This key determines the idle timeout value. When the value is zero then this feature will be disabled. The value is the number of seconds the ODD must be idle (defined as a period of time in which no non-GESN commands are received; minimum value is 30 and maximum value is 300) before the ODD will be powered off. The default value is **60.** If the registry value is set to a value outside this range then the default value of 60 seconds will be used.



13.3 E-mail Notification UI Visible Enable/Disable

By default the e-mail notification feature is visible in the UI under the [Preferences] button. The registry can be modified to allow or disallow the end-user from being able to configure the system to allow email notification. The registry key by default is not populated in the registry. In order to remove the functionality from the UI the registry key has to be created using the following settings:

Open the registry editor and Add the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Intel\IRST

Create a new DWORD (32) value as follows:

DisableEmail

Registry Value	Description
DisableEmail	DWORD(32) = 0 (default): By default or when this value is created and cleared to 0, the UI [Preferences] page displays an item in the left column of the page that allows the user to setup the desired e-mail preferences for the system.
	DWORD(32) = 1: When this value is created and set to 1, the UI will not display a menu item in the 'Preferences' page for the end-user to setup email notification on the system. The feature is disabled.

13.4 Disabling the Intel® SRT Option to Accelerate to Maximized Mode

OEMs have the ability to disable the Accelerate Maximized mode option and limit the Intel® Smart Response Technology to Enhanced mode selection only.

The registry key by default is not populated in the registry. In order to remove the functionality from the UI the registry key has to be created using the following settings:

Open the registry editor and Add the following key:

HKEY_LOCAL_MACHINE\SOFTWARE\Intel\IRST

Create a new DWORD (32) value as follows:



DisablePerformanceMode

Registry Value	Description
DisablePerformanceMode	DWORD(32) = 0 (default): By default or when this value is created and cleared to 0, the UI will allow the end-user to Accelerate disks/volumes in both Maximized and Enhanced mode.
	DWORD(32) = 1: When this value is created and set to 1, the UI will not display the option to Accelerate to Maximized mode and will not allow the user to change from Enhanced mode to Maximized mode



14 Link Power Management with Intel® Rapid Storage Technology

Intel® Rapid Storage Technology implements the Link power management (LPM) feature described by the Serial ATA specification to overcome the power demand of a high-speed serial interface, SATA and providing the capability of SATA at the minimum power cost. LPM, when used in conjunction with a SATA hard drive that supports this feature, enables lower power consumption. LPM was initially enabled by default on mobile platforms starting with ICH6M with Intel® Matrix Storage Manager. Starting with ICH9R this feature has also been supported on desktop platforms with Intel® Matrix Storage Manager 7.5 release but not enabled by default.

Beginning with the Intel® Rapid Storage Technology 10.0 release, LPM support is enabled by default on both mobile and desktop platforms. OEM's who wish to modify the default settings for LPM on their platforms can follow the instructions in the following section(s).

14.1 Instructions to disable/enable LPM

After system is setup with OS and Intel® Rapid Storage Technology installed, follow the below instructions to modify the default LPM support.

NOTE: Beginning with the Intel® Rapid Storage Technology 10.0 release, the registry keys are no longer populated in the Windows registry by default. The RST driver does not require the registry keys to be present to support the default settings.

1. Go to **Start->Run**
2. Type in **RegEdit** and hit the Enter Key.
3. Go to the below mentioned location to insert or configure the registry keys for LPM
   NOTE: *OEM's need to configure the LPM settings per SATA port. Ports are numbered starting with zero (please refer to the desired platform EDS for the number of ports supported on your platform)*

HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\iaStor\Parameters\Port0
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\iaStor\Parameters\Port1
…
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\iaStor\Parameters\Port5

4. Now add the following registry keys under the registry location mentioned in step3, if they are not available (These registry keys are not available by default, they can be added by using automated scripts, .reg files, executable utilities, etc). If you find the below registry keys already available, you can modify the values for desired support. Values are modified on a port by port basis so modify all ports that you wish the changes to be supported on. **



"**LPM**"=dword: 00000001 {dword: 00000000->Disable; dword: 00000001->Enable} [default = Enabled]
"**LPMSTATE**"=dword: 00000000 {dword: 00000000->Partial; dword: 00000001->Slumber} [default = Disabled] (***Note:*** *the driver ignores this key when the LPM key's value is not set to 1. So when LPM value is 0, this value is N/A.*)
"**LPMDSTATE**"=dword: 00000001 {dword: 00000000->Partial; dword: 00000001->Slumber} [default = Enabled]
"**DIPM**"=dword: 00000001 {dword: 00000000->Disable; dword: 00000001->Enable} [default = Enabled]

****Warning:** *If you edit the registry incorrectly, you can cause serious problems that may require you to reinstall your operating system. Intel does not guarantee that problems that are caused by editing the Registry incorrectly can be resolved.*

14.2 Safe Removal with LPM on hot plug capable ports

Please contact your Intel® Rapid Storage Technology representative (AE or FAE) for additional information pertaining to this feature.

NOTE: Request access to *Reference Manual for Safe Removal with Link Power Management (LPM) on Hot Plug Capable Port (HPCP)* white paper (Doc # 445153): (http://www.intel.com/cd/edesign/library/asmo-na/eng/445153.htm)



15 RAID Option ROM and Utilities

There is a unified RAID Option ROM (raid_or.bin) and RAID Utilities (RCfgSata.exe and RCmpSata.exe) beginning with the Intel® Rapid Storage Technology 9.6 release. This unified OROM package is supported on the platforms identified below and all later platform releases when enabled for RAID. With this RAID Option ROM, these platforms now support all RAID functionality based on the hardware capability.

Supported Chipsets					
Intel® 7 Series SATA AHCI / RAID Controller	**Intel® 6 Series Mobile Express Chipset SATA AHCI / RAID Controller**	**Intel® 6 Series Desktop/ Workstation/ Server Express Chipset SATA AHCI / RAID Controller**	**Intel® 5, 3420, and 3450 Series Chipset (PCH)**	**Intel® 4 Series Chipset: ICH10R and ICH10DO**	**Intel® 3 Series Chipset: ICH9M, ICH9ME, ICH9R, ICH9DO, and ICH9DH**
Yes	Yes	Yes	Yes	Yes	Yes



16 HDD Password Support With RAID Volumes

Intel® Rapid Storage Technology supports password protected HDDs to be RAID array member disks and pass-thru disks. The product will rely on the BIOS implementing for most of the ATA Security support. There is a whitepaper available called "Implementing Intel® Matrix Storage Manager Compatible Support for ATA Security in BIOS" available on CDI that describes the necessary BIOS design for compatibility with the Intel Rapid Storage Technology. Rapid Storage Technology product will handle the RAID and hot-plug related behavior with regards to password protected disks.

16.1 HDD Password Use Cases

If at least one unlocked member disk and one locked member disk (with relevant data for the volume) are connected, then a RAID volume will be designated as Locked.

Scenario	Action	Result	Comments
RAID1 Volume Disk 1 – Locked Disk 2 - Unlocked Volume – Locked (Both disks have relevant data)	Remove Disk 1 (locked disk)	Volume becomes unlocked and Degraded. User can rebuild volume unto a new unlocked disk.	The user had authority to access Disk 2 which has the same data as Disk 1, by removing the locked drive the user can access Disk 2.
RAID1 Volume Disk 1 – Locked Disk 2 – Unlocked Volume – Degraded Disk 1 has old data and caused the volume to go Degraded.	None	N/A	The user has access to Disk 2 because the data on Disk 1 is old and irrelevant.



Scenario	Action	Result	Comments
RAID5 Volume Disk 1 – Locked Disk 2 – Unlocked Disk 3 – Unlocked Volume – Locked (All disks have relevant data to Volume)	Remove Disk 1 (locked disk)	Volume becomes unlocked and Degraded. User can rebuild volume onto a new unlocked disk.	The user had authority to access Disk 2 and Disk 3 which has all the data needed to rebuild the volume, by removing the locked drive the user can access Disk 2 and Disk 3 as a Degraded Volume.
Intel® RRT Volume Master Disk – Locked Recovery Disk - Locked (external port docking station) Volume – Locked (Both disks have relevant data)	User connects laptop to docking station and unlocks Recovery disk and Master Disk and boots. Then user takes the laptop from the docking station and leaves the external drive connected to power	The recovery drive can be connected to a new laptop and the information can be used to rebuild an Intel® RRT volume if the power was maintained, because the drive is still in an unlocked state.	Similar situation to a user leaving a laptop unlocked and unattended.

16.2 Unlocking Password Protected Disks in RST UI

By default, the Intel® RST product hides the ability to unlock password protected drives from within the Intel® RST GUI during OS runtime. If the user requires the capability to unlock password protected drives from within the Intel® RST UI this can be enabled via bit6 "HDD Unlock (HDDLK)"of the ***'Intel RST Feature Capabilities'*** register in the SATA controller MMIO space. The default setting is 0 which is hidden (bit6=0). The 'Unlock' option is displayed in the UI when bit6 is set to 1.

17 Intel® Smart Response Technology

¥Note: This feature requires that the SATA controller be set to RAID mode via the system BIOS. There is no support in AHCI mode.

Intel® Smart Response Technology is an Intel® RST caching-related feature that improves computer system performance and lowers power consumption for systems running on battery power while in Maximized mode. It allows OEMs to configure computer systems with an SSD used as cache memory between the hard disk drive and system memory. This provides the advantage of having a hard disk drive (or a RAID volume) for maximum storage capacity while delivering an SSD-like overall system performance experience. Intel® Smart Response Technology caching is implemented as a single drive letter solution; no additional drive letter is required for the SSD device used as cache.

Supported Chipsets					
Intel® 7 Series SATA AHCI / RAID Controller	**Intel® 6 Series Mobile Express Chipset SATA AHCI / RAID Controller**	**Intel® 6 Series Desktop/ Workstation/ Server Express Chipset SATA AHCI / RAID Controller**	**Intel® 5, 3420, and 3450 Series Chipset (PCH)**	**Intel® 4 Series Chipset: ICH10R and ICH10DO**	**Intel® 3 Series Chipset: ICH9M, ICH9ME, ICH9R, ICH9DO, and ICH9DH**
Yes£	Yes£	Yes£	No	No	No

£Note: This feature is only supported on designated RAID enabled SKUs

17.1 Overview

17.1.1 Warnings!

17.1.1.1 Driver/OROM updates of Accelerated systems:

- Updating a **production version** of the driver to a newer **production version** of the driver requires a **production version** of the 10.5.0 or later OROM (***Note: the system will be limited to the features available to the OROM that is installed on the system***)



- Updating the OROM to a newer version requires that the driver version be updated to a driver version from the same release package of the OROM or newer.

17.1.1.2 Hot removal (hot unplug) of Maximized Accelerated components:

- **DO NOT** physically hot unplug any component of the Accelerated system that is in Maximized mode. Before removal of either the Accelerated disk/volume or the "Cache SSD" the user must disable (turn-off) Acceleration. This will eliminate the potential loss of any data that is being accessed or has not been flushed from the cache to the Accelerated disk/volume.

Table 17-1: Example of driver/OROM compatibility

	OROM Version			
Driver Version	10.5.0 PV	10.5.1 PV	10.6.0 PV	11.0.0 PV
10.5.0 PV	O	**X**	**X**	**X**
10.5.1 PV	S	O	**X**	**X**
10.6.0 PV	S	S	O	**X**
11.0.0 PV	S	S	S	O

- X = this configuration is not supported
- O = this configuration is supported and is optimal for the driver and OROM
- S = this configuration is supported; however, it is limited to the features of the driver that was originally released with that OROM version. E.g. if the 11.0.0 PV driver is installed/updated to a system running the 10.5.0 PV, the system will be limited to the features of the 10.5.0 OROM. Any new features associated with the 11.0.0 PV release will not be enabled with this configuration.

17.1.2 Requirements and Limitations

System Requirements:

For a system to support Intel® Smart Response Technology it must have the following:

- One of the following platforms (RAID enabled SKUs):
  - Intel® 7 Series Chipset SATA RAID Controller:
    - Mobile: QM77, HM77, QS77, UM77
    - Desktop: Q77, Z77, H77



        - WS/Server: C216
    - Intel® Mobile Express Chipset SATA RAID Controller:
        - QM67, HM67
    - Intel® Desktop/Workstation/Server Express Chipset SATA RAID Controller:
        - Z68
- CPU:
    - Intel® Core™ i3, i5, or i7 processor (only works on production processors that have the 'Core' or 'Xeon' band name)
    - Intel® Xeon® brand processors (only works on production processors that have the 'Core' or 'Xeon' band name)
- Operating system:
    - Windows* 7 (32/64 bit)
    - Windows* Server 2008 R2 x64
    - Windows* Server 2008 (32/64 bit) and Data Center
    - Windows Vista* (32/64 bit): ***Note that Vista is not supported on Intel® 7 Series and later chipset-based platforms***.
- System BIOS with the Intel® Smart Response Technology caching bit (bit 9) of the ***'Intel RST Feature Capabilities'*** register in the SATA controller MMIO space enabled (set to 1; the default setting is 0).
- System BIOS with SATA mode set to RAID enabled
- System BIOS that is PCI 3.0 and PMM (POST Memory Management) compliant to allow the OROM to handle dirty shutdowns of Accelerated disks/volumes; ***PMM must be able to allocate a minimum of 130MB of temporary, non-aligned, extended memory to the Legacy OROM***
- Intel® RST driver and OROM installed from the production Intel® RST 10.5 version release or later production releases
- Flash part must budget the following space for the Intel® RST OROM:
    - Image file size ~119KB
    - Runtime size ~41.5KB

"Cache SSD" Requirements:

For an SSD to meet the Intel® Smart Response Technology "Cache SSD" criteria it must have the following:

- 18.6GB minimum capacity

The Intel RST product recognizes a device as an SSD if its IDD data structure word
217 = 0x01

Feature Limitations of Intel® Smart Response Technology:

- No support for this feature on Windows XP (all versions)



- A total of only one pass-through disk (RAID Ready system) or RAID volume can be 'Accelerated' per computer system
- Accelerated Volume criteria:
    - Accelerated RAID volumes are limited to 31.5TB or less
    - RAID levels 0, 1, 5, and 10
        - No IRRT volumes allowed to be accelerated
    - All array member disks must be HDDs
        - No arrays with SSD member disks allowed to be accelerated
- No Acceleration of a RAID volume that is one of the volumes of a Matrix array (multiple volumes on a single array)
- The SSD to be used as the "Cache SSD" must be attached to an internal SATA port of the computer

17.1.3 Acceleration Modes

There are three 'Acceleration' modes of operations for Intel® Smart Response Technology caching.

Acceleration Mode Values and Limitations

Mode	Performance	Separation Safe	Data Caching
Off	No Acceleration	N/A	N/A
Enhanced	Boot-time, Run-time reads, Paging	Yes	Write-Through
Maximized	Boot-time, Run-time reads & writes, Paging	No	Write-Back

17.2 Build a New System with Disk/Volume Acceleration Enabled

Note: This section requires that the computer system's Intel® SATA AHCI/RAID controller be set to RAID mode via the system BIOS. This applies to both the 'New System' and the 'Build System' when using RAIDCfg32.

Note: There is no ability to enable Acceleration while in the OROM UI. Acceleration must be enabled either in the Intel® RST UI during OS runtime or the RCfgSata



utility (as described below) if required to do so pre-OS. The OROM UI only allows disabling of Acceleration.

17.2.1 Prepare New Computer

The following instructions are for preparing a brand new computer system to be built with disk/volume Acceleration enabled.

17.2.1.1 Meet System Requirements

For Acceleration requirements, see section Requirements and Limitations. Ensure that the targeted system meets all the requirements for Acceleration.

17.2.1.2 Determine Type of System Configuration

Although the OEM can configure a system any number of possible configurations, the following are four of the most common configurations. Config1 will be used in this setup and configuration example detailed in the next few sections.

1. ***Config1: OS installed on an Accelerated pass-through disk; RAID Ready system***
2. Config2: OS installed on an Accelerated RAID volume (RAID 0, 1, 5, or 10)
3. Config3: OS installed on a non-Accelerated pass-through disk with an Accelerated data disk; RAID Ready system
4. Config4: OS installed on a non-Accelerated RAID volume with an Accelerated data disk

NOTE: For systems that will share the SSD between Intel® Smart Response Technology cache and Intel® Rapid Start Technology, skip to section 17.2.1.4

17.2.1.3 Setup the HW for Installing the OS to an Accelerated Disk/Volume

17.2.1.3.1 Install HDD(s)

1. Select the required HDD(s) needed for the type of system configuration
2. Locate SATA port(s) and attach HDD(s). (**Note** the port number of the pass-through HDD to be used for the OS system disk.)
3. Install any other HW peripheral desired for the system configuration (e.g. ODD)



17.2.1.3.2 Install SSD Meeting Intel® Smart Response Technology Caching Criteria

1. Select an SSD that has a minimum capacity of 18.6GB.
2. Locate a SATA port that is configured as 'Internal' and attach the SSD. Note the port number of the SSD.

17.2.1.3.3 Configure the SSD to be the "Cache SSD"

This section will step through the process for setting up the SSD and HDD for Config1 where the OS will be installed on a single pass-through HDD that will be accelerated in Maximized mode.

- **If Using RCfgSata**

1. Copy the RCfgSata tool from the RST 10.5 release or later to a DOS bootable media (e.g. USB thumb drive) and attach the media device to the targeted new system

2. Boot to the DOS bootable media

3. Setup the HDD:

   At the command line type: **rcfgsata /c Sys_Vol /ds 0** (where 'Sys_Vol' is the logical name of the single pass-through disk and '0' is the Internal configured port where the single physical disk is located)

4. Setup the SSD to be the "Cache SSD":

   At the command line type: **rcfgsata /c Cache_Dev /ds 3** (Where 'Cache_Dev' is the logical name representing the "Cache SSD" and '3' is the Internal configured port location of the physical SSD. (**Note**: if the SSD is larger than 64GB, the following command will be required: **rcfgsata /c Cache_Dev /ds 3 /s 18.6**; where **/s** is for the size of the caching region in GB (18.6 GB is the minimum and 64 GB is maximum supported size))

   OR

   4a. Configure the SSD with both Cache and User Data regions (or volumes):

   Cache region: **rcfgsata /c Cache_Dev /ds 3 /s 64** (where 64 is max size of 64GB and 18.6 is the minimum)

   User Data region: **rcfgsata /c SSD_UserData /ds 3** (this will default to the remaining capacity of the SSD)

5. Pre-configure the pass-through disk (OS system disk) for Acceleration:



At the command line type: **rcfgsata /accel Sys_Vol Cache_Dev max** (Where Sys_Vol is the single disk to be Accelerated, 'Cache_Dev' is the "Cache SSD", and 'max' indicates the Acceleration mode is 'Maximized' .

Note: *The disk/volume is not actually Accelerated until the system is booted and the pre-configured Accelerated disk/volume and the pre-configured "Cache SSD" are enumerated by the Intel® Rapid Storage Technology driver.*

6. Reboot; the 'New System' is now prepared and ready for Windows OS installation to a pre-configured Accelerated pass-through disk.

- **If Using RAIDCfg32/64 (compatible with WinPE)**

1. Locate a Widows system that meets all the system requirements for Acceleration; let's call this the 'Build System'. (**Note**: The 'Build System' should not have any SSD's or HDD's already configured for Acceleration. If so, remove Acceleration and reset any SSD configured as a "Cache SSD" to an Available SSD.)

2. Locate the HDD that will be used as the single pass-through disk that will have the OS installed and Accelerated on the new system and attach it to an unused SATA port on the 'Build System'

3. Locate the SSD that will be used as the "Cache SSD" on the new system and attach it to an 'Internal' configured SATA port on the 'Build System'. (**Note**: the 'Build System' cannot already have an SSD configured as a "Cache SSD")

4. Boot the 'Build System' into Windows (login as administrator) and launch a DOS prompt command line. If not already done so, copy a version of the RAIDCfg32/64 application from the RST 10.5 Release or later to a directory on the 'Build System'.

5. At the DOS prompt command line (note that **raidcfg32** and **raidcfg64** are interchangeable in the below example).

Accelerate a single pass-through HDD using the /ACCEL**D** command:

- **Setup the SSD to be the "Cache SSD":**

At the DOS Prompt, type:
raidcfg32 /c Cache_Dev /ds Y /s 00.0 /nvcache

(Where 'Cache_Dev' is the logical name for the "Cache SSD" targeted for the 'New System,

replace 'Y' with the port number on the 'Build' system where the SSD is physically located,



where '/s 00.0' represents the Cache volume size (replace 00.0 with any number between 18.6 and 64),

and **/nvcache** designates this SSD as the cache device.)

To use the maximum capacity of the SSD for the cache memory, up to a maximum of 64GB, remove the **/s 00.0** from the command line. To use a custom size use **/s 18.6** and replace 18.6 with any number between 18.6 and 64.

- **Accelerate the pass-through disk (this is the disk planned to be the OS system disk for the 'New System'):**

  At the DOS Prompt, type:
  raidcfg32 /acceld # /mode max

  (Where the single pass-through disk is located a SATA port number '#' and 'max' represents that it is in Maximized Acceleration mode.)

Skip to step 6 below:

Accelerate a multi-disk RAID volume (i.e. RAID1) using /ACCEL command:

- **Setup the SSD to be the "Cache SSD":**

  At the DOS Prompt, type:
  raidcfg32 /c Cache_Dev /ds Y /s 00.0 /nvcache

  (Where 'Cache_Dev' is the logical name for the "Cache SSD" targeted for the 'New System,

  replace 'Y' with the port number on the 'Build' system where the SSD is physically located,
  where '/s 00.0' represents the Cache volume size (replace 00.0 with any number between 18.6 and 64),

  and **/nvcache** designates this SSD as the cache device.)

  To use the maximum capacity of the SSD for the cache memory, up to a maximum of 64GB, remove the **/s 00.0** from the command line. To use a custom size use **/s 18.6** and replace 18.6 with any number between 18.6 and 64.

- **Accelerate the RAID volume (this is the RAID volume planned to be the OS system disk for the 'New System'):**

  At the DOS Prompt, type:
  raidcfg32 /accel Volume_Name /mode max

  (Where the RAID volume to be accelerated is named 'Volume_Name' and '/mode max' represents that it is in Maximized Acceleration mode.)



6. Power down the 'Build System' and physically remove the "Cache SSD" and the Associated single pass-through disk that are targeted for the new system. (**Note:** To remain valid, the preconfigured "Cache SSD" and Accelerated HDD must be installed as a pair in a system that has no Accelerated Disk/Volume or "Cache SSD" already installed.)

7. In the 'New System', install the "Cache SSD" and the Associated pass-through disk onto the desired SATA ports (**Note**: the SSD must be installed to an Internal configured port).

8. The 'New System' is now prepared for OS installation to an Accelerated single pass-through disk

- **If Using the Intel® RST UI**

  This process is similar to using RAIDCfg32/64

  1. Repeat the steps 1 – 4 from previous section "If Using RAIDCfg32/64"

  2. Launch the Intel® RST UI

  3. Click on the 'Accelerate' tab at the top of the UI and click the Enable acceleration link

     a. If multiple SSDs on the build system, select the SSD that will be used in the 'New System'

     b. Select the size to be allocated on the SSD for cache memory (18.6GB or Full disk capacity(maximum 64 GB))

     c. Select the HDD to be accelerated and will be used in the 'New System'

     d. Select the mode of Acceleration and click [OK]

  4. Power down the 'Build System'. Remove the "Cache SSD" and the Associated single pass-through disk that are targeted for the 'New System'.

  5. On the 'New System', install them into the desired SATA ports (**Note**: the SSD must be installed to an Internal configured port).

  6. The 'New System' is now prepared for OS installation to an Accelerated single pass-through disk.

  7. Skip the next section and go to section 17.2.2 for OS installation instructions.



17.2.1.4 Setup the HW for Installing the OS to an Accelerated Disk/Volume with FFS

NOTE: For systems not using Intel® Rapid Start Technology, go to previous section 17.2.1.3

This section has detailed information regarding configuring systems so that a single SSD can co-exists between the Intel® Smart Response Technology caching solution and the Intel® Rapid Start Technology solution. This document only covers configuration information as it is related to Intel® Smart Response Technology.

The following configuration instructions make the assumption that the OEM/user has properly met all the requirements for Intel® Rapid Start Technology. There are NO instructions in this document for configuring the platform for Intel® Rapid Start Technology. For detailed configuration and setup requirements for Intel® Rapid Start Technology please contact your Intel representative for assistance.

You will also see the acronym **FFS** to represent the portion of the SSD that is required for Intel® Rapid Start Technology.

17.2.1.4.1 Install HDD(s)

1. Select the required HDD(s) needed for the type of system configuration
2. Locate SATA port(s) and attach HDD(s). (**Note** the port number of the pass-through HDD to be used for the OS system disk.)
3. Install any other HW peripheral desired for the system configuration (e.g. ODD)

17.2.1.4.2 Install SSD Meeting Intel® Smart Response Technology Caching and FFS Criteria

1. Select an SSD that has a minimum capacity of 24 GB (18.6GB cache portion + 4GB FFS portion).
2. Locate a SATA port that is configured as 'Internal' and attach the SSD. Note the port number of the SSD.

17.2.1.4.3 Configure the SSD to be the "Cache SSD" with FFS

This section will step through the process for setting up the SSD and HDD for Config1 where the OS will be installed on a single pass-through HDD that will be accelerated in Maximized mode.

- **If Using RCfgSata**

    1. Copy the RCfgSata tool from the Intel® RST production release to a DOS bootable media (e.g. USB thumb drive) and attach the media device to the targeted new system



2. Boot to the DOS bootable media

3. Setup the HDD:

   At the command line type: **rcfgsata /c Sys_Vol /ds 0** (where 'Sys_Vol' is the logical name of the single pass-through disk and '0' is the Internal configured port where the single physical disk is located)

4. Setup the SSD to be the "Cache SSD" and the FFS:

   At the command line type:

   For the FFS region:

   rcfgsata /c FFS /ds 3 /s 4

   - **FFS** is the name for the volume
   - **3** is the SATA port location of the SSD
   - **4** is the size (GB) of the FFS partition; **CAUTION: the minimum size requirement is based on system DRAM and other Intel® Rapid Start Technology criteria. Before selecting the size, first consult your Intel® Rapid Start Technology representative for guidance)**

   For the Cache region:

   rcfgsata /c Cache_Dev /ds 3

   (this will default to the remaining capacity of the SSD and **must be 18.6GB or greater, if not the cache will not be functional**)

5. Pre-configure the pass-through disk (OS system disk) for Acceleration:

   At the command line type: **rcfgsata /accel Sys_Vol Cache_Dev max** (Where Sys_Vol is the single disk to be Accelerated, 'Cache_Dev' is the "Cache SSD", and 'max' indicates the Acceleration mode is 'Maximized' .

   Note: *The disk/volume is not actually Accelerated until the system is booted and the pre-configured Accelerated disk/volume and the pre-configured "Cache SSD" are enumerated by the Intel® Rapid Storage Technology driver.*



6. Reboot; the 'New System' is now prepared and ready for Windows OS installation to a pre-configured Accelerated pass-through disk.

- **If Using RAIDCfg32/64 (compatible with WinPE)**

  1. Locate a Widows system that meets all the system requirements for Acceleration; let's call this the 'Build System'. (**Note**: The 'Build System' should not have any SSD's or HDD's already configured for Acceleration. If so, remove Acceleration and reset any SSD configured as a "Cache SSD" to an Available SSD.)

  2. Locate the HDD that will be used as the single pass-through disk that will have the OS installed and Accelerated on the new system and attach it to an unused SATA port on the 'Build System'

  3. Locate the SSD that will be used as the "Cache SSD" on the new system and attach it to an 'Internal' configured SATA port on the 'Build System'. (**Note**: the 'Build System' cannot already have an SSD configured as a "Cache SSD")

  4. Boot the 'Build System' into Windows (login as administrator) and launch a DOS prompt command line. If not already done so, copy a version of the RAIDCfg32/64 application from the RST 11.0 or later Release to a directory on the 'Build System'.

  5. At the DOS prompt command line (note that **raidcfg32** and **raidcfg64** are interchangeable in the below example).

  Accelerate the single pass-through HDD using the /ACCEL**D** command in the following order:

  - Setup the SSD to be the "Cache SSD" and FFS device:

    Cache region:

    At the DOS Prompt, type: **raidcfg32 /c Cache_Dev /ds Y /s 00.0 /nvcache**

    - where 'Cache_Dev' is the logical name for the "Cache SSD' targeted for the 'New System,
    - replace 'Y' with the port number on the 'Build' system where the SSD is physically located.
    - Replace **00.0** with a number between 18.6 and 64. This number will represent the desired size of the cache volume. The remaining capacity can be used for the FFS partition.

    FFS region:

    The FFS region will be the remaining capacity of the SSD once the cache region is determined and created.



E.g. for an SSD that is 30 GB, if you select the cache region to be 22GB (replace 00.0 with 22), the user data region will be 8GB which can be used for the FFS partition

- Accelerate the pass-through disk (this is the disk planned to be the OS system disk for the 'New System'):

  At the DOS Prompt, type: **raidcfg32 /acceld T /mode max**
  (Where the single pass-through disk is located a SATA port number 'T' and 'max' represents that it is in Maximized Acceleration mode.)

6. Power down the 'Build System' and physically remove the "Cache SSD" and the Associated single pass-through disk that are targeted for the new system. (**Note:** To remain valid, the preconfigured "Cache SSD" and Accelerated HDD must be installed as a pair in a system that has no Accelerated Disk/Volume or "Cache SSD" already installed.)

7. In the 'New System', install the "Cache SSD" and the Associated pass-through disk onto the desired SATA ports (**Note**: the SSD must be installed to an Internal configured port).

8. The 'New System' is now prepared for OS installation to an Accelerated single pass-through disk with an FFS partition that is ready to be configured for Intel® Rapid Start Technology.

- **If Using the Intel® RST UI**

  This process is similar to using RAIDCfg32/64. However, the UI is limited to selecting the minimum size of the Cache volume of 18.6GB. The remaining capacity on the SSD can be used for the FFS partition requirements.

  1. Repeat the steps 1 – 4 from previous section "If Using RAIDCfg32/64"

  2. Launch the Intel® RST UI

  3. Click on the 'Accelerate' tab at the top of the UI and click the Enable acceleration link

     a. If multiple SSDs on the build system, select the SSD that will be used in the 'New System'

     b. Select the size to be allocated on the SSD for cache memory (18.6GB or Full disk capacity(maximum 64 GB)). **Note that all remaining capacity on the SSD can be used for the FFS partition**; **ensure that capacity will meet the minimum FFS capacity requirements for your system**



configuration (contact your Intel® Rapid Start Technology representative for guidance).

c. Select the HDD to be accelerated and will be used in the 'New System'

d. Select the mode of Acceleration and click [OK]

4. Power down the 'Build System'. Remove the "Cache SSD" and the Associated single pass-through disk that are targeted for the 'New System'.

5. On the 'New System', install them into the desired SATA ports (**Note**: the SSD must be installed to an Internal configured port).

6. The 'New System' is now prepared for OS installation to an Accelerated single pass-through disk with an FFS partition that is ready to be configured for Intel® Rapid Start Technology.

17.2.2 Installing the OS to a New System Prepared for Disk/Volume Acceleration

To install the OS to a disk/volume previously configured for Acceleration, the RST driver installation files will be required. Download the **f6** zipped archive file (from the RST 10.5 Release or later kit) and extract the files to a media that can be accessed during OS installation (e.g. USB thumb drive).

17.2.2.1 Prepare the New System for OS Installation

17.2.2.1.1 For Acceleration Components Pre-configured Via RCfgSata (DOS)

If the Accelerated disk and the Cache SSD" were prepared using the RCfgSata utility then all that should be required is to boot the system to the Windows OS installation media in the ATAPI ODD.

17.2.2.1.2 For Acceleration Components Pre-configured Via RAIDCfg32/64 (Windows)

1. If not done already, locate the desired SATA port(s) of the 'New System' and install the disk/volume to be Accelerated that was previously configured on the 'Build System'.



2. Locate the desired Internal configured SATA port of the 'New System' and install the SSD that was previously configured on the 'Build System'

17.2.2.2 OS Installation

1. Boot to the Windows OS installation media (ensure that the media with the RST driver, e.g. USB thumb drive, is not installed in the system during boot)

2. When prompted to load driver, insert the RST driver installation media and click the Load Driver link.

3. Load the RST driver for your computer system

4. The disk/volume with Acceleration enabled, should now be available in the list of storage drives.

5. Select the Acceleration enabled disk/volume and continue with the normal OS installation procedure from this point.

Once the installation is complete nothing else is required to enable Acceleration. Acceleration should be enabled on the disk/volume that was configured via the RCfgSata/RAIDCfg32 tool or RST UI.

17.3 OEM System Manufacturing and Intel® Smart Response Technology

Note: This section requires that the computer system's Intel® SATA AHCI/RAID controller be set to RAID mode via the system BIOS. This applies to the 'Build System' when using RAIDCfg32/64. In the example below, RAIDCFG32 is used but can be interchangeable with RAIDCFG64

Note: The following procedures are for setting up a single disk (RAID mode pass-through disk) in Acceleration mode with an OS image pre-configured and installed. OEMs need only move and install the imaged Accelerated disk along with its Associated "Cache SSD" as a pair to a properly configured computer system. The system will boot up with the Acceleration mode enabled.

17.3.1 Imaging an OS onto a Pre-Configured Acceleration-enabled HDD

Note: This section allows the OEM to pre-configure an HDD/SSD Accelerated pair and image the OS onto that Accelerated HDD. The HDD and SSD pair can then be



removed from the manufacturing system and assembled together in the final system to be shipped to the end-user.

Note: The HDD and SSD must remain together at all times as an Accelerated pair.

17.3.1.1 What You Need (Build Environment)

1. Computer platform (call it the 'Build System') that supports Intel® Smart Response Technology
2. HDD (targeted for Acceleration and OS image) and SSD (targeted for "Cache SSD")
3. Media with bootable WinPE environment (USB thumb drive, ODD, etc) that has:
   a) Intel® RST driver (compatible with RAID OROM on 'Build Platform'); driver must be loaded during WinPE boot or loaded using the ***drvload*** command.
   b) RAIDCfg32/64 executable
   c) The OS image

17.3.1.2 Bring up the Build Environment

1. Attach SSD to an Internal SATA port on the 'Build System'
2. Attach the HDD to a SATA port on the 'Build System'
3. Boot into WinPE with Intel® RST driver, RAIDCfg32/64, and OS image
4. If the Intel® RST driver is not loaded, load the driver

17.3.1.3 Pre-Configured the SSD/HDD for Acceleration

1. Prepare the SSD to be the "Cache SSD": at the command prompt/> type

   raidcfg32 /c CacheDeviceName /ds 1 /s 18.6 /nvcache
   *(where **1** represents the port location of the targeted SSD, **18.6** is the desired size of the cache volume)*

2. Enable Acceleration on the single HDD: at command prompt/> type

   raidcfg32 /acceld 2 /mode max
   (where 2 = the SATA port location of the HDD to be Accelerated, and max = Maximized mode, (enh = Enhanced mode))



17.3.1.4 Complete the Pre-Configuration Process

1. Use your imaging program (e.g. Ghost) to transfer the OS image to the Accelerated HDD that you just setup in the steps above (**Note**: the OS image must have Acceleration enabled)

2. Once the imaging process has completed, the Accelerated pair (HDD + SSD) can be moved to a supported platform to build a fresh new system with Intel® Smart Response Technology already pre-configured

3. This completes the process

17.3.2 Enabling Acceleration post end user OOBE

OEM's may wish to enable acceleration of a system preconfigured for Intel® Smart Response Technology as the last step of the customer OOBE. This process allows the OEM to make minimal changes to their manufacturing environment. All the OEM would have to do is:

1. Setup the system for Acceleration of the OS system disk

    a. Make sure that along with the HDD for the system disk, that an SSD is installed in the system

    b. System must meet chipset and CPU requirements for Intel® Smart Response Technology

    c. System BIOS is properly configured with the Intel® RST RAID OROM and the SATA controller set to RAID mode

    d. Image the HDD with Intel® RST RAID driver installed

2. Sysprep the system for OOBE

3. Once system boots for first time after OOBE automatically run a script in the background to enable Acceleration of the system disk; see an example batch file in the next bullet:

17.3.2.1 Step 1: Setup the system for Acceleration of the OS Disk

17.3.2.1.1 Prepare system HW

1. At least one HDD (for the OS image) and one SSD (for the cache device; must be installed to an Internal configured SATA port)

2. Ensure that system BIOS has a properly integrated Intel® RST OROM that supports Intel® Smart Response Technology. Set the SATA controller to RAID mode via the system BIOS



17.3.2.1.2 Transfer OS Image the HDD

1. Prepare the final master OS image with Intel® RST RAID driver (the OS image system must have been built on a system with the RAID mode set in the BIOS so that the Intel® RST RAID driver is the installed storage driver)

2. Transfer the RST RAID-enabled master OS image to the HDD in the new system (use whatever is current OEM process, e.g. Ghost)

17.3.2.2 Step 2: System preparation of the system for OOBE

OEM should run their normal process to prepare the system for the end-user's initial boot-up and login of the system (OOBE).

17.3.2.3 Step 3: Auto-run a script to enable Acceleration of the OS disk

Acceleration script for post OOBE using the RAIDCFG32/64

1. Because the RAIDCFG32/64 utilities are not supported in the end-user environment, the script should delete the utility upon completion or the utility shall be located on the computer in a location that is not accessible to the end user

2. Once the system boots for first time after OOBE, automatically run a script in the background to enable Acceleration of the system disk; see an example batch file in the next bullet:

3. **Accel_SysDisk.bat** (for example this batch file with the following commands included in it will Accelerate the system disk):

   3a: RAIDCFG32 /C CACHEDEV /DS 1 /S 18.6 /NVCACHE /Q /Y :
   *This command line prepares the SSD on port 1 as the cache device; **CACHEDEV** = the volume name for the cache device, **1** = the SATA port location of the SSD, and **18.6** = the minimum size of the cache volume on the SSD*

   3b: RAIDCFG32 /ACCELD 0 /MODE MAX /Q /Y *: This command line accelerates the system disk on port **0** using the SSD on port 1 as the cache device; **MAX** = Maximized mode*

   EOF

Once the script using the 2 bolded command lines above completes, the system should be in Maximized Acceleration mode with a cache size of 18.6GB.



17.4 OEM System Manufacturing and Cache Pre-load for Intel® Smart Response Technology

The cache pre-load feature was developed to help OEMs mass produce systems configured with Intel® Smart Response Technology enabled with a cache that is out-of-the-box in a near optimal performance state.

The Intel® Smart Response Technology caching solution is a learning solution. This means that when the cache is initially enabled, there is little to no data being cached. This initially results in many cache misses causing the host to have to access the HDD for I/O requests. However, over time, the caching policies of Intel® SRT places data in the cache that is accessed often. So after some time the cache will be loaded with often used data giving the system its optimal or maximum performance configuration.

The problem with this is that when a new system is first used by an end user, the system will have no data cached and thus the performance gains expected of the cache will be small. Depending on use, it could take days of use before the end user starts seeing the expected performance as the cache learns what data should be stored in the cache.

To overcome this initial poor performance gain is to ship the system in the box with the cache already loaded with user data. This could be accomplished in one of two ways. The OEM could configure the system then spend weeks using the system so that it learns what data to load into the cache, or the OEM could preload data into the cache that is likely to be used immediately by the end user.

The following sections dictate the process that the Intel® SRT solution uses to pre-load the cache to be shipped in the box ready for an optimal OOBE performance-wise. It will detail the steps for setting up a system with an Accelerated single pass-through disk with a pre-loaded cache.

17.4.1 Requirements

1. Must meet all Intel® Smart Response Technology requirements
2. Intel® RST 11.1 production release or later (includes 11.1 driver and 11.1 OROM)
3. Optional: The NVCacheScripts.zip file from the RST SW kit on VIP

17.4.2 Process

The NV cache loading process is a two stage process (assuming that SRT caching has already been enabled in Maximized mode):

1. Registry and RST driver setup, reboot.

2. NV cache content loading, cleanup, and shutdown.



17.4.2.1 Setup system for Cache Loading

1. Configure an Accelerated system in Maximized mode and install the OS and desired applications
2. Download NvCacheScripts archive (.zip) and RAID configuration utility from Intel VIP site (use same kit as the Intel® RST driver and OROM)
3. If not already exist, create the C:\Intel directory on the system disk
    a. Copy the NvCacheScripts.zip file to this directory.
4. Unzip the archive (double click on it and select to extract it to the default directory which should be C:\Intel\NvCacheScripts\). The following files should be extracted into the directory:
    - cache_cleanup.reg
    - cache_insert.reg
    - Readme.txt
    - step1_RegistrySetup.bat
    - step2_LoadNVCache.bat

    a. Copy into the C:\Intel\NvCacheScripts\ directory either of the following RAID configuration utilities, depending upon your Intel® RST driver version; RAIDCFG32.exe, RAIDCFG64.exe, RSTCLI.exe, or RSTCLI64.exe
5. Open a DOS prompt command window and change directories to C:\Intel
6. At the DOS prompt type raidcfg32 /st <enter> and note the port number of the system HDD and the SSD (e.g. assume port = 0 for the HDD and port = 1 for the SSD)
7. Edit the step1_registrySetup.bat file by replacing the following command line:
    a. Raidcfg32.exe /ACCEL:SN# /MODE:MAX /Y

    With the following command line:

    b. raidcfg32.exe /acceld 0 /mode max /Y (***Note**: this line can be removed if you always build your systems in Maximized mode*)
8. Run the script **step1_registrysetup.bat** (this will change the caching policy for cache loading then reboot the system for the new policy to take effect)
9. Once the system reboots, it is ready for loading of the cache

17.4.2.2 Cache Loading and Cleanup

1. Edit the step1_registrySetup.bat file by replacing the following command line:



a. Raidcfg32.exe /ACCEL:SN# /MODE:Enh /Y

With the following command line:

b. raidcfg32.exe /acceld 0 /mode max /Y (***Note****: select the mode that you plan to ship the system in*)

2. Run the script **step2_LoadNVCache.bat,** this will:
   a. Copy selected files to the cache
   b. Change Acceleration mode if required
   c. Cleanup the registry to return to the default SRT caching policy
   d. Cleanup the cache loading directory and reboot the system
3. The system now is ready to be shipped with the cache loaded for optimal performance out of the box

17.4.3 Replication for Mass Production of Accelerated HDD and SSD

Now that system has pre-loaded the cache, the OEM can consider the images on the HDD and SSD to be the gold images for that Accelerated system. These disks can now be mass produced (replicated) and paired into systems and the HDD and SSD will be Associated by the OROM upon first boot up as an accelerated pair.

17.4.3.1 Rules and Limitations

The disk image replication process should not itself cause data to be loaded into the cache. This could cause the cache to become full and evict cached data from the gold cache image. To prevent this, the OEM must employ image copy tool that only uses "block copying" and not "file copying".

Once the disks are duplicated they can be paired at a later time and installed in systems on the manufacturing line. Any SSD and any HDD can be paired because the OROM will fix the metadata on both the SSD and the HDD the very first time they are booted together after the replication process.

17.5 HDD ATA Power-Up in Standby (PUIS) Supporting Intel® Smart Connect Technology‡

Note: This '*ATA PUIS for HDDs*' feature requires:



- **that the SATA controller be set to RAID mode via the system BIOS (there is no support in AHCI mode)**
- **that the OS be installed to a PUIS compatible HDD**
- **that Intel® Smart Response Technology be enabled and in Maximized mode**
- **and that the Intel® Smart Connect Technology[‡] Version 2.0 feature be enabled (for information on installation, configuration, and usage of this feature, please refer to its documentation CDI/IBP document # 482930 or contact your Intel® Smart Connect Technology representative)**

The Intel® Smart Connect Technology (Intel® SCT) is a feature of the platform in which the software on the platform and combination of NIC (LAN/WLAN/WWAN) features provides content updates during periods when the user is away from the PC and the PC is in a power saving mode. In addition, if the system has the Intel® Rapid Storage Technology driver installed with the Intel® Smart Response Technology caching feature enabled in Maximized mode and the pass-through HDD supports the PUIS feature of the ATA specification, then there is an additional power savings benefit during the periods of content update. The PUIS feature allows the HDD to stay in a powered-down state (the drive does not spin-up) during the periods of content update.

[‡]***The Intel® Smart Connect Technology*** *feature is not a part of the Intel® RST SW suite. Contact your Intel field representative for more details regarding this feature.*

Supported Chipsets					
Intel® 7 Series Chipset SATA AHCI / RAID Controller	**Intel® 6 Series Mobile Express Chipset SATA AHCI/RAID Controller**	**Intel® 6 Series Desktop/ Workstation/ Server Express Chipset SATA AHCI/RAID Controller**	**Intel® 5, 3420, and 3450 Series Chipset (PCH)**	**Intel® 4 Series Chipset: ICH10R and ICH10DO**	**Intel® 3 Series Chipset: ICH9M, and ICH9ME**
Yes[£]	No	No	No	No	No

[£]Note: This feature is only supported on the designated chipsets' RAID-enabled SKUs.

17.5.1 Overview

17.5.1.1 System Requirements

For a system to support ATA Power-Up in Standby (PUIS) for HDDs with Intel® Smart Connect Technology it must:

- Include one of the following Intel chipsets (RAID enabled SKUs):
  - Intel® 7 Series Chipset SATA RAID Controller; following SKUs only:



- Mobile: QM77, HM77, QS77, or
- Desktop: Q77, Z77, H77, or
- Workstation/Server: C216, and

- Meet the Intel® Smart Connect Technology system requirements (refer to CDI/IBP document # 482930 for all requirements)
- Have the Intel® Rapid Storage Technology driver and OROM installed (the 11.0 or later production versions of each)
- Meet the Intel® Smart Response Technology system requirements

17.5.1.2 Feature Limitations of ATA Power-Up in Standby (PUIS) for HDDs with Intel® Smart Connect Technology

- OS must be installed to a bootable HDD (the targeted PUIS volume) (**Systems configured with only SSDs are not supported**)
- Intel® Smart Response Technology must be enabled in **Maximized** mode
- Operating system:
  - Windows* 7 (32/64 bit)
  - Windows* Server 2008 R2 x64
  - Windows* Server 2008 (32/64 bit) and Data Center

17.5.2 Theory of Operation

17.5.2.1 External Dependencies Required for Activation

This feature enhancement to the Intel® RST driver has two main external dependencies:

1. The presence of the Intel® Smart Connect Technology Agent to initiate activation of the feature in the driver (if no Agent present, the feature remains inactive)
2. OS installed on a PUIS compatible* HDD (once the Agent communicates with the driver to activate the feature, the RST driver checks for the presence of required HW configuration before activating the feature)

Note: Only HDDs that DO NOT require a jumper to emulate the feature are supported.

17.5.2.2 Activation Process

Step 1: Communication is established between Intel® Smart Connect



Technology Agent and the Intel® RST driver where the Agent informs the driver to enable PUIS support for HDDs during Agent triggered power state changes

Step 2: Intel® RST driver checks for the following and if all are present, enables PUIS support per the Agent's request:

- OS system disk (boot volume) is a single pass-through hard disk drive (HDD) AND
- The HDD supports the ATA PUIS feature (jumpered drives are not supported) AND
- The system disk is Accelerated in Maximized mode AND
- The HDD is Writeable

If all of these are true, then the feature enhancement is active in the Intel® RST driver.

17.5.2.3 PUIS Feature Enable Override

Registry key A Windows registry key is provided to disable the PUIS feature

When the key is present and set to '1', the RST driver does the following:

- Disable the PUIS feature on all enumerated HDDs that support the PUIS feature
- Enumerate all HDDs as non-PUIS HDDs
- Return an 'Invalid' status to the Intel® SCT Agent
- Thereafter have no special handling related to the PUIS feature

Sub Key	Value	Type	Default	Description
N/A	DisablePuis	dword	**0**	This key will determine whether or not the PUIS feature is actively disabled on the system. The feature will be enabled if there no key or the key is set to a value of '0'. Default value will be enabled (**0**).

[KEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\]

"DisablePuis"=dword(0, 1)

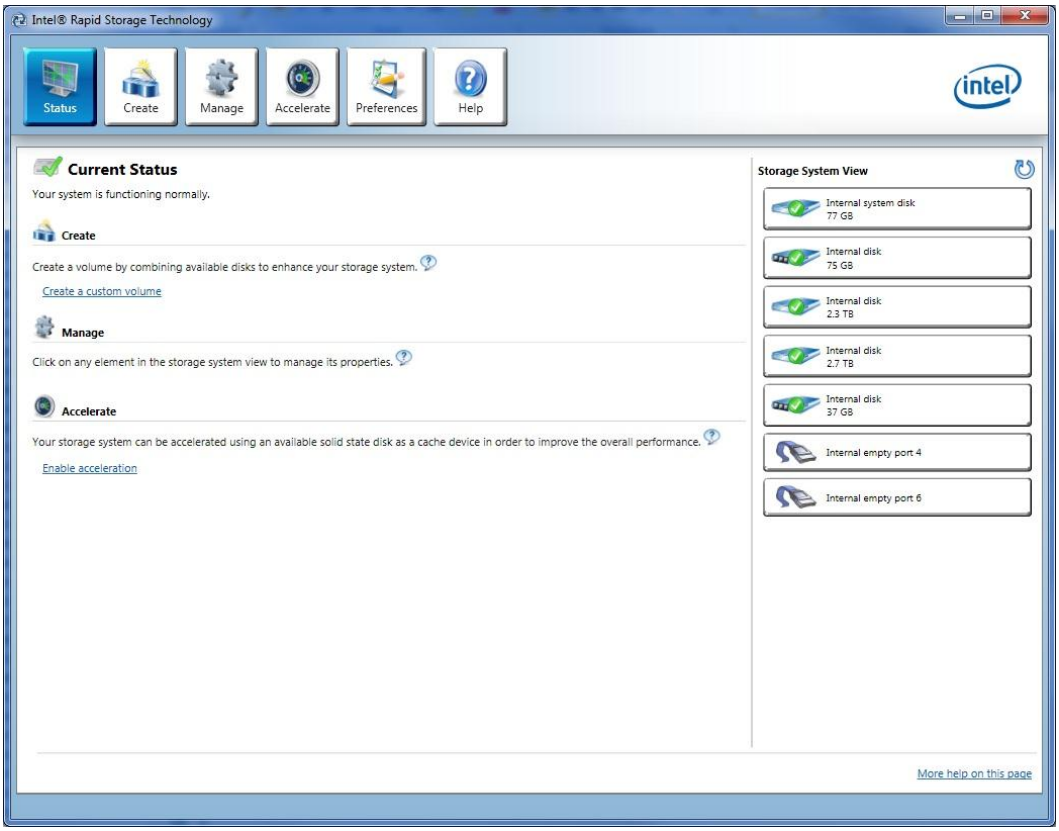


17.5.2.4 Performance

HDD Spin-up Frequency	During the period when the Intel® Smart Connect Technology Agent cycles the system to the S0-ISCT power state to synch the application data, the Intel® RST driver is expected to keep the PUIS HDD in the ATA Power Standby Mode approximately 90% of the synch cycles.

18 Intel® Rapid Storage Technology UI

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



18.1 Introduction

The Intel® Rapid Storage Technology UI is a Windows*-based application that provides users monitoring and management capabilities for the Intel® RST storage subsystem. It offers a wide range of monitoring and management activities for the Intel® RST RAID subsystem (***In AHCI mode there are no management or monitoring capabilities offered by the UI application***).



18.1.1 Getting Started

The Intel® Rapid Storage Technology software package provides high-performance SATA AHCI and SATA RAID capabilities for supported operating systems.

Refer to the System Requirements and the online user's manual to set up your system's configuration and feature support level. You can also review the Readme file installed with this software or visit Intel's online support to learn more about the full system requirements and RAID BIOS configuration.

RAID enabled systems

Redundant Array of Independent Disks (RAID) refers to multiple independent disks combined to form one logical drive. The main objective of this technology is to improve storage system performance, data protection, and increase fault-tolerance.

This technology provides support for the following features:

- **Intel® Rapid Recover Technology**
  This technology provides full data redundancy by copying data from a designated source drive (i.e., master disk) to a designated destination drive (i.e., recovery disk). Data updates of recovery volumes can be continuous or on request.
- **Intel® Rapid Storage Technology RAID**
  This technology provides the ability to create RAID 0, RAID 1, RAID 5, and RAID 10 volumes on desktop and mobile platforms. Data is distributed across two or more disks to provide data redundancy (exception of RAID0) or to enhance data storage performance.
- **Intel® Matrix RAID Technology**
  This technology allows two independent RAID volumes to be created on a single array. The first volume occupies part of the array, leaving space for the second volume. The array may consist of two to six SATA disks depending on the volume types.
- **Hot plug**
  Also referred to as hot swap, this feature allows SATA disks to be removed or inserted while the computer is turned on and the operating system is running. As an example, hot plugging may be used to replace a failed external disk.
- **Intel® Smart Response Technology caching**
  This feature allows you to use a non-system solid state disk and configure it as a non-volatile intelligent cache device in order to accelerate a disk or volume that is part of the storage system. This configuration helps improve the overall system performance.
- **Volume migration**
  This feature provides support for converting system data into a high-performance or protection RAID configuration.
- **Volume size increase**
  This feature allows you to increase the data storage capacity of a volume by using 100% of the available array space or by adding one or more SATA disk to an existing volume.
- **Password-protected disks**
  This feature provides high-level security and protection for the data on your disks with a password, denying access from any unauthorized user.

Additional features and technology supported by the driver although not directly accessible via the Intel® RST UI:



- **TRIM (Microsoft Windows 7* only)**
  This feature provides support for all solid state disks (SSDs) in your storage system that meet the ATA-8 protocol requirements and are not part of an array. This feature optimizes write operations, helps devices reduce wear, and maintains unused storage area on devices as large as possible.
- **ODD power optimization (Microsoft Windows Vista* and higher)**
  This feature allows an unused optical disk drive (ODD) to be automatically powered off when media such as a compact disk, a DVD, or Blu-ray disk are not present in the drive and the tray** is closed. The ODD is powered back on by the operating system or user interaction with the device, including when the eject button is pressed. ODD power optimization is particularly valuable for mobile computers as battery life is negatively affected when the ODD is powered on and in an idle state. This feature is only supported on the following system configurations: Intel® 6 Series Chipset or later, compatible motherboards, and compatible ODDs. For more information about compatibility requirements, refer to the SATA specifications available at www.sata-io.org
  **For slot-loadable drives, the tray condition does not apply.
- **Native command queuing**
  A feature that allows SATA disks to accept more than one command at a time. When used in conjunction with one or more disks that support NCQ, storage performance is increased on random workloads by allowing the disk to internally optimize the order of commands.
- **Disks of more than two terabytes**
  This feature provides support for hard disks and solid state disks with a capacity greater than 2 TB that are reported as pass-through devices (available) or used in a RAID configuration. In addition, booting from a system disk greater than 2 TB is allowed as long as the version of the option ROM in your system supports this feature.

  Note: If a source disk is greater than 2TB and using the MBR partitioning scheme, the application will not allow data preservation in order to create a volume. Instead, a new volume will be created with no partition on it. Also, if the operating system is Windows* XP, capacity expansion operations will not be allowed for volume sizes equal or greater than 2TB.

AHCI-enabled systems

Advanced Host Controller Interface (AHCI) is an interface specification that automatically allows the storage driver to enable advanced SATA features, such as Native Command Queuing and Native Hot Plug, on the SATA disks connected to your computer. The following features are supported on AHCI-enabled systems:

- **Native command queuing**
- **Hot plug**
- **Disks of more than two terabytes (if that size is supported by the option ROM)**
- **Password-protected disks**
- **ODD power optimization (Microsoft Windows Vista* and higher)**



18.1.2 System Requirements

Intel® Rapid Storage Technology provides enhanced management capabilities and detailed status information for Serial ATA AHCI and RAID subsystems. This application is supported on the following operating systems and hardware components:

New platforms/chipsets

- Chipset: Intel® Mobile Express Chipset SATA AHCI/RAID Controller
  - Platform: Huron River
    - **SKU**: QM67, HM67, QS67, UM67[1], HM65[1]
- Chipset: Intel® Desktop/Workstation/Server Express Chipset SATA AHCI/RAID Controller
  - Platform: Sugar Bay (desktop)
    - **SKU**: Q67, Z68, H67, P67, Q65[1], B65[1] **(Note: For the H61 SKU – Intel® RST is not supported)**
  - Platform: Bromolow (workstation/server)
    - **SKU**: C206, C204, C202

Legacy platforms/chipsets

- Intel® 5 Series Chipset-based (PCH) platforms (Ibex Peak)
  - Mobile
    - **SKU**: QM57, HM57, PM55, QS57, HM55[1]
  - Desktop
    - **SKU**: Q57, H57, P55, B55[1], H55[1]
- Intel® ICH7 and newer ICH chipsets (mobile and desktop)[2]
  - Intel® 4 Series Chipset, ICH10R/ICH10DO (Eagle Lake)
  - Intel® ICH9M/ICH9ME (Montevina)
  - Intel® ICH7MR/ICH7R

Supported operating systems

- Microsoft Windows* XP Professional
- Microsoft Windows* XP Professional 64-bit Edition
- Microsoft Windows* XP Home
- Microsoft Windows* XP Media Center Edition 2004
- Microsoft Windows* XP Media Center Edition 2005
- Microsoft Windows* XP 64-bit Edition
- Microsoft Windows* XP
- Microsoft Windows* Vista
- Microsoft Windows* Vista x64-bit Edition
- Microsoft Windows* 7
- Microsoft Windows* 7 x64-bit Edition
- Microsoft Windows* Server 2003
- Microsoft Windows* Server 2003 Enterprise x64-bit Edition
- Microsoft Windows* Server 2003 Standard x64-bit Edition
- Microsoft Windows* Server 2008
- Microsoft Windows* Server 2008 x64-bit Edition
- Microsoft Windows* Server 2008 R2 x64

---

[1] Indicates that the SKU is supported in AHCI mode Only. All other SKUs are support in both RAID and AHCI modes. Note: Intel® RST is not support on any SKUs in IDE mode.

[2] The 10.0 OROM will display the text "ISM Unsupported HW…" if installation attempted on any chipset older than ICH8 (mobile or desktop)



18.1.3 Understanding the Application

The Intel® Rapid Storage Technology application allows you to optimize and maintain a healthy storage system by creating volumes, customizing performance settings and managing storage system elements. This section provides you with a general overview of a storage system configuration and an individual review of all the areas contained in this application.

18.1.3.1 Storage System Configuration

The storage system combines hardware capabilities with RAID technology to provide flexible data storage units on your computer. Each data storage unit, or RAID configuration, consists of three elements that include physical SATA disks, one or two volumes, and one array. When at least one volume is present on the system, these elements are represented in the storage system view of the Status and Manage areas.

In this section, we describe each of these RAID configuration elements and explain how they relate to each other.

- **Array**
  An array is a collection of two or more SATA disks in a RAID configuration and is the highest element in the hierarchy of a storage system. Once a volume is created, the disks you used to create that volume form an array. Refer to the Creating Additional Volumes topic for details on how you can create two volumes across the same disks. An array can include one or two RAID volumes if the hardware allows it.
- **Volume**
  A volume is the storage area on two or more disks whose type dictates the configuration of the data stored. If you created a volume for data protection, then your storage system may include a RAID 1 volume spanning two SATA disks, which mirrors data on each disk.
- **Disks**
  A disk (i.e., hard disk or hard disk drive) physically stores data and allows read/write data access. If a disk is used to create a volume, it becomes an array disk because it has been grouped with other disks to form an array.

The storage system can also include ATAPI devices, which cannot be used to create a volume. They are a mass storage device with a parallel interface, such as CD-ROM, DVD/Blu-ray disc, or tape drive.

18.1.3.2 Navigation

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

**Status**

The 'Status' area provides a general state of health of your storage system. If a status other than normal is reported, the Manage sub-section will be available to provide you with basic information and actions links necessary to return the status to normal.



Create

The 'Create' area allows you to create different types of volumes to protect data, enhance disk performance, optimize disk capacity, or create a custom volume to combine benefits.



Note

The ‘Create’ area is only available if your computer supports RAID technology, and if the volume requirements are met. Refer to the Volume Requirements topic for an exhaustive list of storage system conditions to create a volume.



Manage

The 'Manage' area combines the logical and physical view of your storage system. The area displays detailed information about each element that is part of the storage system, such as volumes and disks; the storage system view shows how the selected element relates to others. Each element has its own 'Manage' area which is accessible by clicking any element displayed in the storage system view under 'Status' or 'Manage'.

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



Accelerate

The ‘Accelerate’ area allows you to manage the cache memory configuration using a non-system solid state disk as a cache device. If the cache is reported in an abnormal state, detailed information and troubleshooting actions will display. The Acceleration View is specific to the ‘Accelerate’ area and only displays in this location.



Preferences

The 'Preferences' area allows you to customize system settings by enabling the display of the notification area icon, and by selecting the type of notifications that you want the application to display.



Storage System View

The storage system view has two functions:

- It is a simplified representation of your storage system and displays graphic elements, such as arrays, volumes, devices, and ports. Each element provides general attribute information, such as status, name and size. Hovering over each element provides additional attribute details.
- You can also use the graphical view to access 'Manage' by clicking the storage system element you want to work with. For example, if an array is present, clicking the volume opens Manage Volume and clicking one of the array disks will open Manage Disk for the selected disk.



Acceleration View

The Acceleration View is a graphical representation of the acceleration configuration, and only displays the devices (disks and volumes) included in this

particular configuration. You can use this view to access the 'Manage' page specific to each represented device by clicking the storage system element for which you want more detailed information.

18.1.4 Notification Area

The notification area (also called the system tray) is located on your desktop. The taskbar contains the notification area icon for Intel® Rapid Storage Technology. The icon provides storage system status and notifications such as volume and disk events based on a change of state.

The notification area icon will automatically display in the notification area once Intel Rapid Storage Technology is installed. Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Settings changes are applied on a per user basis, and do not affect other users' settings.

Opening the application from the notification area

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

The notification area icon can be in the following states:

Icon	Description
	The storage system is reported in a **normal** state and your data is protected from a disk failure.
	The storage system is reported in a **warning** state and data may be at risk. We recommend that you open the application to review and resolve the reported issues.
	The storage system is reported in an **error** state and data may be lost. We recommend that you open the application to review and resolve the reported issues as soon as possible.
	The storage system is reported in a **busy** state while an operation is in progress. Once the operation is complete, all actions will be available again, allowing you to manage the storage system as long as it is reported in a normal state. You can follow the progress of the operation by hovering over the icon.
	This icon is displayed while you are attempting to open the application, but the Intel® Rapid Storage Technology service has not started running yet. The service is expected to start automatically with a delay when you launch Windows. This icon appears if you attempt to launch the application before the delay period ends. If the application fails to open, try starting the service manually using Microsoft Windows* Services.

Selecting system notifications

1. Right-click the icon.
2. Select the types of notifications you want to receive. The notification area menu allows you to select or deselect one option at a time. Repeat this procedure until you are finished with your selection. The same operation can also be completed using the application, from the 'Preferences' area.


Note

To hide the notification area icon, deselect ‘Show the notification area icon’ under ‘System Preferences’.

Reviewing notifications

- Hover over the icon at any time to view the storage system status or the progression of an operation.
- Small pop-up windows will display for a short time to notify you of specific events, such as a missing disk or the completion of an operation.
- Open the application to view more details about storage system events in the 'Status' or 'Manage' areas.

18.2 Storage System Status

Anytime Intel® Rapid Storage Technology is launched, the application opens to the 'Status' area. This is where the general state of health of your storage system is reported, both in the storage system view and in details. Depending on the status, volume creation and management options may be available in order to enhance or repair your storage system.

18.2.1 Understanding the Status

To get the full benefits of what Intel® Rapid Storage Technology has to offer, it is critical to maintain a healthy storage system. The application helps you track and reports any disk or volume related problems that could put the safekeeping of your data at risk.

The storage system can be in the following states:


Normal

Reports that the system is functioning as expected, SATA disks are present and connected to the computer. If an array is present, volume data is fully accessible.


The Create subsection is only available if the storage system meets the minimum requirements to create a volume. Depending on the available hardware, you may be given the option to create a volume to protect data, optimize the disk performance, or create a custom volume.


The Manage subsection is only available if the storage system reports atypical conditions in a normal state. Typically, details or a recommended action are provided to help you rectify any storage system conditions. For example, if a recovery volume was reported as read-only, we would inform you that disk files must be hidden prior to requesting updates.


The Accelerate subsection is only available if a solid state disk can be used as a cache device and an eligible disk or volume can be accelerated. This area typically provides the option to enable acceleration and reports the cache and accelerated device health state, as well as the current acceleration mode.


Warning

Reports that storage system data may be at risk due to a problem detected on one or more SATA disks.


The Manage subsection displays any SATA disk or volume states reported by the storage system that may require your attention in order to keep data fully protected and accessible. Details or a recommended action

are provided to help you fix any storage system problems. For example, if the master disk in a recovery volume is reported as failed, we would recommend that you rebuild the volume to another disk.

Note

In this state, we recommend that you backup any accessible data before taking action

In this state, the Accelerate subsection typically reports that the cache volume is failing possibly because the solid state disk is reported at risk of failing (smart event). Details and a recommended action are provided to help you fix the problem reported on the solid state disk.

Error

Reports that storage system data may be lost due to a problem detected on one or more SATA disks.

The Manage subsection displays any SATA disk or volume states reported by the storage system that require your immediate attention in order to keep data fully protected and accessible. Details or a recommended action are provided to help you fix any storage system problems. For example, if the data on a RAID 1 volume appears inaccessible due to a failed array disk, we would recommend that you rebuild the volume to another disk.

Note

In this state, we recommend that you backup any accessible data before taking action

In this state, the Accelerate subsection typically reports that the cache volume has failed possibly because the solid state disk has also failed and there is data loss. Details and a recommended action are provided to help you fix the problem reported on the solid state disk.

18.2.2 Storage System View

The storage system view provides a visual representation of your storage system and displays arrays, volumes, devices, and ports. Volumes and SATA disks graphics reflect their current states, which allows you to quickly identify the element that is causing the storage system to be in a state other than normal.

Note:
Hovering over a designated element in the storage system view provides a snapshot of its properties. Clicking allows you to access and manage its properties.

Overview of SATA disks states

State	Description	Recommendation
	An internal hard disk is reported normal.	None
	An external hard disk is reported normal.	None
	An internal solid state disk is reported as normal.	None
	An external solid state disk is reported as normal.	None
	An internal disk is reported missing.	Ensure that the disk is securely connected to the SATA port and that the SATA cable is functioning properly. Refer to the Troubleshooting section for more information.
	An internal disk is reported at risk or Incompatible.	Back up your data and replace the disks as soon as possible. Refer to the

		Troubleshooting section for more information.
	An external hard disk is reported at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.
	An internal solid state disk is reported as being at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.
	An external solid state disk is reported at risk or incompatible.	Back up your data and refer to the Troubleshooting section for more information.
	An internal disk is reported offline.	Unlock all array disks to unlock the volume. Refer to the Troubleshooting section for more information.
	An internal recovery disk is reported offline.	• The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or, • Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.
	An external disk is reported offline.	Unlock all array disks to unlock the volume. Refer to the Troubleshooting section for more information.
	An external recovery disk is reported offline.	• The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or, • Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.
	An internal disk is reported normal and locked.	Unlock the disk to access more options.
	An external disk is reported normal and locked.	Unlock the disk to access more options.
	An internal hard disk is reported failed.	Refer to the Troubleshooting section for more information.
	An external hard disk is reported failed.	Refer to the Troubleshooting section for more information.
	An internal solid state disk is reported as failed.	Refer to the Troubleshooting section for more information.
	An external solid state disk is reported as failed.	Refer to the Troubleshooting section for more information.

Volume states

Volume type	Normal	Degraded	Failed
		Refer to Troubleshooting Degraded Volumes and	Refer to Troubleshooting Failed Volumes and Caching Issues for

		Caching Issues for more information.	more information.
RAID 0		Not applicable	
Single-disk (cache)			
Single-disk (data)		Not applicable	
RAID 1/Recovery			
RAID 5			
RAID 10			

Other storage system elements

Element	Description	Recommendation
	A port that has no devices connected to it.	None
	An ATAPI device is present, such as CD-ROM, DVD/Blu-ray disc, or tape drive.	None

18.3 Creating a Volume

You can combine SATA disks to create a volume in order to enhance your storage system. Based on the available hardware and your computer's configuration, you may be able to create a volume by selecting an enhancement goal, such as 'Protect data' under 'Status', or by selecting a volume type under 'Create'. We recommend you get familiar with the minimum requirements in this section before starting the volume creation process.

Warning
Performing this action will permanently delete any existing data on the disks used to create a volume, unless you choose to keep the data when selecting array disks. Backup all valuable data before starting this process.

18.3.1 Volume Requirements

Creating a volume is only available as an option if the following requirements are met:

- You are logged on as an administrator.
- The computer is RAID ready (refer to the user's manual available on Intel's online support web site, for assistance on setting up a RAID ready system).
- Two or more SATA disks, including the operating system disk are connected, in a normal state, and unlocked (only applies to password-protected disks).

Enabling more disks
When configuring a volume, the application will only list the SATA disks that meet the requirements listed below. For example, a locked disk connected to your computer will not be listed as an option until it is unlocked.



Based on the first disk selected, some disks may become grayed out if one or more requirements are not met. Selecting a different disk generally helps re-enable disks that were previously grayed out.

- If the first selection is a system disk, any additional SATA disks selected must be of equal or greater size to ensure that all the system files are migrated to the new volume.
- If the first selection is a non-system disk, and a system disk is then selected, the latter must be of equal or smaller size to ensure that all the system files are migrated to the new volume.
- A system volume cannot be greater than 2 TB. If your first selection is a system disk, the total size of the other disks shall not allow the volume size to exceed 2 TB. Exception: If you are creating a volume using disks that have no existing data, and your operating system is XP 64-bit Edition, the application will allow a volume to be greater than 2TB.
- The SATA disks used to create a volume must have the same type of connection, internal or external. An internal disk shall not be paired with an external disk to create a volume. Some systems will support mixed connection types.

Enabling more volume types

Depending on the input/output (I/O) controller hub that your computer is using and the hardware connected to the system, some volume types may not be enabled in the selection list. Refer to the Readme file located in the Program Files directory for this application or to the Device Manager to determine which controller is installed on your computer. Review the controller support table below to determine which volume types you can create.

Note
Intel® 5 Series Chipset applies to both desktop and mobile platforms

Volume type	Number of disks	Controller support
Recovery volume	2	ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series. **Note** No other volumes can be present on the system. The master disk must include 100% of the available disk space and must be less than 1.3125 TB
RAID 0	2	ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 0	3 or 4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series .
RAID 0	5 or 6	ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 1	2	ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 5	3 or 4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.
RAID 5	5 or 6	ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.

RAID 10	4	ICH7R, ICH7DH, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, 5 Series/3400 Series.

18.3.2 Creation Process

Now that you have reviewed the volume requirements, this section will guide you through the three easy steps necessary to create a volume.

18.3.2.1 Selecting a Volume Type

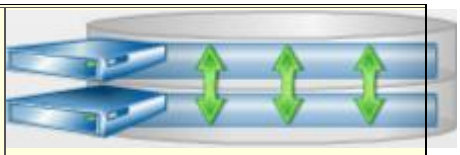
Before you can create a volume, you need to decide how you want to enhance your storage system based on your needs. Depending on the available hardware, you may have the option to combine volume types by creating more than one volume on a single array. Refer to 'Creating Additional Volumes' for more information on this type of configuration. Below is an overview of the five volume types that you can create.

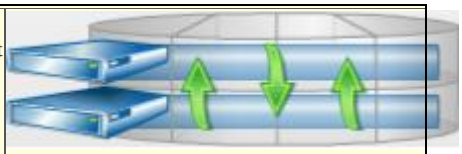
Creating a two-disk volume from 'Status'
This option displays if only two disks are available, one has data such as system files, the second one doesn't, and the latter has a size that is equal or greater than the other. Based on this simple configuration, you can create a volume to protect data or optimize disk performance by clicking one of the two options listed in the Create subsection. When choosing this option, the application automatically configures the volume using the only two disks available and assigns default settings. Refer to the applicable procedure described in Completing the Process for details.

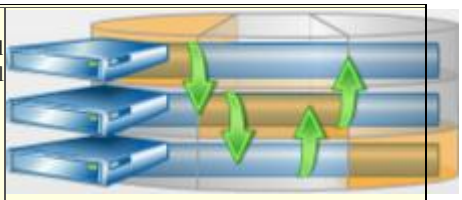
Creating a custom volume

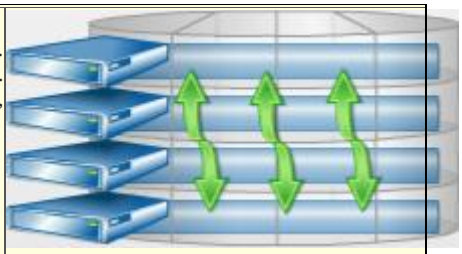
1. Click 'Create' or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click 'Next'.

Recovery volume: **Flexible data protection** Combines two SATA disks and utilizes RAID 1 functionality to copy data from a designated master disk to a designated recovery disk. Data updates on the volume can be continuous or on request. In 'on request' mode, you can request data updates that copy changes from the master disk to the recovery disk since the last update. No other volumes can be present on the system. The master disk must include 100% of the available disk space and must be less than 1.3125 TB.	
Disks required	2
Advantage	Full data redundancy; more control over how data is copied between master and recovery disks; fast volume updates in on request mode; master and recovery disk files can be viewed in Windows Explorer*.
Disadvantage	Storage capacity is only as large as the smallest disk.
Application	Critical data protection for mobile systems; fast restoration of the master disk to a previous or default state. Available in specific mobile configurations.

RAID 1: **Real-time data protection** Combines two SATA disks where each stores an exact copy of the data to appear as a mirror of each other.	
Disks required	2
Advantage	Full data redundancy and excellent fault-tolerance; increased read transfer rate.
Disadvantage	Storage capacity is only as large as the smallest disk; slight decrease in write transfer rate.
Application	Typically used in workstations and servers to store critical data. Available in specific mobile configurations.

RAID 0: **Optimized disk performance** Combines two to six SATA disks and breaks down data into units that are spread across the array disks.	
Disks required	2 to 6
Advantage	Increased data access and storage performance; no loss in data capacity
Disadvantage	No data redundancy (if one disk fails, all data on the volume is lost).
Application	Typically used in desktops and workstations to store high performance, temporary data and software. Various RAID 0 volume configurations available in specific mobile configurations.

RAID 5: **Efficient data hosting and protection** Combines three to six SATA disks where data and parity are striped across the array disks in a rotating sequence. Parity is a mathematical method for recreating lost data to a single disk.	
Disks required	3 to 6
Advantage	Data redundancy; improved storage performance and capacity; high fault-tolerance and read performance.
Disadvantage	Time-consuming to rebuild and decreased performance during the process.
Application	Good choice for large amounts of critical data, such as file and application servers; Internet and Intranet servers. Available in mobile configurations that include the Intel® 5 Series Chipset which supports up to six SATA ports.

RAID 10 : **Balanced performance and data protection** Combines four SATA disks to create a combination of RAID types 1+0. The data is striped across a two-disk array forming a RAID 0 component. Each disk in the RAID 0 array is mirrored by a disk in the RAID 1 array, storing an exact copy of all the data.	
Disks required	4
Advantage	Combines the read performance of RAID 0 with the fault-tolerance of RAID 1, resulting in increased data access and full data redundancy, and increased storage capacity.
Disadvantage	4 disks are required, resulting in increased cost.
Application	High performance applications and high load database servers requiring data protection, such as video editing. Available in mobile configurations that include the Intel® 5 Series Chipset which supports up to six SATA ports.

18.3.2.2 Configuring the Volume

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

Recovery volume

1. Type a new volume name if you want to change the default name.
2. Select the master disk.
3. Select the recovery disk.
4. Select a different update mode, if desired.
5. Click 'Next'. This button will not be active until all the required selections have been made.

Advanced configuration settings:

- Enable or disable the volume write-back cache.
- Select the check box to initialize the volume. You can choose to perform this action at a later time.

RAID Volume

1. Type a new volume name if you want to change the default name.
2. Select the required number of disks.
3. Select the disk from which you want to keep data, if desired. You can only keep data from one disk. If you want to keep data from more than one disk, you must back up all valuable data prior to creating a volume.
4. Click 'Next'. This button will not be active until all the required selections have been made.

Advanced configuration settings:

- Select the array allocation by using the slider.
- Select a data strip size.
- Enable or disable the volume write-back cache.

- Select the check box to initialize the volume. You can choose to perform this action at a later time.

Note
Currently, the application does not allow the creation of greater than 2TB volumes where the source disk is greater than 2TB and data on that disk is preserved (e.g. system volume). Target disks can be greater than 2TB but such volumes cannot. This limitation results from the lack of GPT partition scheme support. Note that volumes greater than 2TB that include member disks greater than 2TB are supported as long as array disks are unpartitioned or that no data is preserved at volume creation.

18.3.2.3 Completing the Process

If you are creating a custom volume, and have configured the volume with the disk selection and other settings, you are ready to review the projected configuration and complete the volume creation process.

If you are creating a two-disk volume for data protection or disk optimization from 'Status', you can follow the procedure provided below.

Creating a two-disk volume from 'Status'

1. Under 'Status', in the Create sub-section, select the type of volume you want to create.
2. In the 'Confirm Volume Creation' dialog, review the volume configuration. Note that the volume name is the only setting that can be changed.
3. Click 'Create Volume' to confirm. The process starts immediately.
4. Once completed, a dialog displays to notify you that the volume was successfully created. Click 'OK' to close the dialog.
5. The page refreshes and displays the new volume in the storage system view as well as the data migration progress.

Creating a custom volume

Warning
You can only keep existing data from one of the disks you select to create a volume. We recommend that you backup all valuable data before proceeding.

If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

1. Under 'Confirm', review the selected configuration.
2. Click 'Create Volume' if you want to create the volume using the selected configuration. Otherwise, click 'Back' and make any necessary changes. The process starts immediately.
3. Click 'OK' to confirm.
4. Once completed, a dialog displays to notify you that the volume was successfully created. Click 'OK' to close the dialog.



5. The ‘Status’ area displays the new volume in the storage system view as well as the data migration progress.

If the size of the new volume is larger than the size of the source drive, the following steps apply:

6. Once the migration status reports 100% complete, restart your computer for the operating system to recognize the new volume size.
7. Create a new partition or extend the existing partition to utilize the new volume space using Windows Disk Management*. If your system is running Microsoft XP*, you may only have the option to create a new partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management.

18.3.3 Creating Additional Volumes

Creating multiple volumes on a single array

You can add a volume to an existing RAID array by creating another volume that uses the available space on the array. This feature allows you to combine different volume types and their respective benefits. For example, a configuration with RAID 0 and RAID 1 on two SATA disks provides better data protection than a single RAID 0 and higher performance than a single RAID 1.

The first RAID volume occupies part of the array, leaving space for the other volume to be created. After creating the first volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to add a second volume to that array.

Note
The configuration is only available if the array allocation for the first volume created is less than 100%, and space is available on that array. The application currently supports an array to include a maximum of two RAID volumes.

1. Click ‘Create’ or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click ‘Next’.
4. Select 'Yes' to add the volume to an existing array.
5. Make any necessary changes in the Advanced section.
6. Click ‘Next’.
7. Review the selected configuration. Click 'Back' or an option in the left pane if you want to make changes.
8. Click 'Finish' to start the creation process.

Supported RAID volume combinations on a single array:

Combine	With
2-disk RAID 0	2-disk RAID 0

	2-disk RAID 1
2-disk RAID 1	2-disk RAID 0 2-disk RAID 1
3-disk RAID 0	3-disk RAID 0 3-disk RAID 5
3-disk RAID 5	3-disk RAID 0 3-disk RAID 5
4-disk RAID 0	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
4-disk RAID 5	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
4-disk RAID 10	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
5-disk RAID 0	5-disk RAID 0 5-disk RAID 5
6-disk RAID 0	6-disk RAID 0 6-disk RAID 5

Visit our Online Support for additional information on RAID type combinations for each I/O controller hub.

Creating additional volumes on a new array

You can choose to create two or more volumes on two different arrays, as long as the volume requirements are met.

1. Click ‘Create’ or 'Create a custom volume' under 'Status'.
2. Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.
3. Click ‘Next’.
4. Select 'No' in order to add a volume to a new array.
5. Select the required number of disks.
6. Select the disk from which you want to keep data, if desired. You can only keep data from one disk. If you want to keep data from more than one disk, you must back up all valuable data prior to creating a volume.
7. Make any necessary changes in the Advanced section.
8. Review the selected configuration. Click 'Back' or an option in the left pane if you want to make changes.
9. Click ‘Next’.
10. Click 'Finish' to start the creation process.

18.4 Managing the Storage System

The 'Manage' area combines the logical and physical view of your storage system. The area displays detailed information about each element that is part of the storage system, such as volumes and disks; the storage system view shows how the selected



element relates to others. Each element has its own 'Manage' area which is accessible by clicking any element displayed in the storage system view under 'Status' or 'Manage'.

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

18.4.1 Managing Arrays

You must be logged on as an administrator to perform the actions listed in this section.

You can manage arrays by clicking a selected array in the storage system view under 'Status' or 'Manage'. This allows you to review the properties and access all actions associated with that array, such as adding a disk or increasing a volume size.

18.4.1.1 Array Properties

An array is a logical grouping of physical SATA disks. The array properties listed below display to the left of the storage system view under Manage Array and report values specific to the element selected in the view.

Parameter	Value
Name	Reports the name of the array. The array name is automatically assigned and cannot be changed.
Size	Reports the total capacity of the array in megabytes (MB).
Available space	Reports the unallocated space on the array that can be used.
Disk data cache	Reports whether the data cache is enabled for all array disks.

18.4.1.2 Adding a Disk to an Array

You can add one or more SATA disks to an existing array to increase the system storage capacity. This feature can be useful if you want to change to a volume type that requires additional disks.

This option is only available if:

- A RAID 0 and/or a RAID 5 volume is present,
- One or more SATA disks are connected to the computer and available,
- The available disk matches the internal or external connection type of the existing array disks. You cannot add an external disk to an array that includes internal disks, and vice versa.

Refer to Connecting a Disk under Managing Disks for more information on installing SATA disks on your computer.

Warning

Any existing data on the available disk used to increase the array size will be permanently deleted. Backup all the data you want to preserve prior to executing this action.

If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize



the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

This action can also be performed from Manage Volume. Refer to the Adding a Disk to a Volume section for more information.

1. Under 'Status' or 'Manage', in the storage system view, click the array to which you want to add a disk. The element properties are now displayed on the left.
2. Click 'Add disk'.
3. Select the disk you want to use to increase the array capacity.
4. Click 'Add Disk'. Caution: Once the data migration starts, the operation cannot be canceled.
5. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes for which a disk was added, or add another partition.

**Note**
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

18.4.1.3 Adding a Volume

You can add a volume to an existing RAID array by combining different volume types and their respective benefits. For example, a configuration with RAID 0 and RAID 1 on two SATA disks provides better data protection than a single RAID 0 and higher performance than a single RAID 1.

The first RAID volume occupies part of the array, leaving space for the other volume to be created. After creating the first volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to add a second volume to that array.

**Note**
This configuration is only available if the array allocation for the first volume is less than 100%, and space is available on that array. The application currently supports an array to include a maximum of two RAID volumes on a single array.

You can also complete this action using the 'Create' area.

1. Under 'Status' or 'Manage', in the storage system view, click the array to which you want to add a volume. The array properties are now displayed on the left.
2. Click 'Create additional volume'.
3. In the 'Create Additional Volume' dialog, type a new name if you want to change the default name.



4. Select the volume type, and then click 'OK'. Only the volume types available for the current configuration will display. Refer to the table below for more information.
5. The page refreshes and the array now displays the additional volume.

Supported RAID volume combinations on a single array:

Combine	With
2-disk RAID 0	2-disk RAID 0 2-disk RAID 1
2-disk RAID 1	2-disk RAID 0 2-disk RAID 1
3-disk RAID 0	3-disk RAID 0 3-disk RAID 5
3-disk RAID 5	3-disk RAID 0 3-disk RAID 5
4-disk RAID 0	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
4-disk RAID 5	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
4-disk RAID 10	4-disk RAID 0 4-disk RAID 5 4-disk RAID 10
5-disk RAID 0	5-disk RAID 0 5-disk RAID 5
6-disk RAID 0	6-disk RAID 0 6-disk RAID 5

Visit our Online Support for additional information on RAID type combinations for each I/O controller hub.

18.4.1.4 Increasing Volume Size

You can increase the size of a RAID volume by using the remaining available space on the array. A minimum of 32 MB must be available for this action to be available. Hovering over the array name in the storage system view displays the amount of available space in MB.

After creating a volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to increase the volume size by the amount of available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size.

This option is only available if:

- A RAID 0, RAID 1, RAID 5 and/or RAID 10 volume is present,
- The array allocation for the volume is less than 100% and space is available on the existing array.



Warning
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

Increasing the volume size from Manage Array

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The array properties are now displayed on the left.
2. Click 'Increase size' next to the volume name. If more than one volume is present on a single array, you will need to increase the size of each volume one at a time.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Increasing the volume size from Manage Volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the left.
2. Click 'Increase size'.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

18.4.1.5 Enabling Disk Data Cache

Enabling the disk data cache for all SATA disks on the array allows you to enable cache memory physically present on the disks and use it to speed up data access. This action is only available from Manage Array because the data cache must be in the same state across all disks that are part of a single array.

Under Manage Array, the disk data cache is reported as enabled or disabled for all SATA disks in the array. Under Manage Disk, the disk data cache is reported as enabled or disabled for a specific disk that is part of that array. The option to change this setting is only available from Manage Array.

Warning



Enabling the disk data cache increases the cache size and the amount of cached data that could be lost in the event of a power failure. The risk can be decreased if your computer is connected to an uninterruptable power supply (UPS).

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The element properties are now displayed on the left.
2. In the Advanced section, click 'Enable' or 'Disable' depending on the option available.
3. Click 'Yes' to confirm.
4. The page refreshes and now displays the new setting.

18.4.2 Managing Volumes

You must be logged on as an administrator to perform the actions listed in this section.

You can manage existing volumes by clicking a volume in the storage system view under 'Status' or 'Manage'. This allows you to review the volume properties and access all actions associated with that volume, such as renaming, changing type, and deleting.

18.4.2.1 Volume Properties

A volume is an area of storage on one or more SATA disks used within a RAID array. A volume is formatted by using a file system and has a drive letter assigned to it. The volume properties listed below display to the left of the storage system view under 'Manage' and report values specific to the element selected in the view.

RAID volume status table

Status	Description
Normal	Indicates that volume data is fully accessible.
Locked	Indicates that at least one array disk is locked with a password. The volume is visible because at least one other array disk is unlocked. Refer to Unlocking Password-Protected Disks for instructions on unlocking disks.
Degraded	Indicates that one array disk is missing or has failed. A RAID 0 volume cannot be in this state because of the striping configuration.
Failed	• RAID 0 volume: indicates that one or more array disks are missing or have failed. • RAID 1 volume: indicates that both array disks are missing or have failed. • RAID 5 or 10 volume: indicates that two or more array disks are missing or have failed.
Incompatible	Indicates that the volume was moved to another system that does not support the volume type and configuration.
Inaccessible	Indicates that data on the accelerated volume cannot be accessed because it is missing, or that the accelerated volume data is not synchronized with the data on the

	cache volume.
Unknown	Indicates that an unknown error was detected.

Recovery volume status table

Status	Description
Normal	Indicates that volume data is fully accessible.
Locked	Indicates that at least one array disk is locked with a password. The volume is visible because at least one other array disk is unlocked. Refer to Unlocking Password-Protected Disks for instructions on unlocking disks.
Degraded	• The recovery disk has failed, or • The master disk is missing or has failed and the volume is running off the recovery disk.
Failed	Indicates that both array disks have failed.
Incompatible	Indicates that the volume was moved to another system that does not support the volume type and configuration.
Unknown	Indicates that an unknown error was detected.
Power-saving mode	Indicates that the computer is running on battery power. If the volume is in continuous update mode, data updates are paused and will resume as soon as the computer is reconnected to the power supply.
Data update needed	Indicates that the recovery disk does not have a redundant copy of the data on the master disk, and you should request an update.
Running off recovery disk	Indicates that the recovery disk is the designated source drive in the volume.
Master disk read-only	Indicates that the recovery disk is the designated source drive in the volume, and that the master disk files are accessed. In this state, data recoveries from the recovery disk are not available.
Recovery disk read-only	Indicates that the recovery disk files are accessed. In this state, data updates are not available.

Busy volume states table

Status	Description
Initializing	Indicates that data on a volume is being synchronized. This step is required prior to verifying or verifying and repairing data on a volume.
Verifying	Indicates that the volume is being scanned to detect data inconsistencies.
Verifying and repairing	Indicates that the volume is being scanned to detect data inconsistencies, and errors are being repaired. This state does not apply to a RAID 0 volume because errors cannot be repaired.
Migrating data	Indicates that data is being reorganized on the volume. This state displays when a system volume is created, the volume size is increased, or the type is changed to different RAID configuration.
Rebuilding	Indicates that data redundancy is being restored across all disks associated with the volume. A RAID 0 volume cannot be in this state because of the striping configuration.
Recovering data	Indicates that data on the master disk is being overridden by all the data on the

	recovery disk. This state only applies to recovery volumes.
Updating data	Indicates that the latest master disk changes are being copied to the recovery disk. This state only applies to recovery volumes.

General parameters table

Parameter	Value
Details	Provides detailed information if a volume is in a state other than normal.
Type	Reports the volume type.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device.
Size	Enhanced: Indicates that the disk or volume is accelerated for optimized data protection.
	Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance.
	None: Indicates that no disk or volume is accelerated.
	Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.
	Reports the total capacity of the volume in gigabytes (GB) in the storage system view and in megabytes (MB) in the volume properties under Manage Volume.
Data stripe size	Reports the size of each logical contiguous data block used in the volume for RAID 0, 5, and 10 volumes. The strip size is indicated in kilobytes (KB).
Write-back cache	Reports whether the write-back cache feature is enabled for the volume.
System volume	Reports whether the volume contains system files that are required to start and run the operating system.
Initialized	Reports whether the volume is initialized.
Verification errors found	Reports the number of inconsistencies found during the last volume data verification.
Block with media errors	Reports the number of blocks with media errors found during the last volume data verification.
Physical sector size	Reports the size of each sector that is physically located on the disk.
Logical sector size	Reports the size of data collection blocks.
Details	Provides detailed information if a volume is in a state other than normal.

18.4.2.2 Renaming a Volume

You can change the name assigned to a volume present in your storage system at any time. The name change will take effect immediately.

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to rename. The volume properties are now displayed on the left.
2. Click 'Rename'.
3. Type a new volume name, and then click 'OK'.

Note
Volume names are limited to 16 English alphanumeric and special characters including spaces, but cannot include a backslash "\".

18.4.2.3 Rebuilding a Volume

When a volume is reported as degraded because of a failed or missing disk, the disk must be replaced or reconnected and the volume be rebuilt in order to maintain fault-tolerance. The option to rebuild is only available when a compatible disk is connected, available and normal. If a spare disk is available, the rebuild process will start automatically when a disk fails or is missing. For RAID 0 volumes, the rebuild process will start automatically only when one of its members is reported as at risk.

Warning
Completing this action will permanently delete existing data on the new disk and make any other volume on the array inaccessible. We recommend you backup valuable before continuing.

Rebuilding from 'Status' (manually)

1. Verify that the volume is reported as degraded in the Manage subsection. If you have more than one volume listed in this section, you will need to fix the issues reported one at a time.
2. Click 'Rebuild to another disk' next to the volume you want to rebuild.
3. In the Rebuild Volume dialog, select the disk that will replace the failed disk. Only compatible disks in a normal state will be displayed. Refer to Volume Requirements for more information.
4. Click 'OK' to confirm.
5. The volume starts rebuilding and the page refreshes displaying the progress of the operation. You can use other applications during this time and you will be notified when the process has successfully completed.

Rebuilding from 'Manage' (manually)

1. Under 'Status' or 'Manage', in the storage system view, click the volume you want to rebuild. The element properties are now displayed on the left.
2. Click 'Rebuild to another disk', and then follow the procedure described above.

18.4.2.4 Recovering Data

Recovering data to the master disk allows you to maintain full data redundancy on the recovery volume and keep the volume data healthy. This action is only available if a recovery volume is present and running off the recovery disk.

You may have to recover data if:

- Data on the recovery and master disk is not synchronized and full data redundancy is at risk.



- Data on the master disk is invalid or inaccessible.

Warning
Completing the action will override existing data on the master disk and update it with the data on the recovery disk. Backup all valuable data before continuing.

1. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'.
2. Click 'Yes' to confirm.
3. The recovery operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.

Note
If master disk is removed while the data recovery is in progress and is then reconnected, the operation will resume automatically from where it stopped as long as the volume is in on request update mode. If the volume is in continuous update mode, you will need to restart the operation by following the procedure described above,

18.4.2.5 Resetting Volume to Normal

This action is only available when a volume is reported as failed, but both array disks are present and normal, and allows you to access and try recovering healthy volume data.

In most cases, this situation will occur after one or more array disks was reported as failed or at risk, and then reset to normal.

Completing this action resets the volume state by ignoring previous events and does not repair data. Any data loss or corruption that may have occurred as a result of prior hardware failure or change of state remains. We recommend that you back up accessible data and replace failed hardware as soon as possible to prevent further data loss.

1. Under 'Status', in the Manage subsection, click 'Reset volume to normal'. You can also perform this action from Manage Volume, which is accessible by clicking the RAID 0 volume in the storage system view.
2. Click 'Yes' to confirm.
3. The page refreshes and the volume displays as normal. If the operation failed to return the volume to a healthy state, visit Intel's online support web site for more options.

18.4.2.6 Changing Volume Type

You can choose to change the type of an existing volume based on your storage system needs. The following configurations are possible:

Change type from	To
2-disk recovery volume	2-disk RAID 1

Note Only available if the recovery volume is in continuous update mode	
2-disk RAID 1	2-disk recovery volume **Note** No other volumes can be present on the system. The RAID 1 volume must be less than 1.3125 TB and include 100% of the available space on the array
2-disk RAID 1	2-disk RAID 0 3, 4, 5 or 6-disk RAID 0 3, 4, 5 or 6-disk RAID 5
2-disk RAID 0	3, 4, 5 or 6-disk RAID 5
3-disk RAID 0	4, 5 or 6-disk RAID 5
4-disk RAID 0	5 or 6-disk RAID 5
4-disk RAID 10	4, 5 or 6-disk RAID 5

Note
Before starting, refer to the system and volume requirements to determine which RAID types are supported by your computer and make sure the required number of SATA disks are connected. The Intel® Chipset provides support for the creation of all RAID volume types and for up to six SATA ports on a mobile platform. Changing volume type does not require re-installation of the operating system

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to modify. The volume properties are now displayed on the left.
2. Click 'Change type'.
3. In the 'Change Volume Type' dialog, type a new name if you want to change the default name.
4. Select the new volume type, and then click 'OK'. Caution: Once the data migration starts, the operation cannot be canceled.
5. Once the migration has completed, the 'Manage' page refreshes and reports the new volume type.

Warning
All applications and existing volume data remain intact, but any existing data on the disks added to enable this operation will be permanently deleted. Backup data before adding these disks.

If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.



18.4.2.7 Increasing Volume Size

You can increase the size of a RAID volume by using the remaining available space on the array. A minimum of 32 MB must be available for this action to be available. Hovering over the array name in the storage system view displays the amount of available space in MB.

After creating a volume with an array allocation set to less than 100% in the Configure Volume step, you will be able to increase the volume size by the amount of available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size.

This option is only available if:

- A RAID 0, RAID 1, RAID 5 and/or RAID 10 volume is present,
- The array allocation for the volume is less than 100% and space is available on the existing array.

Warning
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

Increasing the volume size from Manage Array

1. Under 'Status' or 'Manage', in the storage system view, click the array you want to manage. The array properties are now displayed on the left.
2. Click 'Increase size' next to the volume name. If more than one volume is present on a single array, you will need to increase the size of each volume one at a time.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Increasing the volume size from Manage Volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume whose size you want to increase. The volume properties are now displayed on the left.
2. Click 'Increase size'.
3. Click 'Yes' to confirm. Caution: Once the data migration starts, the operation cannot be canceled.

4. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes, or add another partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

18.4.2.8 Adding a Disk to a Volume

You can add one or more SATA disks to an existing array to increase the system storage capacity. This feature can be useful if you want to change to a volume type that requires additional disks.

This option is only available if:

- A RAID 0 and/or a RAID 5 volume is present,
- One or more SATA disks are connected to the computer and available,
- The available disk matches the internal or external connection type of the existing array disks. You cannot add an external disk to an array that includes internal disks, and vice versa. In specific advanced system configurations, this condition may not apply.

Refer to Connecting a Disk under Managing Disks for more information on installing SATA disks on your computer.

Warning
Any existing data on the available disk used to increase the array size will be permanently deleted. Backup all the data you want to preserve before completing this action.
If you perform a driver upgrade or downgrade while the data migration is in progress and then restart your computer, the driver will not be able to recognize the volume or the data on it. If you are migrating a system volume, you will not be able to restart your system because the operating system cannot load. If you are migrating a data volume, you will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state

This action can also be performed from Manage Array. Refer to the Adding a Disk to an Array section for more information.

1. Under 'Status' or 'Manage', in the storage system view, click the volume to which you want to add a disk. The element properties are now displayed on the left.
2. Click 'Add disk'.
3. Select the disk you want to use to increase the array capacity.
4. Click 'Add Disk'. Caution: Once the data migration starts, the operation cannot be canceled.
5. Once the migration has completed, restart your computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes for which a disk was added, or add another partition.

Note
To open Windows Disk Manager, click Start, right click My Computer, select Manage, then in the console tree select Disk Management

18.4.2.9 Changing Update Mode

A recovery volume gives you the flexibility to choose between updating data on the recovery disk continuously or on request.

In continuous update mode, the latest master disk changes are copied to the recovery disk automatically, as long as both disks are connected to the computer. In on request mode, the latest master disk changes are copied to the recovery disk only when you request a data update.

The current update mode is reported in the volume properties under Manage Volume. By default, the recovery volume is created in continuous update mode.

Note
This action is only available if a recovery volume is present and in normal state. If the recovery volume is read-only because the master or recovery disk files are accessed, you will need to hide the files before the update mode can be changed.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Change mode’, and then click 'Yes' to confirm.
3. The page refreshes and the volume properties report the new update mode.

18.4.2.10 Updating Data

You can manually copy the latest master disk changes to the recovery disk at any given time; this action allows you to synchronize data on the recovery volume, improving data protection and lowering the risk of losing valuable data in the event of a disk failure. When you request an update, only changes since the last update are copied.

Note
This action is only available if a recovery volume is present, and in ‘on request’ update mode.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Update data’.
3. The update process can be instantaneous or may take a while depending on the amount of data being copied. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.



Note
You can follow the progress of the update by hovering over the notification area icon or under ‘Status’ or Manage Volume.

18.4.2.11 Accessing Master or Recovery Disk Files

This action is only available if a recovery volume is present, in a normal state, and in on request update mode.

You can view the recovery or master disk files using Windows Explorer* depending on the designated source drive of the recovery volume. This feature can be useful when a data recovery from or to the master disk is necessary.

Accessing recovery disk files
This action is only available if the master disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click 'Access recovery disk files'.
3. Windows Explorer opens and displays the files located on the recovery disk.

Accessing master disk files
This action is only available if the recovery disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click 'Access master disk files'.
3. Windows Explorer opens and displays the files located on the master disk.

Note
When files have been accessed, the disk is displayed as missing from the array, and becomes available. Also, the volume is set to read-only and data updates are not available in this state. Hiding disk files will make the volume writable and allow data updates.

You can also access master or recovery disk files from Manage Disk.

18.4.2.12 Hiding Master or Recovery Disk Files

This action is only available if a recovery volume is present and disk files have been accessed.

When you are done viewing master or recovery disk files, you can hide the display of the files from Windows Explorer*. Once the disk files are hidden, the disk becomes writable, and data updates on the volume are available.

Hiding recovery disk files



This action is only available if the master disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Hide recovery disk files’.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Hiding master disk files

This action is only available if the recovery disk is the designated source drive and the volume is running off that disk.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. Click ‘Hide master disk files’.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Note

You can also hide master or recovery disk files from Manage Disk

18.4.2.13 Deleting a Volume

Use caution: you cannot recover data once a volume is deleted.

When a volume is deleted, you create available space that can be used to create new volumes. Note that you cannot delete a system volume using this application because the operating system needs the system files to run correctly. Also, if the volume is a recovery volume and the master or recovery disk files are accessed, you will need to hide these files before the volume can be deleted.

Warning

When a volume is deleted, all existing data on all disks that are a part of the selected volume is permanently lost. It is recommended to complete a backup of all valuable data before continuing.

1. Under ‘Status’ or ‘Manage’, in the storage system view, click the volume you want to delete. The volume properties are now displayed on the left.
2. Click ‘Delete volume’.
3. Review the warning message, and click ‘Yes’ to delete the volume.
4. The ‘Status’ page refreshes and displays the resulting available space in the storage system view. You can now use it to create a new volume.

18.4.2.14 Setting the Data Strip Size

You can assign a data strip size to a volume while creating a new volume or while changing the type of an existing volume. You cannot change the strip size of an existing volume without changing its type.



The strip size refers to each logical contiguous data block used in a RAID 0, RAID 5, or RAID 10 volume. This setting is not available for RAID 1 or recovery volumes, due to their redundant configuration. The default value is the recommended strip size based on the system configuration and the volume type selected; changing the pre-selection is best suited for advanced users.

The following table describes the usage scenarios for the typical strip sizes.

Usage scenarios for supported strip sizes*

Strip Size	Description	RAID Types
4 KB	Best for Web Servers (fast read transfer rate with slow write transfer rate).	RAID 0, 10
8 KB	Best for databases (fast read transfer rate with faster write transfer rate than with 4KB strips).	RAID 0, 10
16 KB	Good for sequential transfers.	RAID 0, 5, 10
32 KB	Best for sequential transfers.	RAID 0, 5, 10
64 KB	Best general purpose strip size.	RAID 0, 5, 10
128 KB	Best for audio and video editing.	RAID 0, 5

*Disclaimer: The data provided in this table may vary based on the brand, type, size, and speed of the disks used.

Setting the strip size when creating a volume

1. Under 'Status', click ‘Create’ or ‘Create a custom volume’.
2. Select the volume type, and then click 'Next'.
3. Make the required disk selection, and then select a new data strip size from the drop-down list in the Advanced section.
4. Complete the volume creation process as described in the Creation Process topic.

Setting the strip size when changing volume type

1. Under 'Status' or 'Manage', in the storage system view, click the RAID volume that you want to modify. The volume properties are now displayed on the left.
2. Click 'Change type'.
3. Make the necessary volume type and disk selections, and then select a new data strip size.
4. Click ‘OK’ to change the type of the existing volume.
5. The ‘Manage’ page refreshes and reports the new volume configuration.

Available Strip Size Configurations

	RAID 0	RAID 5	RAID 10
Default			
SATA disks	128 KB	64 KB	64 KB
Solid state disks	16 KB	128 KB	16 KB
Options	4 KB, 8 KB, 16 KB, 32 KB, 64 KB, 128 KB.	16 KB, 32 KB, 64 KB, 128 KB.	4 KB, 8 KB, 16 KB, 32 KB, 64 KB.



18.4.2.15 Enabling Volume Write-back Cache

You can improve the read/write performance of a RAID or recovery volume by enabling the write-back cache on one or all volumes on an array. When this feature is enabled, data may be temporarily stored in the cache memory before being written to the physical disks. Multiple I/O requests may be grouped together to improve performance. By default, the write-back cache is disabled.

Warning
While this feature highly improves the volume and array performance, it also increases the amount of cached data that could be lost in the event of a power failure. This risk can be lowered if your computer is connected to an uninterruptable power supply (UPS)

Enabling the volume write-back cache

1. Under 'Status' or 'Manage', in the storage system view, click the volume for which you want to enable the write-back cache. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Enable', and then click 'Yes' to confirm.
3. The page refreshes and the write-back cache is now enabled.

Note
If your computer is running on battery and a recovery volume is present, the option to enable the write-back cache is not available because the recovery disk is offline and data updates are not available. If this feature was enabled prior to running the battery, write-back cache activity would be temporarily disabled until you reconnect your computer to the power supply.

Disabling the volume write-back cache

1. Under 'Status' or 'Manage', in the storage system view, click the volume for which you want to disable the write-back cache. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Disable', and then click 'Yes' to confirm.
3. The page refreshes and the write-back cache is now disabled.

18.4.2.16 Initializing a Volume

Initializing a volume is the process of synchronizing all redundant data on a volume prior to verifying or verifying and repairing that data. If you attempt to start a verification process for a volume that has not been initialized, you will be prompted to do so.

Initializing a volume

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to initialize. The volume properties are now displayed on the left.
2. Click 'Initialize'.
3. Click 'OK' to start the initialization process. Caution: Once the data migration starts, the operation cannot be canceled.

Initializing a volume when verifying data



1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the left.
2. Click 'Verify'.
3. When prompted to initialize the volume before verifying data, click 'OK' to start the initialization process. Caution: Once the data migration starts, the operation cannot be canceled.
4. Once complete, click 'Verify' to start the verification process.

Note
While inititialization is in progress, you can view the status in the notifications area by hovering over the Intel(R) Rapid Storage Technology icon, or in the application under Staus or Manage Volume.

Warning
The initializating process could take a while depending on the number and size of the disks. You can continue using array disks and other applications during this time. Closing the application, or powering off and restarting your computer will not disrupt the progress of this operation.

18.4.2.17 Verifying and Repairing Data

You can verify data on an existing volume by identifying and repairing inconsistencies. Running this operation on a regular basis helps you keep valuable data and the overall storage system healthy.

1. Under 'Status' or 'Manage', in the storage system view, click the volume that you want to verify. The volume properties are now displayed on the left.
2. Click 'Verify'.
3. Select the check box if you want errors found to be repaired automatically during the verification process.
4. Click 'OK' to start the verification process.

Note
Data on a volume cannot be verified and repaired unless the volume has been initialized first. If you attempt to start a verification process for a volume that is not initialized, you will be prompted to first initialize the volume. Based on its configuration, a RAID 0 volume cannot be repaired because of the lack of redundancy.

18.4.2.18 Swapping Disks

You can change the order of designation for array disks in a recovery volume by setting the master disk as the destination drive and the recovery disk as the source drive. This action is best suited for advanced users.



Note
This action is only available if a recovery volume is present, normal, and in continuous update mode.

Swapping disks can be useful if:

- You selected the wrong disk as the master disk when you created the recovery volume,
- You think one of the disks is failing,
- You replaced the recovery disk with a faster, newer disk, and want to run off that device once it has been updated.

1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.
2. In the Advanced section, click 'Swap master and recovery disks'.
3. Click 'Yes' to confirm.
4. Hover over each disk in the storage system view to review their new usage.

18.4.3 Managing Disks

You must be logged on as an administrator to perform the actions listed in this section.

You can manage disks by clicking a selected disk in the storage system view under 'Status' or 'Manage'. This allows you to review the properties and access all actions associated with that disk, such as unlocking a password-protected disk or marking a disk as spare. Depending on their usage or status, some actions may not be available.

18.4.3.1 Disk Properties

The disk properties listed below display to the left of the storage system view under 'Manage' and report values such as usage and status that are specific to the disk selected in the view. Based on the detailed information provided, you can make changes to the way each disk is configured, or take action on one or more disk to keep your overall storage system healthy.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.
Usage	Array disk: a disk that has been grouped with other disks to form an array containing RAID volumes.
	Master disk: the disk that is the designated source drive in a recovery volume.
	Recovery disk: the disk that is the designated destination drive in a recovery volume.
	Spare: the disk has been designated as the destination drive for automatic volume rebuilds in the event of a failed, missing or at risk array disk. For RAID 0 volumes, automatic rebuilds will only occur when one of its array disks is reported as at risk.

	Available: the disk is physically connected to the computer, healthy, and available to be used in an array or as a spare disk. **Warning** Assigning an available disk to an array or marking it as a spare will permanently delete any existing data on that disk. Unknown: the disk is available but contains metadata that cannot be displayed in the operating system. Even though the disk is reported as normal, you will need to clear and reset the disk to make the disk available.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device. Enhanced: Indicates that the disk or volume is accelerated for optimized data protection. Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance. None: Indicates that no disk or volume is accelerated. Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.
Status	Normal: the disk is present, functioning as expected, and unlocked. Locked: the disk is password-protected. Note: if a volume includes at least one locked disk, the volume will display as locked. At risk: an impending error condition was detected on the disk and it is now at risk of failure. Missing: the disk is not present or physically connected to the computer. Failed: the disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold. Offline: indicates that an array disk is locked, that the recovery volume is in on request update mode, or that your computer is running on battery and data updates to the recovery volume are not available.
Size	Reports the total capacity of the disk in megabytes (MB) in the disk properties and in gigabytes (GB) in the storage system view.
Serial number	Reports the manufacturer's serial number for the disk.
Model	Reports the model number of the disk.
Firmware	Reports the version of the firmware found in the disk.
System disk	Reports whether the disk contains system files that are required to start and run the operating system.
Password protected	Reports whether the disk is protected with a password.
Disk data cache	Reports whether the data cache is enabled on this disk. This feature is controlled

	at the array level.
Native command queuing	Reports whether the disk supports this feature.
SATA transfer rate	Reports the data transfer rate between the SATA controller and the SATA disk. The supported rates are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.
Physical sector size	Reports the size of physical sectors on the disk (bytes).
Logical sector size	Reports the size of logical sectors on the disk (bytes).

18.4.3.2 Unlocking Password-Protected Disks**

You can unlock a password-protected disk by entering the password which allows you to access data or use that disk to create a volume. The password is setup through the system BIOS. Locked disks can be identified with the lock icon appended to them and display a 'Locked' status in the disk properties.

1. Under 'Status' or 'Manage', in the storage system view, click the disk you want to unlock. The disk properties are now displayed on the left.
2. Click 'Unlock'.
3. Enter the password, and then click 'Unlock'.

Note
After the fifth failed attempt to enter a valid password, you will need to power and power on your computer to reset the lockout on the disk you are attempting to unlock. Refer to the disk manufacturer or your IT administrator if you did not setup the password originally or you need to retrieve the password.

** This capability is only available if the system BIOS has it enabled. See section Unlocking Password Protected Disks in RST UI for details.

18.4.3.3 Marking a Disk as Spare

This action is only available for non-system disks in a normal state. Also, unless your mobile computer is equipped with the Intel® 5 Series Chipset or later, which provides support for up to six SATA ports, you will not be able to mark a disk as a spare. Most mobile computers are limited to one internal and one external disk, which are used to create the volume.

Marking a disk as a spare allows you to designate an available SATA disk as the default destination for automatic volume rebuilds in the event of a failed, missing or at risk array disk. However, for RAID 0 volumes, automatic rebuilds will only occur if one of its members is reported at risk.



1. Under 'Status' or 'Manage', in the storage system view, click the disk that you want to mark as a spare. The volume properties are now displayed on the left.
2. Click 'Mark as spare'.
3. Click 'OK'.

Note
RAID 1, 5, 10, and recovery volumes can use one or more spares.

Warning
When marking a disk as a spare, any existing data on that disk is permanently deleted. Back up all data you want to preserve before starting this action.

If you system is running a version of the RST OROM that does not support disks that are 2TB or larger, you can reset such a disk to available, but disallow the marking of it as a spare.

18.4.3.4 Resetting a Disk to Available

After a disk was marked as spare, you can choose to make that spare disk available again and use it differently. Once available, the disk can be used to create a volume or be added to an existing volume if all other requirements are met.

1. Under 'Status' or 'Manage', in the storage system view, click the disk that you want to reset to available. The volume properties are now displayed on the left.
2. Click 'Reset to available'.
3. The page refreshes and the disk usage is now reported as available.

18.4.3.5 Resetting a Disk to Normal

You can reset a SATA disk to normal when the storage system reports one of the following disk statuses:

At risk
A disk is reported at increased risk of failing in the near future that could be due to a slow degradation over time. You can choose to ignore this alert at this time by resetting the disk to normal, but it may re-appear if the disk continues to assert this condition. We recommend that you contact the manufacturer for more information to prevent potential data loss.

Failed
A SATA disk has failed to properly complete read and write operations in a timely manner, and data may be lost. We recommend that you replace the failed disk as soon as possible to return the overall storage system to normal. In this state, data may be lost, but you can try resetting the disk to normal and attempt a data recovery. If the disk operations continue to fail, the disk will return to a failed state immediately.



If the failed disk is an array disk, refer to the Troubleshooting section for guidelines on rebuilding a failed or degraded volume.

1. Under ‘Status’, in the Manage subsection, locate the disk reported as at risk or failed. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view.
2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.

Note
Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss.

18.4.3.6 Accessing Disk Files

This action is only available if a recovery volume is present, in a normal state, and is on request update mode.

This feature allows you to view the files on the designated destination drive in a recovery volume using Windows Explorer*. For example, you may want to review the recovery disk files prior to starting a data recovery in the event that data on the master disk is inaccessible or corrupted.

When the volume status is normal, the recovery disk is the designated destination drive and files are accessible. When the volume status is running off the recovery disk, the master disk is the designated destination drive and files are accessible. You can review the usage of each disk by hovering over the array disks in the storage system view or by clicking one of the disks to review its properties under Manage Disk.

1. Under ‘Status’ or ‘Manage’, in the storage system view, click the recovery or the master disk depending on the volume status. The disk properties are now displayed on the left.
2. Click ‘Access files’.
3. Windows Explorer opens and displays the files located on the disk.

Note
When files have been accessed, the disk is displayed as missing from the array, and becomes available. Also, the volume is set to read-only and data updates are not available in this state. Hiding disk files will make the volume writable and allow data updates.

Warning
Windows Explorer will not open if the disk does not have any partitions on it.



18.4.3.7 Hiding Disk Files

This action is only available if a recovery volume is present and disk files have been accessed.

When you are done viewing master or recovery disk files, you can hide the display of the files from Windows Explorer*. Once the disk files are hidden, the disk becomes writable, and data updates on the volume are available.

1. Under 'Status' or 'Manage', in the storage system view, click the disk whose files are accessed. The disk properties are now displayed on the left.
2. Click 'Hide files'.
3. The disk files no longer display in Windows Explorer.
4. The page refreshes and data updates on the volume are now available.

Note
You can also hide master or recovery disk files from Manage Volume.

18.4.3.8 Connecting a Disk

Installing new hardware is one of the steps you may have to take to keep you storage system healthy or to extend the life of a computer that is running out of storage space.

Intel® Rapid Storage Technology provides hot plug support, which is a feature that allows SATA disks to be removed or inserted while the computer is turned on and the operating system is running. As an example, hot plugging may be used to replace a failed external disk.

Our application provides support for SATA 1.5 Gb/s (generation 1), SATA 3 Gb/s (generation 2), and 6 Gb/s (generation 3) data transfer rates. The rate support depends on the Intel® Chipset and SATA disks present in your system. Visit our Online Support for additional information on chipset features and benefits.

Follow these procedures to replace or connect a disk in case you need to power off your computer:

Replacing a disk

1. Power off your computer.
2. Replace the disk that reports a problem.
3. Turn your computer back on. If the replaced disk was part of an array, you will need to follow the procedure provided in the Troubleshooting section based on the volume state and type.

Note
To install an external disk, plug it into you computer and connect the power cord.



To remove and install an internal disk, you should be comfortable opening your computer case and connecting cables. Follow the manufacturer's installation guide to complete this procedure. If you are replacing the system disk, you will have to re-install the operating system after you connect the disk because the system disk contains the files required to start and run your computer.

Installing a new disk (to increase storage space)

1. Power off your computer.
2. Connect the new disk.
3. Turn your computer back on. During the system startup, the application's option ROM should automatically detect the new disk if it is installed correctly. Once you open the application, verify under 'Status', in the storage system view, that the new disk displays. You can then access management options by clicking that disk.

18.4.4 Managing Ports

A port is a connection point on your computer where you can physically connect a device, such as a SATA disk or ATAPI device. A port transfers I/O data between the device and the computer.

If a port is reported as empty in the storage system view, you can use that port to connect a new device in order to increase the storage system capacity. Currently, the maximum number of internal ports that can be used to connect devices is six.

The port properties listed below display to the left of the storage system view under 'Manage' and report values specific to the element selected in the view.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.

18.4.5 Managing ATAPI Devices

An ATAPI device is a mass storage device with a parallel interface such as a CD-ROM, DVD/Blu-ray disc, tape drive, or solid-state disk. The ATAPI properties listed below display to the left of the storage system view under 'Manage' and report values specific to the selected element.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.
Serial number	Reports the manufacturer's serial number for the device.



Model	Reports the model number of the device.
Firmware	Reports the version of the firmware found in the device.
SATA transfer rate	Reports the transfer mode between the SATA controller and the ATAPI device. The typical values for this parameter are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.

18.5 Accelerating the Storage System

You can configure an internal solid state disk to be used as a non-volatile intelligent caching for a system or non-system disk or volume that's present on your system. Moving frequently accessed data over to the cache allows you to improve overall system performance, increase read/write access times, and reduce start up times without adding more system memory.

This feature also increases the power efficiency of a mobile computer by retaining stored data and reading data from the cache instead of the SATA disk itself.

Accelerate is only available if the requirements listed in this section under Cache Device Properties are met.

18.5.1 Cache Device Properties

The Accelerate tab and page are only available if the following requirements are met:

- Processor: Intel® Core™ i3, Intel® Core™ i5, or Intel® Core™ i7
- Operating system: Microsoft Windows* Vista 32-bit Edition and 64-bit Edition, Microsoft Windows* 7 32-bit Edition and 64-bit Edition or later, and Microsoft Windows* Server 2008 32-bit Edition and 64-bit Edition.
- Controller: Intel® Desktop/Workstation/Server Express Chipset SATA RAID controller hub and Intel® Mobile Express Chipset SATA RAID controller hub.
- BIOS: RAID-Ready system and Accelerate feature bit is set.
- An internal SATA solid state disk is present with a minimum capacity of 18.6 GB.
- A hard disk or volume (array members must all be hard disks) is eligible for acceleration.
- No recovery volume is present.

Limitations

- The maximum cache size is 64 GB.



- Only one disk or volume at a time can be accelerated per system.
- If two volumes are present on a single array (they share the same array of disks), neither volume can be accelerated.
- Once a volume is accelerated, a second volume cannot be added to the same array.
- Once a solid state disk is configured to be used as a cache device, the option to create a recovery volume is no longer available. Recovery volumes do not support system configurations with multiple volumes.

Solid State Disk Properties

Parameter	Value
Port	Reports the port number to which the solid state disk is attached.
Port location	Reports that the solid state disk is internal.
Status	Reports the state of health of the internal solid state disk present in the system.
	Normal: Indicates that the solid state disk is present, functioning as expected, and unlocked.
	Failed: Indicates that the solid state disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.
	At risk:
	Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that cache data is being deleted in order to disable acceleration. In some cases, these transitions will start automatically in the event that errors are detected and a risk of data loss is identified.
Type	Reports that the device is a solid state disk.
Usage	Reports that the solid state disk is configured to be used as a cache device.
Size	Reports that the solid state disk is configured to be used as a cache device.
Serial number	Reports the manufacturer's serial number for the internal solid state disk.
Model	Reports the model number of the solid state disk.
Firmware	Reports the version of the firmware found in the solid state disk.
Password protected	Reports whether the solid state disk is password-protected.
Disk data cache	Reports that the data cache is enabled on the solid state disk. When a solid state disk is configured as a cache device, this setting can only be changed at the operating system level.
Native command queuing	Reports whether the solid state disk supports this feature.
SATA transfer rate	Reports the data transfer rate between the SATA controller and the SATA solid state disk. The supported rates are: • SATA 1.5 Gb/s (generation 1) • SATA 3 Gb/s (generation 2) • SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in your system.
Physical sector size	Reports the size of physical sectors on the solid state disk (bytes).



Logical sector size	Reports the size of logical sectors on the solid state disk (bytes).
Accelerated device	Indicates the location of the disk or the name of the volume that is currently accelerated by the cache device.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device . Enhanced: Indicates that the disk or volume is accelerated for optimized data protection. Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance. None: Indicates that no disk or volume is accelerated.

Cache and Simple Data Volume Properties

Parameter	Value
Name	Reports the name of the volume.
Status	Normal: For simple data volumes, indicates that volume data is fully accessible. For cache volumes, indicates that cache data is fully accessible, and that caching activity with the associated disk or volume is occurring under healthy conditions. Failing: Indicates that a SMART event was detected on the solid state disk that is used as a cache device. Failed: Indicates that the cache volume has exceeded its recoverable error threshold, and that read and write operations are no longer occurring.
Type	Indicates that the single-disk RAID 0 volume is a cache volume.
Data stripe size	Indicates that the single-disk RAID 0 volume is a cache volume.
Allocated cache size	Reports the volume capacity used for cache memory.
Write-back cache	Reports whether the write-back cache feature is enabled for the volume.
Physical sector size	Reports the size of each sector that is physically located on the disk.
Logical sector size	Reports the size of data collection blocks.

Accelerated Disk or Volume Properties

Parameter	Value
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the cache device. Enhanced: Indicates that the disk or volume is accelerated for optimized data protection. Maximized: Indicates that the disk or volume is accelerated for optimized input/output performance. Busy: Indicates that acceleration is transitioning from maximized to enhanced mode, or that acceleration is being disabled from maximized mode. In the event that errors are detected and a risk of data loss is identified, transitions from maximized mode will start automatically.



18.5.2 Enabling Acceleration

You can enable acceleration in order to improve the performance for a SATA hard disk or a RAID volume that includes only SATA hard disks. This operation caches its contents using a non-volatile memory device (a solid state disk) that is attached to an AHCI port.

Enabling acceleration allows you to:

- Use a solid state disk as a cache device: The maximum cache size is 64 GB.
- Accelerate one system or non-system disk or volume present in the system by associating it with the cache volume, and subsequently caching its content.
- Configure acceleration in enhanced (optimized for data protection) or maximized mode (optimized for input/output performance). See below for more information on the acceleration modes.
- If the solid state disk used as a cache device is larger than 64 GB and has a minimum of 4 GB of additional space, a second single-disk RAID 0 volume will be automatically created which can be used for simple data storage.

Enabling Acceleration

Follow these steps to enable acceleration:

1. Click 'Enable acceleration' either under 'Status' or 'Accelerate'.
2. Select the solid state disk you want to use as a cache device.
3. Select the portion of the solid state disk you want to use to store non-volatile cache memory. Any remaining space on the solid state disk may be used for data storage using the simple data single-disk RAID 0 volume that is automatically created.
4. Select the disk or volume you want to accelerate. We highly recommend that you accelerate the system volume or system disk for maximum performance.
5. Select the acceleration mode you want to use, and then click 'OK'. By default, enhanced mode is selected.
6. The page refreshes and reports the new acceleration configuration in the Acceleration View.

Acceleration modes

Non-volatile cache memory can be enabled in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**

  This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the data in the cache memory. For data safety reasons, this mode is the default acceleration setting.

- **Maximized mode: Acceleration optimized for input/output performance.**



This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn't synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

18.5.3 Disabling Acceleration

You can disable acceleration on a disk or volume if you want to:

- Enable acceleration on a different disk or volume,
- Return the solid state disk to pass-through,
- Physically move an accelerated disk or volume to another computer.

Completing this action makes any cached data associated with the accelerated disk or volume immediately inaccessible. If the current acceleration mode is maximized, disabling acceleration may take a while to complete, depending on the cache and the solid state disk size. You can use other applications during this time.

1. Click 'Accelerate', and then click 'Disable acceleration'.
2. In the dialog, click 'Yes' to confirm.
3. The page refreshes and reports the acceleration as disabled.

In the event that you are unable to open or access Intel® Rapid Storage Technology due to an application error or operating system issue, you will need to disable acceleration using the option ROM user interface.

1. Restart your computer.
2. Press Ctrl-I to access the main menu of the option ROM user interface.
3. Select 'Acceleration Options' from the main menu.
4. Select the accelerated disk or volume.
5. If acceleration is in maximized mode, type 's' to synchronize data from the flash memory to the accelerated disk or volume. Otherwise, go to step 7.
6. Press 'Y' to confirm.
7. Type 'r' to remove acceleration.
8. Press 'Y' to confirm.

18.5.4 Changing Acceleration Mode

This action is only available if a disk or volume is currently accelerated. A disk or volume can be accelerated in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**
  This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the



data in the cache memory. For data safety reasons, this mode is the default acceleration setting.

- **Maximized mode: Acceleration optimized for input/output performance.**
  This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn't synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

By default, acceleration is enabled in enhanced mode due to the lower risk of data loss, but you can change acceleration mode at any time as long as the cache volume and accelerated device are in a normal state and caching activity is occurring.

Follow these steps to change the acceleration mode:

1. Click 'Accelerate', and then click 'Change mode'.
2. Click 'Yes' to confirm the mode change to either enhanced or maximized, depending on the current acceleration mode.
3. The page refreshes and the new acceleration mode displays under the Acceleration Configuration subsection and the Acceleration View.

Warning
When a device is accelerated in Maximized mode, performance is highly improved but cached data is at higher risk of being lost in the event of a power failure or under other atypical conditions.

Acceleration in a busy state

The acceleration mode will display as busy under the following conditions (by user interaction or automatic transition):

- When changing acceleration mode from maximized to enhanced.
- When disabling acceleration while in maximized mode.

The transition time varies based on the cache and disk sizes. Disk and volume actions will not be available until the acceleration transition has completed, except for renaming and deleting volumes.

18.5.5 Accelerating a Disk or Volume

Once a solid state disk is configured to be used as a cache device, you can choose to accelerate any disk or volume in a normal state that is located on your storage system. We recommend that you accelerate the system disk or volume in order to get the full benefits of the non-volatile cache memory configuration.

Follow these steps to accelerate a disk or volume:

1. Click 'Accelerate', and then click 'Select device'.
2. In the 'Accelerate Disk or Volume' dialog, select the device you want to accelerate.



3. Select the acceleration mode you want to use, and then click 'OK'. By default, enhanced mode is selected.
4. The page refreshes and reports the newly accelerated device in the Acceleration Configuration subsection. The Acceleration View also indicates the accelerated device with the acceleration icon appended to it.

Non-volatile cache memory can be enabled in either of the following modes:

- **Enhanced mode (default): Acceleration optimized for data protection.**
  This mode uses the write-through cache method to write data to the cache memory and the disk simultaneously. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is no risk of data loss because data on the disk is always synchronized with the data in the cache memory. For data safety reasons, this mode is the default acceleration setting.
- **Maximized mode: Acceleration optimized for input/output performance.**
  This mode uses the write-back cache method where data is written to the disk at intervals. In the event that the accelerated disk or volume becomes inaccessible, fails, or is disconnected, there is a high risk of data loss. In most cases, data on the disk wasn't synchronized with the data in the cache memory when the event occurred, or new data was written to the disk after the event occurred and it can no longer be synchronized with the cache memory.

18.5.6 Resetting a Cache Device to Available

This action is only available if a solid state disk is configured as a cache device and there is no accelerated disk or volume present (no association with the cache device). In this situation, you have two options:

- Reset the solid state disk to available and use that device for other purposes.
- Accelerate a disk or volume that is eligible and available for acceleration. Refer to Cache Device Properties for a detailed list of eligibility requirements.

Warning
In the event that a single-disk RAID 0 data volume was created along with a cache volume, resetting the solid state disk to available will delete both volumes. Data on the RAID 0 data volume will be permanently erased. Backup all valuable data before beginning this action.

1. Click 'Accelerate'.
2. Click 'Reset to available'.
3. In the dialog, select the check box to confirm that you understand that data on the data volume will be permanently deleted.
4. Click 'Yes' to confirm.
5. The 'Accelerate' page refreshes. Under 'Status', the storage system view displays the solid state disk usage as available. The device can now be used for any purpose.



18.5.7 Disassociating the Cache Memory

This action is only available if an issue is reported on the accelerated disk or volume that is associated with the cache device and it is missing. In this state, the acceleration mode is typically reported as unavailable and caching activity is no longer occurring.

If you are unable to resolve the reported issue on the accelerated disk or volume (e.g., degraded or failed volume due to a missing array disk), the only option will be to remove the association between the cache device and the disk or volume.

Once the association between the cache and the accelerated disk or volume is removed, all cache metadata and data is deleted from the cache device. You can then reset the solid state disk to available or accelerate a different disk or volume, as long as the cache device is healthy.

Follow these steps to disassociate the cache memory and the accelerated device:

1. Click ‘Accelerate’.
2. Click ‘Disassociate’.
3. In the 'Disassociate' dialog, click 'Yes' to confirm.
4. The page refreshes and the Acceleration View reports the new configuration. Options to reset the solid state disk to available or to select a new device to accelerate (as long as an eligible disk or volume is available) are now available.

Note
You can also perform this action using the option ROM user interface.

18.6 Preferences

System preferences allow you to decide whether you want the notification area icon to display. In addition, you can select the types of notifications you want to receive, such as storage system warnings or errors, and be notified of any reported problems while the application is closed.

Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Settings changes are applied on a per user basis, and do not affect other users’ settings.

Showing system notifications

By default, System preferences are set to show the notification area icon. If you previously chose to hide the notification area icon, follow these steps to display the icon again:

1. Under ‘Preferences’, select ‘Show the notification area icon’.
2. Click ‘Apply Changes’. Verify that the icon is now displayed in the notification area.



Hiding system notifications

Once you hide the notification area icon, the service no longer reports storage system information, warnings, or errors through the notification area. You will need to use the application to monitor the health of the storage system. Follow these steps to hide the notification area icon:

1. Under ‘Preferences’, deselect 'Show the notification area icon'.
2. In the ‘Hide Notification Area Icon’ dialog, click ‘Yes’ to confirm.
3. Verify that the icon is no longer displayed in the notification area.

Selecting system notifications

1. Under ‘Preferences’, make sure that ‘Show the notification area icon’ is selected.
2. Select the types of notifications you want to receive.

**Note**

Storage system information provides details on any changes of state other than warnings or errors, such as new disks being detected or locked.

Storage system warnings report the cause for the overall warning state of the storage system, such as a degraded RAID volume due to a missing disk.

Storage subsystem errors report the cause for the overall error state of the storage system, such as a failed volume due to a failed disk.



19 Using the BCFS to Differentiate Platform SKUs

Beginning with the Intel® RST 10.x Release and the Intel® 5 Series Express Chipset (codename Ibex Peak), the BIOS Control Feature Set (BCFS) has been enabled to give OEMs the opportunity to customize the Intel® RST features offered on any particular Intel® 5 Series Express Chipset model/SKU and later. OEMs no longer need special Intel® RST OROM images from Intel in order to enable/disable certain desired features for a platform SKU. OEMs can now enable/disable the desired features per platform SKU directly in their BIOS code. By clearing or setting the corresponding bits of the '***Intel RST Feature Capabilities***' register in the Intel chipset's SATA controller MMIO space, OEMs now have greater flexibility in determining what Intel® RST features will be supported per platform model/SKU.

The following sections explain the use of each of the bits in the BCFS, also known as the Software Feature Mask bits.

Note: This document does not cover details on how to setup a system BIOS. For that level of information please contact your platform's BIOS vendor or your Intel field representative to put you in contact with the appropriate Intel BIOS support personnel.

19.1 Configuring the Platform's RAID Related Features

When the BIOS has set the SATA Controller's mode to RAID, the following bits of the 'Intel RST Feature Capabilities' register in the Intel chipset's SATA controller MMIO space will determine what RAID levels will be supported on the platform SKU:

Note: Clearing all RAID level related bits to '0' (that includes the Intel® RRT bit) is an unsupported configuration. The Intel® RST OROM will ignore the BIOS settings and enable all RAID levels (Intel® RRT inclusive).

19.1.1 Configuring the Standard Supported RAID Levels

There are four (4) bits that control the 4 standard RAID levels supported by Intel® Rapid Storage Technology:

Bits	Type	Reset/ Default	Description



3	RWO	1h	**RAID 5 Enable (R5E):** If set to '1', then RAID5 is enabled
2	RWO	1h	**R10 Enable (R10E):** If set to '1', then RAID10 is enabled
1	RWO	1h	**RAID 1 Enable (R1E):** If set to '1', then RAID1 is enabled
0	RWO	1h	**R 0 Enable (R0):** If set to '1', then RAID0 is enabled

19.1.1.1 Example Configuration

To configure a platform SKU that offers **only** RAID levels 0 and 10, the bits must be configured as follows:

Bit 0 == 1 (default)

Bit 1 == 0

Bit 2 == 1 (default)

Bit 3 == 0

19.1.2 Configuring Intel® RRT Related RAID Features

There are two (2) bits that control two capabilities/features related to the Intel® RRT feature:

Bits	Type	Reset/ Default	Description
8	RWO	0h	**Intel® RRT Only on ESATA (ROES):** If set to '1', then only Intel® RRT volumes can span internal and eSATA drives. If cleared to '0', then any RAID volume can span internal and eSATA drives.
4	RWO	1h	**Intel® RRT Enable (RSTE)**: If set to '1', then Intel Rapid Recovery Technology is enabled.

19.1.2.1 Example Configuration

To configure a platform SKU that offers Intel® RRT, only RAID level 5, and allows both Intel® RRT and RAID5 Volumes to span disks on both internal and external (eSATA) ports, the bits must be configured as follows (**Note**: an Intel® RRT volume cannot coexist with another RAID level volume at the same time on the platform):

Bit 0 == 0



Bit 1 == 0

Bit 2 == 0

Bit 3 == 1 (default)

Bit 4 == 1 (default)

Bit 8 == 0 (default)

19.1.3 Configuring the Behavior of the OROM UI and Banner

There are three (3) bits that control the behavior of the Intel® RST OROM UI and the Banner Splash Screen that are displayed during POST at system boot-up. Use the following bit configurations to determine whether or not the splash screen will be displayed during post and if so, how long the delay will be before the system continues the boot process:

Bits	Type	Reset/ Default	Description
11:10	RWO	0h	**OROM UI Normal Delay (OUD)**: Values of these two bits specify the delay of the OROM UI Splash Screen in a normal status. 00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs **Note:** If bit 5 == 0, then these values are disregarded **Comment:** Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.
5	RWO	1h	**Intel RST OROM UI (RSTOROMUI):** If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.

19.1.3.1 Example Configuration

To configure a platform SKU that enables the OROM Banner Splash Screen to be displayed for 6 seconds, the bits must be configured as follows:



Bit 5 == 1 (default)

Bit '10' == 0 (default)

Bit '11' = = 1

19.1.4 Configuring Intel® RST UI Capabilities

There is one capability within the Intel® RST UI that is controlled by the BCFS bits. To enable/disable the ability for the Intel® RST UI to unlock password protected disks, use the following bit configurations:

Bits	Type	Reset/ Default	Description
6	RWO	0h	**HDD Unlock (HDDLK):** If set to '1', then HDD password unlock is enabled in the OS.

The default settings for these two features are:

Bit 6 == 0

When this bit is cleared the Intel® RST UI does not display any option to use this feature.

19.1.4.1 Example Configuration

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

Bit 6 == 0 (default)

19.1.5 Configuring the Platform to Support Intel® Smart Response Technology

The BCFS bit is only one of the platform parameters (see section 17.1.1 for all requirements) that determines whether Intel® Smart Response Technology can be enabled on a platform, however, unless this bit is set there will be no support for this feature on the platform.

Bits	Type	Reset/ Default	Description
9	RWO	0h	**Intel® Smart Response Technology Enable**

			Request (SEREQ): If set to '1', then Smart Response Technology is enabled. If cleared to '0', Smart Response Technology is disabled.

Bit 9 == 1

This enables the Intel® Smart Response Technology feature on the platform SKU.

19.1.6 BIOS Control Feature Set

Bits	Type	Reset	Description
15:14	RO	0h	Reserved.
13:12	RWO	0h	Reserved
11:10	RWO	0h	**OROM UI Normal Delay (OUD)**: Values of these bits specify the delay of the OROM UI Splash Screen in a normal status. 00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs If bit 5 == 0, then these values are disregarded **Comment: Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.**
9	RWO	0h	**Intel® Smart Response Technology Request Enable (SEREQ):** If set to '1', then Smart Response Technology is enabled. If cleared to '0', Smart Response Technology is disabled.
8	RWO	0h	**Intel® RRT Only on ESATA (ROES):** If set to '1', then only Intel® RRT volumes can span internal and eSATA drives. If cleared to '0', then any RAID volume can span internal and eSATA drives.
7	RWO	0h	**Not implemented in the RST driver**
6	RWO	0h	**HDD UNLOCK (HDDLK):** If set to '1', then HDD password unlock is enabled in the OS.
5	RWO	1h	**Intel® RST OROM UI (RSTOROMUI):** If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.
4	RWO	1h	**Intel® RRT Enable (RSTE)**: If set to '1', then Intel Rapid Recovery



			Technology is enabled
3	RWO	1h	**RAID 5 Enable (R5E):** If set to '1', then RAID5 is enabled
2	RWO	1h	**RAID 10 Enable (R10E):** If set to '1', then RAID10 is enabled
1	RWO	1h	**RAID 1 Enable (R1E):** If set to '1', then RAID1 is enabled
0	RWO	1h	**RAID 0 Enable (R0E):** If set to '1', then RAID0 is enabled



20 Testing, Certification Notes

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

For Internal SATA ports with interlock switches, the RST driver will set Removable=TRUE in the IRP_MN_QUERY_CAPABILITIES handler. This causes Win7 to show the internal device in its own "container" which is used to describe devices that are external to the system. For example, a CD-ROM on an interlocked switch in a Win7 system, under 'Devices and Printers',
You can see that the CD-ROM on the interlocked SATA port shows up separately in the top-level 'Devices' view. This can result in a platform WHQL test failure.

There is a whitepaper describing use of Removable device capability bits on Win7 by bus drivers:
http://www.microsoft.com/whdc/Device/DeviceExperience/ContainerIDs.mspx.

In order to correct this issue to pass the platform WHQL test, RST recommends the OEM to take the following action:

In the system BIOS, define an _EJ0 ACPI method on the interlocked port. _EJ0 will signal to the ACPI driver to set Removable for the RST driver and still mark the device as internal to the system such that it does not show in its own container. The implementation is to use a registry key for each port to tell RST whether to set Removable bit or not. If _EJ0 ACPI method is defined in the system BIOS by the manufacturers, they can tell RST not to set the Removable bit. For example:

[HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaStor\Parameters\Port1]
"EJ0IsDefined"=dword:1
If 1, _EJ0 will set Removable bit instead of RST. If 0, no _EJ0 defined so RST will set Removable bit. The default value is 0.



21 Glossary

Term	Definition
ATA	Advanced Technology Attachment
BIOS	Basic Input/Output System
BOM	Bill Of Materials
CD	Compact Disc
Chipset	Term used to define a collection of integrated components required to make a PC function.
Hard drives	Physical hard drives attached to a RAID controller
DOS	Disk Operating System
GB	Giga-byte
HDD	Hard Drive
I/O	Input/Output
ICH	I/O Controller Hub
ICH9	Intel® 82801IR/DO SATA RAID Controller
IDE	Integrated Drive Electronics
INF	Information file (.inf) used by Microsoft operating systems that support the Plug & Play feature. When installing a driver, this file provides the OS needed information about driver filenames, driver components, and supported hardware.
Intel® Option ROM (OROM)	Standard Plug and Play option ROM that provides a pre-operating system user interface for the Intel RAID implementation.
MB	Mega-byte
Migration	Term used to describe the movement of data from one configuration or usage model to another.
OEM	Original Equipment Manufacturer
Option ROM	A code module built into the System BIOS that provides extended support for a particular piece of hardware. For this product, the Option ROM provides boot support for RAID 0/1/5/10 volumes, and provides a user interface for configuring and managing RAID 0/1/5/10 volumes.
OS	Operating System
PATA	Parallel ATA
PCI	Peripheral Components Interconnect
PFW	Package for the Web
PIO	Programmed Input Output
PnP	Plug and Play



Term	Definition
Port 0..3	Term used to describe the point at which a SATA drive is physically connected to the SATA Controller. Port n is the nth of the four available ports in ICH9 systems, where n=0..3
RAID	Redundant Array of Independent Disks
RAID 0	A RAID level where data is striped across multiple physical hard drives (aka striping)
RAID 1	A RAID level where data is mirrored between hard drives to provide data redundancy (aka mirroring)
RAID 5	A RAID level where data and parity are stripped across the hard drives to provide good read/write performance and data redundancy. The parity is stripped in a rotating sequence (aka Stripping and rotating parity).
RAID 10	A RAID level where information is striped across a two disk array for system performance. Each of the drives in the array has a mirror for fault tolerance. (aka Stripping and mirroring)
RAID volume	A block of capacity allocated from a RAID Array and arranged into a RAID topology. Operating Systems typically interpret a RAID volume as a physical hard drive.
RAM	Random Access Memory. Usually refers to the system's main memory
ROM	Read Only Memory
SATA	Serial ATA
SCSI	Small Computer System Interface
SP#	Service Pack (number)
Strip	Grouping of data on a single physical hard drive within a RAID volume
Stripe	The sum of all strips in a horizontal axis across physical hard drives within a RAID volume
UI	User Interface

Glossary

A B C D E F H I L M N O P R S U V W

- A
  - **Accelerated disk or volume**
    A disk or RAID volume that has its non-volatile cache enabled in either maximized or enhanced mode.
  - **Access master or recovery disk files:**
    Action to view the files located on the master or recovery disk of a recovery volume using Windows Explorer*. Selecting this option sets the volume in read-only and volume updates are not allowed. This action is only available when the recovery volume is in on request update mode or running off the recovery disk.
  - **Activate port LED:**



Action to locate the port connected to a disk present on the system by activating the Light Emitting Diode (LED) light.

- **AHCI:**
  An interface specification that allows software to communicate with SATA devices such as host bus adapters, and enables advanced SATA features such as Native Command Queuing, native hot plugging, and power management. Advanced Host Controller Interface (AHCI).
- **Array:**
  An abstraction layer or collection of two or more disks used to manage RAID volumes existing on a same collection of disks. RAID arrays are not visible to the operating system.
- **At risk:**
  Status indicative that a disk or device has experienced a SMART event, and that an impending error condition was detected and the disk or device is now at risk of failure.
- **ATAPI device:**
  A mass storage device with a parallel interface such as CD-ROM, CD-RW, DVD-ROM, Blu-ray Disc, and tape drives. Advanced Technology Attachment Packet Interface (ATAPI).

· B

- **Blocks with media errors:**
  Number of inconsistencies found during the data verification of a RAID volume. This feature only applies to the verification process or the verification and repair process.

· C

- **Cache**
  A resource allocation on a storage component used for temporary data operations. Cache can be allocated in components such as RAM or non-volatile memory.
- **Cache device**
  The selected solid state disk used for cache storage.
- **Cache volume**
  The portion of the cache device (a solid state disk) that holds the non-volatile cache data. That portion is configured into a single-disk RAID-0 volume.
- **Change volume type:**
  Action to change the volume from one RAID configuration to another, and move data from one RAID volume to another. A RAID 1 volume can also be converted to a recovery volume, and vice versa.
- **Continuous update mode:**
  Update mode assigned to a recovery volume, where data on the master disk is copied to the recovery disk automatically, as long as both disks are connected to the system.

· D

- **Data stripe size:**
  Size of a grouping of data on a single physical disk within a RAID volume. Reported in kilobytes (KB).
- **Data volume**
  The portion of extra space on the cache device (a solid state disk) that can be used for data storage. That portion is configured into a single-disk RAID-0 volume.
- **Degraded:**
  Volume status indicative that one member has failed or is missing. This status only applies to recovery, RAID 1, RAID 5, and RAID 10 volumes.
- **Disassociating a cache**
  The action of removing the association between the non-volatile cache and the accelerated disk or volume.
- **Disk data cache:**
  A cache memory within a hard drive that temporarily stores frequently used data sectors for faster access. As a result, overall hard drive performance is improved.
- **Disk:**
  A hard or floppy disk. Also known as hard drive or hard disk drive.

· E



- **E-mail notification:**
  Alert mechanism that allows the user to receive storage system information, warning, and error notifications by e-mail via SMTP. By default, this feature is disabled and requires configuration settings such as the SMTP host and e-mail addresses to be set up.
- **Enhanced Mode**
  An acceleration mode that uses write-through, non-volatile cache to improve performance. The mode also is known as “separation safe” because all host-write requests are written to the accelerated disk or volume and possibly to the non-volatile cache.

F

- **Failed:**
  Volume and disk status indicative that one or more array members are missing or have failed.
- **Firmware:**
  Permanent instructions and data programmed directly into the read-only memory (ROM) for controlling the operation of the computer. Firmware usually requires updates to fix defects or add features to the hardware.

H

- **Hide master or recovery disk files:**
  Action to close the display of files located on a master or recovery disk in a recovery volume after viewing them in Windows Explorer*. This option is only available when ‘Access recovery disk files’ or ’Access master disk files’ was previously selected. Once disk files are hidden, volume updates can resume.
- **Hot plug:**
  Action to remove or insert a SATA disk when the system is powered on.

I

- **Increase volume size:**
  Action to expand the data storage capacity of a volume by utilizing the available array space on a RAID 0, RAID 1, RAID 5, or RAID 10 volume.
- **Initialize:**
  Process of synchronizing all redundant data on a volume prior to creating a volume, verifying and repairing data, or changing volume type. Initialization is still required for non-redundant volumes such as RAID 0 to ensure that data is readable before starting the verification process.
- **Intel® Rapid Recover Technology:**
  Official name for Intel’s technology that allows the user to copy data from a master disk (source) to a recovery disk (destination) either continuously or on request.
- **Intel® Rapid Storage Technology:**
  Official name for Intel's Windows-based software to provide support for high-performance, fault-tolerant, and capacity SATA RAID arrays on select supported chipsets. Intel Rapid Storage Technology also provides support for Intel® Rapid Recover Technology, AHCI Native Command Queuing, and matrix RAID for two RAID volumes on a single array.

L

- **Locked:**
  Volume and disk status indicative that the data is protected with a password and cannot be accessed until disks are unlocked.

M

- **Mark as spare:**
  Action to designate an available and compatible SATA disk as the default destination for automatic rebuilds in the event that an array member fails or is missing.
- **Master disk:**
  The disk that is the designated source drive in a recovery volume.
- **Maximized Mode**
  An acceleration mode that uses write-back, non-volatile cache to improve performance better than the enhanced Mode. This mode is optimized for input/output performance and power savings.



- **Migrating:**
  Volume status indicative that data is being moved/transferred across selected storage devices due to a change request in the storage system configuration, such as changing volume type, creating a volume preserving existing data, increasing the volume capacity, or changing data stripe size.

- N
  - **Native Command Queuing:**
    Command protocol in SATA that allows multiple commands to be outstanding within a disk at the same time. The commands are dynamically reordered to increase disk performance.
  - **Normal:**
    Volume, disk, and device status indicative that they are in a healthy state, functioning as expected, disks are properly connected, and data is fully accessible.

- O
  - **Offline:**
    Disk status indicative that an array disk is locked (the volume status displays as locked), that the recovery volume is in on request update mode, or that your computer is running on battery and data updates to the recovery volume are not available.
  - **On request update mode:**
    Update mode assigned to a recovery volume, where data on the master disk is copied to the recovery disk when the user requests it. Only changes since the last update process are saved to the recovery disk.
  - **Option ROM:**
    Firmware that is called by the system BIOS in order to communicate and provide support for a hardware device. For this product, the option ROM provides boot support for RAID volumes as well as a user interface for configuring and managing RAID volumes. Also known as OROM.

- P
  - **Port:**
    An internal or external data connection of a computer (e.g., SATA controller) to which a peripheral device (e.g., SATA disk) can be attached.

- R
  - **RAID 0:**
    A RAID type or configuration where data is striped across multiple physical disks. Data is split into manageable blocks called strips across array members. Striping does not create data redundancy but improves read/write performance.
  - **RAID 1:**
    A RAID type or configuration where data is mirrored across a second physical disk in the array. Mirroring is a key feature that ensures real-time data redundancy and increased fault tolerance. There is no striping.
  - **RAID 10:**
    A RAID type or configuration that uses four disks to create a combination of RAID type 0 and 1. The data is striped across a two-disk array forming a RAID 0 component. Each of the disks in the RAID 0 array is mirrored by a disk in the RAID 1 array.
  - **RAID 5:**
    A RAID type or configuration where data and parity are striped into manageable blocks called strips across three or more physical disks. This type is a preferred configuration as it combines efficiency, fault-tolerance, and data performance.
  - **RAID:**
    Redundant Array of Independent/Inexpensive Disks is a technology used for computer data storage schemes that divide and/or replicate data among multiple disks. RAID can be designed to provide increased data reliability or increased I/O (input/output) performance, or both. A number of standard schemes have evolved which are referred to as levels or types. Intel® Rapid Storage Technology software supports RAID 0, RAID 1, RAID 5, and RAID 10 configurations (refer to each RAID type definition for more information).
  - **Read-only:**

Recovery volume status indicative that the recovery or master disk files are accessed, allowing files on the disk to be read or copied, but not changed or saved.

- **Rebuild:**
  The process of restoring a recovery, RAID 1, RAID 5, and RAID 10 volume in the event that a volume disk has failed or is missing. If a spare disk is present and compatible, the application will automatically use it as a replacement for the failed disk. An automatic rebuild process will also occur if a RAID 1 member is removed and then reconnected, in order to re-establish the mirroring. This process does not apply to RAID 0 volumes.
- **Recover data:**
  The action of retrieving data in the event that a recovery volume has failed due to a missing or failed master disk. Data from the recovery disk is copied to a new or healthy master disk, restoring data redundancy. Selecting this option will overwrite all master disk data with data on the recovery disk.
- **Recovery disk:**
  The disk that is the designated destination drive in a recovery volume.
- **Recovery volume:**
  A two-disk redundant volume that includes a master disk (source) and a recovery disk (destination) and uses Intel® Rapid Recover Technology. This configuration provides flexibility of volume data updates and maximum data redundancy.
- **Repair:**
  The process of fixing verification errors and blocks with media errors found during the verification process. This feature is only available for volumes with a normal or at risk status. A RAID 0 cannot be repaired due to its non-redundant configuration.
- **Reset disk to normal:**
  Action to return a failed or at risk disk to a normal state. We recommend that you contact the manufacturer for more information to prevent potential data loss.
- **Reset to available:**
  Action to return a disk previously marked as a spare to an available state.
- **Reset volume to normal:**
  Action to return a failed volume where both array disks are present and normal to a normal state. This feature allows you to access and attempt a recovery of healthy volume data.

· S

- **SATA disk:**
  A disk with an interface that transmits data using a serial protocol in order to communicate with the SATA controller.
- **SATA transfer rate:**
  Rate at which the SATA controller and SATA disk communicate with each other. Transfer rates are important when large contiguous blocks of data are being used, such as video and image files. Reported in gigabits/seconds (Gb/s).
- **SATA:**
  A successor to ATA and PATA, SATA is a computer bus technology primarily designed for transfer of data between storage devices such as hard drives or optical devices, and a computer. Benefits of this technology are: usage of high-speed serial cables, air cooling to work more efficiently, faster transfers, the ability to remove devices while operating (hot plugging), enables more reliable operation with tighter data integrity checks. Also known as Serial Advanced Technology Attachment or Serial ATA.
- **Single-disk RAID 0**
  A RAID-0 volume that has one (1) disk as its array disk. This is the volume type used to create the non-volatile cache region on a solid state disk being used for caching. This volume type creates an OS-visible volume that enables access to extra space on the solid state disk that is not being used for caching.
- **Size:**
  Reports the total capacity of a physical device such as a SATA disk, a volume, or an array.
- **Storage system:**
  One or more physical disks or devices that act as a unit for data storage.



- **System disk:**
  A disk that contains system files required to start and run the operating system. By default, the Windows operating system files are in the WINDOWS folder, and the supporting files are in the WINDOWS\System32 folder.
- **System volume:**
  A volume that refers to the disk(s) volume that contains the hardware-specific files that are needed to start Windows, such as Ntldr, Boot.ini, and Ntdetect.com, as well as the Windows operating system files and supporting files. The system volume can be the same volume as the boot volume.

· U

- **Unknown:**
  Disk status indicative that its usage could not be determined, due to a possible incompatibility between this software version and the disk configuration, or a virus. Also a volume status indicative that the volume is in an unexpected state due to a configuration error. Data on the volume can no longer be accessed.
- **Update mode:**
  Type of update assigned to a recovery volume. The update mode can be set to continuous, where master disk changes are automatically saved to the recovery disk, or set to on request, where updates of the recovery disk can be requested immediately.

· V

- **Verify:**
  Action of scanning data to detect any types of data damage, disk read errors, and volume data inconsistencies. As an option, errors found can be corrected on redundant RAID volumes. This feature is only available for volumes with a normal or at risk status.
- **Volume size:**
  Amount of data that can be stored on a volume; reported in bytes (B) or %.
- **Volume type:**
  Configuration of a volume which determines how data is stored and managed to improve read/write performance, increase fault tolerance and/or storage capacity.
- **Volume write-back cache:**
  A cache memory used to enhance the read/write performance of a RAID volume by grouping multiple I/O requests into fewer requests and by writing from the cache to the volume at defined intervals.
- **Volume:**
  A fixed amount of space across a RAID array that is structured to emulate a single physical hard drive and appears as such to the operating system. Volumes have drive letters assigned to them and some volumes can span multiple hard disks.

· W

- **Windows Disk Management*:**
  Microsoft Windows* system utility for managing the disks and partitions or volumes that they contain. Disk Management allows the initialization of new disks, volume creation, and formatting. Most disk-related tasks can be performed using this system utility without shutting down or restarting the computer; most configuration changes take effect immediately.
- **Write-back cache allocation:**
  Size of the dynamic random access memory (DRAM) that is allocated for write-back caching on all volumes present on the system. The cache size is set to 16 megabytes (MB) by default.



22 Troubleshooting

This section explains how to resolve the most common problems that may occur while using the application. If you have any questions regarding installing, using or maintaining this product, you can also visit Intel's online support site which provides you with self-help resources and electronic problem submission.

22.1 Failed Volumes

RAID 0
A RAID 0 volume is reported as failed when one of its members is disconnected or has failed. In both cases, the volume and its data are no longer accessible.

Cause	Solution
Missing array disk	Follow this procedure to recover data: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer. During the system startup, the volume status will display as 'Normal' in the Intel Rapid Storage Technology option ROM user interface. 3. Once the operating system is running, open Intel Rapid Storage Technology from the Start menu or click the Intel Rapid Storage Technology notification area icon. 4. Under 'Status', verify that the volume and disks status display as 'Normal'. You can also review this information under 'Manage'.
Failed array disk	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disk with a new one that is of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 0 volume using the new disk. If the failed disk was part of the system volume, you will also need to reinstall the operating system.

RAID 5
A RAID 5 volume is reported as failed when two or more of its members have failed.



Cause	Solution
Two or more array disks failed	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 5 volume using the new disks. If the failed disk was part of the system volume, you will also need to reinstall the operating system.

RAID 10

A RAID 10 volume is reported as failed when two adjacent members are disconnected or have failed, or when three or four of its members are disconnected or have failed.

Cause	Solution
Two adjacent array disks missing	1. Power off your computer and reconnect the missing disks. 2. The rebuild operation will start automatically. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.
Three or four array disks missing	In most cases, the volume cannot be recovered and any data on the volume is lost. This procedure deletes the failed volume: 1. Power off your computer and reconnect the missing disks. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. You will then need to reinstall the operating system on the new volume.
Two or more array disks failed	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, you can try resetting the disks to normal, and then attempt a data recovery. If the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. This procedure deletes the failed volume: 1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity. 2. Turn on your computer. During the system startup, the volume status will display as 'Failed' in the Intel Rapid Storage Technology option ROM user interface. 3. Press Ctrl-I to access the main menu of the option ROM user interface. 4. Select Delete RAID Volume from the main menu.

	5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. You will then need to reinstall the operating system on the new volume.

22.2 Degraded Volumes

Recovery Volume
A recovery volume is reported as degraded when the recovery disk has failed or when the master disk is disconnected or has failed. Data mirroring and redundancy are lost because the system can only use the functional member.

Cause	Solution
Recovery disk failed	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, but if the read/write data access consistently fails, the disk will likely return to a failed state immediately. Refer to Troubleshooting Disk Events for instructions on resetting a disk to normal. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the recovery disk and volume status will display as 'Normal'. 5. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. **Note** If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Master disk missing	If you can reconnect the missing master disk, follow this procedure to recover data: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer and the system will automatically boot from the recovery disk. 3. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'. 4. Click 'Yes' to confirm. 5. The recovery operation starts immediately and cannot be canceled. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 6. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. If you cannot reconnect the missing disk and a SATA disk is available and normal,

	follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. 5. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. **Note** If you cannot reconnect the missing disk or rebuild to an available disk, you will need to power off the computer and connect a new SATA disk. Once rebuilt, the recovery volume will be limited to its original size even if the new disk is larger than the original master disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Master disk failed	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, but if the read/write data access consistently fails, the disk will likely return to a failed state immediately. To reset the failed master disk and the volume to normal, follow this procedure: 1. Under 'Status', click 'Reset disk to normal'. Note that the volume is now running off the recovery disk, and that the master disk is reported as offline. 2. Under 'Status', in the Manage subsection, click 'Recover data' or click the recovery volume in the storage system view, and then click 'Recover data'. **Warning** Starting this action will override existing data on the master disk and update it with the data on the recovery disk. Backup all valuable data before continuing. 3. Click 'Yes' to confirm. 4. The recovery operation starts immediately and cannot be canceled. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 5. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the master disk and volume status will display as 'Normal'. 5. Once completed, the volume returns to the last update mode to which it was set before the issue was reported. **Note** If there is no available disk present, you will need to power off your computer and connect a new SATA disk. Once rebuilt, the recovery volume will be limited to its original size even if the new disk is larger than the original master disk. Once your

	computer is back up and running you can follow the rebuild procedure described above.

RAID 1
A RAID 1 volume is reported as degraded when one of its members is disconnected or has failed. Data mirroring and redundancy are lost because the system can only use the functional member.

RAID 5
A RAID 5 volume is reported as degraded when one of its members is disconnected or has failed. When two or more array disks are disconnected or have failed, the volume is reported as failed.

RAID 10
A RAID 10 volume is reported as degraded when one of its members is disconnected or has failed, or when two non-adjacent members are disconnected or have failed. When two or more array disks are disconnected or have failed, the volume is reported as failed.

Cause	Solution
Missing array disk	If you can reconnect the missing disk, follow this procedure to rebuild the volume: 1. Power off your computer and reconnect the missing disk. 2. Turn on your computer and the rebuild operation will start automatically. If you cannot reconnect the missing disk and a SATA disk is available and normal, follow this procedure to rebuild the volume: 1. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. **Note** If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once your computer is back up and running you can follow the rebuild procedure described above.
Failed array disk	We recommend that you rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, you can try resetting the disk to normal, which will prompt the volume to start rebuilding automatically. But if the read/write data access consistently fails, the disk will likely return to a failed state immediately and you will need to rebuild the volume to another disk. If a SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. **Note** If there is no available disk present, you will need to power off your computer and connect a new SATA disk that is equal or greater capacity than the failed disk. Once

	your computer is back up and running you can follow the rebuild procedure described above.

22.3 Other Volume States

Locked	
Cause	Solution
At least one (but not all) disk included in the volume is locked with a password.	In this state, the overall storage system health is still reported as normal, but we recommend that you unlock the disks to make the volume data fully accessible. Follow this procedure to unlock a disk: 1. Under 'Status' or 'Manage', in the storage system view, click the disk you want to unlock. The disk properties are now displayed on the left. 2. Click 'Unlock'. 3. Enter the password, and then click 'Unlock'. Repeat this procedure for all locked disks included in the volume in order to unlock the volume. **Note** If all the disks included in a volume are locked, the volume is no longer displayed
Incompatible	
Cause	Solution
Indicates that the volume was moved to another system that does not support the volume type and configuration.	In this situation, volume data is accessible to the operating system and can be backed up, but the volume cannot operate because your system does not support its RAID configuration. Here are your options: 1. Reconnect the volume to the computer where the volume was originally created, and continue using it. 2. Delete the volume, and then create a new volume with a RAID configuration that is supported by the current system. Follow the procedure described above to delete the volume. **Warning** When a volume is deleted, all existing data on the member disks of the selected volume is permanently erased. It's recommended that you backup all valuable data prior to beginning this action.
Unknown	
Cause	Solution
The volume is in an unexpected state due to a configuration error.	The application is unable to detect the exact nature of the problem. Try restarting your computer. If the error persists, back up all valuable data and delete the volume using the option ROM user interface. Refer to the user's manual accessible from the Online Support area for details on using the option ROM.
Power-saving mode (Recovery volumes only)	
Cause	Solution
Your computer is running on battery and the volume is in continuous update	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. Reconnect your computer to the power supply, and the operation that was in progress prior to running on battery will resume automatically.

mode. Data updates to the recovery disk or a data recovery operation are not occurring.	**Note** If a data recovery was in progress, the overall storage subsystem health is reported as degraded because the operation could not be completed.
Data update needed (Recovery volumes only)	
Cause	Solution
The data on the recovery disk is not synchronized with the data on the master disk.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. Follow this procedure to update data on the recovery disk: 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2. Click 'Update data'. 3. Select the check box if you don't want this confirmation message to display each time you request an update. Click 'Yes' to confirm. 4. The update process can be instantaneous or may take a while depending on the amount of data being copied. You can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.
Running off recovery disk (Recovery volumes only)	
Cause	Solution
Your computer was booted from the recovery disk using the option ROM, and the volume is operating from that disk. With this reverse configuration, the recovery disk is the designated source drive and data updates to the master disk are not available.	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. A data recovery from the recovery disk to the master disk is required to maintain full redundancy. **Warning** Completing this action will overwrite all master disk data with the data on the recovery disk. Backup all valuable data prior to starting this action. 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2. Click 'Recover data', then 'OK' to confirm. 3. Once complete, we recommend that you restart your computer from the master disk using the option ROM user interface to return to a normal state.
Recovery disk read-only (Recovery volumes only)	
Cause	Solution
The recovery disk files have been accessed and display in Windows Explorer*.	In this state, any data written to the master disk is not copied to the recovery disk because it is read-only. Data mirroring and redundancy may be lost and we recommend that you hide the recovery files to resume data updates. Follow this procedure to hide recovery disk files from Manage Disk or from Manage Volume: 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the recovery disk. The element properties are now displayed on the left. 2. Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume. 3. The Windows Explorer window closes. You can resume data updates by clicking 'Update data' under Manage Volume. To copy the latest changes to the recovery disk automatically, change the update mode to continuous from the same area.
Master disk read-only (Recovery volumes only)	
Cause	Solution
Your computer was booted from the recovery disk using the option ROM, and the volume is	Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. We recommend that you hide the master disk files when finished with your review, and proceed with a data recovery to the master disk in order to maintain full redundancy. Follow this procedure to hide master disk files from



operating from that disk. The master disk files have been accessed and are displayed in Windows Explorer*.	Manage Disk: 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the master disk. The element properties are now displayed on the left. 2. Click 'Hide Files' from Manage Disk or 'Hide master disk files' from Manage Volume. 3. The Windows Explorer window closes and the volume is displayed as running off recovery disk. Refer to the 'Running off recovery disk' procedure above to recover data to the master disk.

Missing volume

Cause	Solution
A driver upgrade or downgrade was performed while a data migration was in progress.	The driver cannot recognize the volume or read its data if a driver upgrade or downgrade was performed during a volume migration. Volume migrations occur after one of the following operations was initiated: • Creation of a system volume or data volume while preserving data. • Volume type change combined with disk addition to the new RAID configuration. • Volume size increase. • Disk addition to an existing array. **Troubleshooting a data volume** 1. If the data migration involved a data volume, you will need to reverse the driver upgrade or downgrade operation and return to the original driver version. This will restore driver and volume compatibility. 2. Once the operation has completed, restart your computer. 3. Open the application and make sure that the volume displays again in the storage system view. Data migration operation should resume immediately. **Troubleshooting a system disk** If the data migration involved a system disk or volume, it is highly likely that you will not be able to start your system because the driver cannot read the system files. The following options may allow you to load the operating system again: • Restore a known good configuration. • Boot from a flash drive that supports NTFS partitioning and includes the storage driver files. • Bring the corrupt disk to another system, and then replace the storage driver files from a compatible driver version. Return the disk to the original system and try booting. **Troubleshooting a system volume** If the data migration involved a system disk or volume, it is highly likely that you will not be able to start your system because the driver cannot read the system files. The following options may allow you to load the operating system again: • Restore a known good configuration. • Bring all corrupted array disks to another system, and then replace the storage driver files from a compatible driver version. Return the disks to the original system and try booting.

22.4 Disk Events

State	Cause	Solution
Locked 	An internal or external disk is protected and locked with a password.	In this state, the overall storage system health is reported as normal, but to make the data fully accessible, you will need to follow this procedure to unlock the disk: 1. Under 'Status' or 'Manage', in the storage system view, click the disk you want to unlock. The disk properties are now displayed on the left. 2. Click 'Unlock'. 3. Enter the password, and then click 'Unlock'.
At risk 	An impending error condition was detected on an internal or external disk and is now at risk of failure.	The application is detecting early warning signs of failure with a SATA disk that result from a slow degradation over time. When a disk is reported at risk, you can reset that disk to normal, but we recommend that you contact the manufacturer for more information to prevent potential data loss. Follow this procedure to reset the disk to normal: 1. Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. **Note:** Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss. If the disk reported at risk is included in a RAID volume and a compatible spare disk is available, the rebuild process will start automatically. Once complete, the disk reported at risk becomes available and you can reset it to normal to return to a healthy state.
	An unexpected error was detected on a disk that has RAID configuration data (metadata) on it.	In this state, it is likely that some or all of the disk data is inaccessible. After backing up any accessible data, you will need to clear the metadata and reset the disk to return to a normal state. **Warning**: Completing this action will permanently delete existing metadata. Back up any accessible data before continuing. 1. Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2. Click 'Clear and reset disk', and then click 'Yes' to confirm. 3. Once complete, the page refreshes with the disk returning to a normal state.
Missing 	An array disk is not present or physically connected to the computer.	Ensure that the disk is securely connected to the SATA port and that the SATA cable is functioning properly. If the disk is lost or cannot be reconnected, you will need to connect a new SATA disk, and then rebuild the volume to that new disk. Refer to Degraded or Failed Volumes in this section for instructions on how to rebuild a volume.

	The recovery or master disk files have been accessed and display in Windows Explorer*.	Hide the recovery or master disk files to return the disk status to offline and resume data updates in on request mode. 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume or the recovery disk. The element properties are now displayed on the left. 2. Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume. 3. The Windows Explorer window closes.
Failed 	An internal or external disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.	Back up your data and we recommend that you replace the disk as soon as possible. If the failed disk is an array disk, the volume will be reported as degraded or failed depending on its configuration. Refer to Degraded or Failed Volumes in this section for instructions on resolving the problem. In a failed state, disk data may be lost, but you can try resetting the disk to normal, and then attempt a data recovery. Follow this procedure to reset the failed disk to normal: 1. Under 'Status', in the Manage subsection, locate the disk reported as failed. You can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state. **Note** If the failed array disk is part of a redundant volume, the volume will start rebuilding automatically as soon as the disk is reset to normal.
Offline 	An internal or external array disk is locked and data on that disk cannot be read.	We recommend that you unlock the disk to make the volume data fully accessible. If more than one array disk is locked, unlock all those disks to unlock the volume.
	The recovery volume is in on request update mode.	Change the volume update mode to continuous. 1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left. 2. Click 'Change mode', and then click 'Yes' to confirm. 3. The page refreshes and the volume properties report the new update mode.
	Your computer is running on battery and data updates to the recovery disk are not available as long as that disk is offline.	Reconnect your computer to the power supply in order to return the recovery disk to a normal state.

22.5 Caching Issues

Cache volume is missing
Regardless of what the acceleration mode currently is, cache and volume data is most likely lost. Devices that



are part of the acceleration configuration display in the following states:

- Solid state disk: Inaccessible
- Cache volume: No longer displays
- Data volume (if applicable): No longer displays
- Accelerated volume (if applicable): No longer displays
- Accelerated disk (if applicable): Offline
- Array disks: Offline

Cause

The solid state disk was removed or the disk is present but cannot be detected.

Solution

The application provides the option to clear the metadata on the array disks or previously accelerated disk and reset these disks to a normal state.

1. Under Status, in the Manage subsection, click 'Clear and reset' next to each array disk reported as offline. You can also perform this action under 'Manage' by clicking any offline disk reported in the storage system view.
2. Click 'Yes' to confirm.
3. The array disk now displays as an available disk in a normal state and can be used to create a new volume

Cache volume is failing

Cause

An impending error condition (e.g., SMART events) was detected on the solid state disk that is used as a cache device. As a result, both the disk and cache volume are at risk of failure.

Solution

Early warning signs of failure with the solid state disk are detected that result from a slow degradation over time. When a disk used as a cache device is reported at risk, you can reset that disk to normal or replace the solid state disk after resetting it to available.
Regardless of which option you choose, we recommend that you contact the manufacturer for more information to prevent potential data loss.

Current acceleration mode: Enhanced
Follow this procedure to reset the disk to normal:

1. Under 'Status', in the Manage subsection, locate the disk reported as at risk. You can also perform this action from Manage Disk, which is accessible by clicking the failing disk in the storage system view.
2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.
3. The cache volume should also return to a normal state and caching activity should resume.

Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss.

Follow this procedure to replace the failing solid state disk:

1. Click 'Accelerate', and then click 'Disable acceleration'.
2. In the dialog, click 'Yes' to confirm.
3. The page refreshes and reports the acceleration as disabled.
4. Click 'Reset to available'.
5. In the dialog, select the check box to confirm that you understand that data on the cache and data volumes will be deleted.
6. Click 'Yes' to confirm.
7. The page refreshes and the storage system displays the solid state disk usage as available.
8. Power off your computer and replace the failing solid state disk with a healthy one.

9. Power on your computer. You can enable acceleration again in order to resume caching activity.

Note
If the last acceleration mode was maximized, that application will automatically transition acceleration to enhanced mode in order to avoid data loss. While transitioning, the mode will display as busy and no acceleration actions will be available until the process is complete.

Current acceleration mode: Maximized

1. If a compatible spare is detected, the volume rebuild operation will start automatically. Once the process is complete, the cache volume will display in a normal state and caching activity will resume.
2. If no compatible spare is detected, the acceleration mode will automatically transition to enhanced in order to avoid data loss. You can then follow the procedures described above to return the solid state disk and cache volume to normal.

Cache volume has failed

Cause

The solid-state disk that is used as a cache device has failed to properly complete read and write operations in a timely manner and it has exceeded its recoverable error threshold. In this state, both cache and data volumes are reported as failed and acceleration is automatically disabled.

Solution

Back up any recoverable data and replace the solid- state disk as soon as possible. In a failed state, disk data may be lost, but you can try recovering it by resetting the disk to normal.

1. In the Manage subsection, under 'Status', locate the disk reported as failed. Alternately, perform this action from Manage Disk, accessible by clicking the disk in the storage system view.
2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.

If the disk operations continue to fail, the disk will return to a failed state immediately and should be replaced. Follow this procedure:

1. Click 'Accelerate'.
2. Click 'Reset to available'.
3. In the dialog, select the check box to confirm that you understand that data on the cache and data volumes will be deleted.
4. Click 'Yes' to confirm.
5. The page refreshes and the storage system displays the solid state disk usage as available.
6. Power off your computer and replace the failed solid state disk with an operational one.
7. Power on your computer. To resume the caching activity, enable acceleration again.

If acceleration was in maximized mode prior to being automatically disabled, the disk or volume previously associated with the cache will be reported as failed if the data cleaning was unsuccessful.
If data cleaning was successful, once the mode transition is complete, the accelerated disk or volume previously associated with the cache will be reported as normal.

Accelerated disk or volume is missing

Cause

The accelerated disk or volume cannot be detected (e.g., the device was moved to another system) and is reported as inaccessible. Caching activity is no longer occurring.

Solution

If the disk or volume can be reconnected:



1.Power off your computer and reconnect the missing disk or volume.
2.Restart your computer.
3.Once the operating system is running, open the application.
4.Under 'Status', in the Accelerate subsection, verify the accelerated device is properly reported as well as the acceleration mode. Caching activity should resume immediately.

If the disk or volume cannot be reconnected, follow this procedure to disassociate the cache and the missing device:

1.Click 'Accelerate'.
2.Click 'Disassociate'.
3.Click 'Yes' to confirm.
4. The page refreshes and you can now select another disk or volume to accelerate.

Accelerated disk or volume is reporting an issue
Cause • Accelerated disk is reported as being at risk or failed. • Accelerated volume is reported as degraded or failed.
Solution Refer to Troubleshooting Disk Events, Failed Volumes, or Degraded Volumes for detailed procedure on fixing the issue. If you cannot fix the issue reported on the accelerated disk or volume, follow this procedure to disassociate the cache and the missing device: 1.Click 'Accelerate'. 2.Click 'Disassociate'. 3.Click 'Yes' to confirm. 4. The page refreshes and you can now select another disk or volume to accelerate.

22.6 Software Errors

Message	Cause	Solution
An unknown error occurred while running this application. If the problem persists, please restart your computer or try reinstalling the application.	This error may be related to: sing components rrupted application plication unable to connect to the service plication fails to start.	Restart your computer or try reinstalling the application.
Intel® Rapid Storage Technology is trying to connect to the service.	The application is launched and is attempting to connect to the service in order to run.	If the connection succeeds, the application opens and is fully functional; if the connection fails, the error message described above is displayed. Try starting the service manually using Microsoft Windows* Services, or follow the recommended solution listed above to resolve the problem.
The Intel® Rapid Storage Technology service cannot be started in safe mode.	Your computer was started in safe mode and the operating system is running with a limited set of files and drivers. Intel Rapid Storage Technology cannot start or	Once you are done troubleshooting application or driver problems in safe mode, you will need to exit safe mode, restart your computer, and

	run in safe mode.	then let the operating system start normally. The Intel Rapid Storage Technology service can now be started and open the application.
Multiple users cannot run the application at the same time.	One or more users are attempting to open the application while an instance of the application is already running.	Make sure only one instance of the application is running at a time.
An error occurred due to insufficient resources, and the operation could not be completed. Please try again later.	The Intel® Rapid Storage Technology driver does not have sufficient resources to execute the request. Another operation may be in progress and needs to complete before being able to handle a new request.	Wait a few moments, then try performing the action again.
An unknown error occurred during the volume creation process. Please try recreating the volume.	An unexpected error occurred during the operation, and the application cannot identify its origin. The volume could not be created.	Verify that your hardware is properly connected and try recreating the volume.
An error occurred while an operation was in progress. The operation could not be completed.	An unexpected error occurred during an operation, such as a data migration or a rebuild, and the application cannot identify its origin.	Restart the operation. If the error persists, try restarting your computer and then the operation.
An error occurred and the selected disk or volume could not be accelerated. Please restart your computer, and then try the operation again.	The cache memory allocation was likely increased to use full solid state disk capacity (up to 64 GB) while enabling acceleration.	Follow these steps to accelerate a disk or volume: • Restart your computer to complete the process of allocating the requested cache size. • Launch the application. • Try enabling acceleration again by clicking 'Enable acceleration'.



23 Appendix A: Common Storage Management Interface Support (CSMI)

Intel® RST driver does not support application development that require interface access to the driver via API method. However RST does support applications developed to interface to the driver via the industry standard specification of the Common Storage Management Interface. We support a subset of the total command set. The below table has a list of the commands that are supported by Intel® RST. For more detail information on the specification you can access the http://www.t11.org/ website. The document number for the specification is 04-468v0.

Supported CSMI commands:

CC_CSMI_SAS_GET_DRIVER_INFO
CC_CSMI_SAS_GET_CNTLR_CONFIG
CC_CSMI_SAS_GET_CNTLR_STATUS
CC_CSMI_SAS_GET_RAID_INFO
CC_CSMI_SAS_GET_RAID_CONFIG
CC_CSMI_SAS_GET_PHY_INFO
CC_CSMI_SAS_GET_SATA_SIGNATURE
CC_CSMI_SAS_STP_PASSTHRU