

# Intel<sup>®</sup> Rapid Storage Technology

## OEM Technical Guide

---

*For the Intel<sup>®</sup> Rapid Storage Technology 9.5 Software Package*

*Revision 0.81*

**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 Intel® ICH9M/M-E, ICH10R 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 © 2009, Intel Corporation. All rights reserved.



# Contents

---

1	About This Document.....	8
1.1	Purpose and scope of this document.....	8
2	Intel® Rapid Storage Technology.....	9
2.1	Overview of RAID Levels.....	9
2.2	Typical usage model for RAID levels.....	11
3	Intel® Rapid Storage Technology Suite.....	12
3.1	Intel® Rapid Storage Technology 9.5 software.....	12
3.2	Intel® Rapid Storage Technology Option ROM.....	12
3.3	Intel RAID Configuration Utility.....	13
3.3.1	RAIDCfgr Utility for MS-DOS*.....	13
3.4	RAIDCmd Utility .....	14
4	Creating a RAID Volume.....	16
4.1	Using the Intel Rapid Storage Technology Console.....	16
4.2	Using the Intel Rapid Storage Technology option ROM User Interface.....	17
4.3	Using the RAID Configuration Utility.....	17
5	Deleting a RAID Volume.....	19
5.1	Using the Windows User Interface Utility.....	19
5.2	Using the Option ROM User Interface.....	19
5.3	Using the RAID Configuration Utility.....	20
6	Common RAID Setup Procedures .....	21
6.1	Build a SATA RAID 0, 1, 5 or 10 System.....	21
6.2	Build a SATA "RAID Ready" System.....	22
6.3	Migrate to RAID 0 or RAID 1 on an Existing "RAID Ready" System.....	23
6.4	Migrate an Existing Data Hard Drive to a RAID 0 or RAID 1 Volume.....	23
6.5	Migrating from one RAID volume type to another volume type.....	24
6.6	Create a RAID volume on SATA while booting to PATA.....	25
6.7	Build a RAID 0 or RAID 1 System in an Automated Factory Environment .....	26
6.7.1	Part 1: Create the Master Image.....	26
6.7.2	Part 2: Apply the Master Image.....	26
7	RAID Volume Data Verification and Repair Feature.....	27
7.1	Verify and Repair Volume Feature.....	27
8	Intel® Rapid Recover Technology.....	28
8.1	Creating a Recovery Volume through the RAID Option ROM.....	28
8.2	Creating a Recovery Volume through the Intel® Rapid Storage Technology Console.....	29
8.3	Changing Recovery volume modes.....	29
8.4	Update Recovery Volume in On Request Update Policy.....	30



8.5	<a href="#">Access Recovery Drive Files</a>	30
8.6	<a href="#">Hide Recovery drive files</a>	31
8.7	<a href="#">Scenarios of Recovering Data</a>	31
9	<a href="#">System Running off Recovery Drive</a>	34
9.1	<a href="#">Drive Offline or Missing</a>	34
10	<a href="#">Pre-Installation of the Intel® Rapid Storage Technology Driver</a>	35
10.1	<a href="#">Pre-Installation Using the "F6" Method</a>	35
10.2	<a href="#">Unattended Installations of Microsoft Windows XP</a>	36
10.3	<a href="#">Creating a Floppy Disk containing the Intel Rapid Storage Technology Driver</a>	37
10.3.1	<a href="#">Method 1 – Using Install applications</a>	37
10.3.2	<a href="#">Method 2: Using F6 application</a>	38
11	<a href="#">Determining the version of the RAID driver</a>	39
11.1	<a href="#">Using Intel® Rapid Storage Technology User Interface (UI)</a>	39
11.2	<a href="#">Using Intel Rapid Storage Technology File Properties</a>	39
11.3	<a href="#">Determining the Version of the Option ROM</a>	40
11.3.1	<a href="#">Using the Intel Rapid Storage Technology UI</a>	40
11.3.2	<a href="#">Using the Intel Rapid Storage Technology Option ROM User Interface</a>	40
12	<a href="#">Uninstallation</a>	41
12.1	<a href="#">Uninstalling the Intel Rapid Storage Technology Software (except the RAID Driver)</a>	41
12.2	<a href="#">Disabling the RAID Driver by Disabling the RAID Controller</a>	41
13	<a href="#">Registry Customizations</a>	43
13.1	<a href="#">Tray Icon Configuration</a>	43
13.2	<a href="#">Customize Support URLs in Rapid Storage Technology Console</a>	43
13.3	<a href="#">Disable the pop-up Window with BIOS incompatibility Warning</a>	44
13.4	<a href="#">Native Command Queuing Settings</a>	45
14	<a href="#">HDD password support with RAID volumes</a>	46
14.1	<a href="#">HDD Password Use cases</a>	46
15	<a href="#">Port multiplier (PMP) support with Intel® Rapid Storage Technology</a>	48
16	<a href="#">Intel® Rapid Storage Technology UI</a>	49
16.1	<a href="#">Introduction</a>	49
16.1.1	<a href="#">Getting started</a>	50
16.1.2	<a href="#">Understanding the application</a>	51
16.1.3	<a href="#">Notification area</a>	53
16.2	<a href="#">Storage System Status</a>	53
16.2.1	<a href="#">Understanding the status</a>	54
16.2.2	<a href="#">Storage system view</a>	55
16.3	<a href="#">Creating a Volume</a>	56
	<a href="#">Volume requirements</a>	56
16.3.1	<a href="#">Creation process</a>	57
	<a href="#">Creating Additional Volumes</a>	61
16.4	<a href="#">Managing the Storage System</a>	63
16.4.1	<a href="#">Managing arrays</a>	63
16.4.2	<a href="#">Managing volumes</a>	66



16.4.3	Managing disks.....	77
	Managing ports.....	82
	Managing ATAPI devices.....	82
16.5	Accelerating the Storage System.....	83
	Configuring flash memory.....	83
16.6	Preferences.....	86
16.7	Troubleshooting.....	86
16.7.1	Failed volumes.....	87
16.7.2	Degraded volumes.....	88
16.7.3	Other volume states.....	91
16.7.4	Disk events.....	93
16.7.5	Flash memory states.....	95
16.7.6	Software errors.....	98
17	Braidwood Technology.....	100
17.1	Braidwood Platform.....	100
17.2	Benefits.....	102
17.3	Usage Models.....	103
17.3.1	Cache only (total capacity).....	103
17.4	Braidwood Implementation.....	103
17.4.1	Factory Tools and Process.....	103
17.5	Caching Implementation.....	115
17.5.1	Caching Requirements.....	115
17.5.2	Configuring the NAND module for Caching.....	115
17.6	Caching Features and Configurations.....	115
17.6.1	Configuration Rules.....	115
17.6.2	Cache Modes.....	116
17.7	OROM Implementation Supporting Braidwood.....	118
17.7.1	Core OROM Features and Functionality.....	118
17.7.2	OROM UI Features and Functionality.....	122
18	Glossary.....	133



## *Revision history*

---

<b>Revision Number</b>	<b>Description</b>	<b>Revision Date</b>
0.5	Initial Draft Version	June 2009
0.81	Edits based upon peer review feedback	July 2009



# ***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 9.5* software for the chipset components as listed in the product's Readme.txt file.

This document also describes installation procedures, 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 to enable an optimal PC storage solution. It offers value-add features such as RAID and advanced Serial ATA\* capabilities for the Microsoft Vista\*, Microsoft Windows\* XP, Microsoft Windows Server 2003 and Microsoft Windows Server 2008 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 NV memory HW module.

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).

A configuration supporting two RAID levels can also be achieved by having two volumes in a single RAID array that use Intel® RST. Targeted for desktops and workstations, this 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 or more hard 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® ICH9 I/O Controller and later chipsets (all desktop) 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.

**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 or four hard 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  $(N-1) * \text{smallest size of the hard drives}$ , where  $N \geq 3$  and  $\leq 4$ .

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

- |                |   |
|----------------|---|
| <b>RAID 0</b>  | This provides end-users the performance necessary for any disk-intensive applications; these include video production and editing, image editing, and gaming applications.                |
| <b>RAID 1</b>  | This provides end-users with data redundancy by mirroring data between the hard drives.   |
| <b>RAID 5</b>  | 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. |
| <b>RAID 10</b> | 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 components:

1. Intel® Rapid Storage Technology 9.5 software
  - a. AHCI/RAID driver
  - b. Graphical User Interface
  - c. Monitor service
2. Intel® Rapid Storage Technology Option ROM
3. Intel® Rapid Storage Technology RAID utilities
  - a. Intel® RAIDCmd utility
  - b. Intel® RAIDCfg utility
  - c. Intel® RAIDCply utility

### 3.1 **Intel® Rapid Storage Technology 9.5 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 driver for supported Windows\* operating systems. that support RAID for Windows XP and provides the Intel® Rapid Storage Technology Console (a 32-bit Windows application). The driver supports the Intel® 82801IR/IO Serial ATA RAID Controller and will recognize its 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 Rapid Storage Technology Console (GUI), provides RAID volume management (create, delete, migrate) 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. The uncompressed option ROM image will be less than 64 KB in size. The run-time image that resides in BIOS compatibility space will be less than 20 KB in size and consumes no more than 12 KB of system memory. 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 has two main modes. The first one 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.

The second mode of operation is console mode (applicable for RAIDCfg only). If this utility is executed without any command line flags, a console interface identical to that of the Intel Rapid Storage Technology option ROM is presented and is fully functional within the DOS environment. The mode requires user interaction; however, it may be used to create, delete, and manage RAID volumes from a DOS environment when the Intel Rapid Storage Technology option ROM is unavailable.

#### 3.3.1 RAIDCfg Utility for MS-DOS\*

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

```
=====
RaidCfg.exe [/?] [/Y] [/Q] [/C:vol_name] [/SS:strip_size] [/L:raid_level]
           [/S:vol_size] [/DS:disk_ports] [/D:vol_name] [/X] [/I] [/P]
[/ST] [/STD] [/STV] [/F:vol_name] [CnG:vol_name] [/Sync]

/?      Displays Help Screen.  Other options ignored.
/Y      Suppress any user input.  Used with options /C, /D, /X, & /F.
```



```

/Q      Quiet mode / No output. Should not be used with status commands.
        COMMANDS - Only one at a time.
/C      Create a volume with the specified name.
        /S, /DS, /SS, & /L can be specified along with /C.
/SS     Specify strip size in KB. Only valid with /C
/L      Specify RAID Level (0, 1, 10, or 5). Only valid with /C
/S      Specify volume size in GB or percentage if a '%' is appended.
        Percentage must be between 1-100. Only valid with /C
/DS     Selects the disks to be used in the creation of volume.
        List should be delimited by spaces.
/D      Delete Volume with specified name.
/X      Remove all metadata from all disks. Use with /DS to delete
        metadata from selected disks.
/I      Display All Drive/Volume/Array Information. /P can be specified.
/P      Pause display between sections. Only valid with /I.
/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
/F      Repair failed RAID0 Volume.
/CnG    Create an Intel® Rapid Recover Technology volume. To be used with /
C and /DS.
/Sync   Set sync type for Intel® Rapid Recover Technology volume. Only
valid with /CnG
/M      Choose port number of the Master disk for Intel® Rapid Recover
Technology volume. Only valid with /CnG
=====

```

### 3.4 RAIDCmd Utility

The Intel RAIDCmd utility is an executable It provides customers 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 RAIDCmd utility is shown below:

```

=====
RAIDCmd.exe    [/?] [/h] [/FN:fileName] [/P] [/Y] [/Q]
               [/RD] [/ST] [/STA] [/STD] [/STV] [/I] [/SB:MB|GB|TB]
               [/C:volumeName] [/D:volumeName] [/V:volumeName]
               [/SS:StripSize] [/L:RaidLevel] [/S:volumeSize]
               [/DS:diskPorts] [/IRRT:volumeName] [/R:volumeName]
               [/M:diskPort] [/F:volumeName] [/SDR:DiskList]
               [/Spare:DiskList] [/Sync:A|M|volumeName]
               [/X] [/PL:portList]

? or /H       Displays help screen. Other options are ignored except
option /p.
/Y           Suppress any user input.
/Q           Quiet mode; no information is displayed onscreen.
             Do not use with status commands.
/P           Pauses display between sections.
/FN         Specifies the file name where screen information will be
saved.

```



```
/RD          Inquires if the RAID driver is running.
/C          Creates a volume with the specified name.
           /S, /DS, /SS, & /L must be specified along with /C
           /V can be specified along with /C
/D          Deletes a volume with the specified name.
/V          Specifies the volume name that the command will be applied
to
           (e.g. create a volume, etc.).
/SS         Specifies the volume strip size in KB. Only valid with
/C.
/L          Specifies the volume RAID level (0, 1, 5, 10 or IRRT).
           Only valid with /C .
/S          Specifies volume size in GB or percentage if a '%' is
appended. Percentage
           must be between 1 and 100. TB, GB or MB can be also be
appended
           to specify size in Tera, Giga or Mega bytes. Only valid
with /C.
/DS         Selects the disks to be used with the command.
           List of disks should be delimited by spaces.
/IRRT       Specifies the update mode of the recovery volume.
/M          Specifies the port number of the master disk for the
recovery volume.
           Valid with /IRRT (recovery volumes) only.
/ST         Displays Volume/RAID/Disk Status.
/STA        Displays delimited Array Status
           Index,Size,FreeSize,Status,DiskCount,VolCount,Name
/STD        Displays delimited Disks Status
           Port,Model,SerialNumber,FirmwareVersion,Array,Status,
Size,Free,Type
/STV        Displays delimited Volume Status
           Index,Level,StripSize,Size,Status,Bootable,Array,Name
/F          Repairs a failed RAID 0 Volume.
/I          Displays All Disk/Volume/Array Information.
/SDR        Converts a RAID disk to non RAID. Use the disk list as the
argument.
/SPARE      Marks a disk as spare.
/R          Rebuilds a degraded volume or array.
/PL         Enables Port LED functionality to easily locate the port
light.
           on the system hardware. Port must be equipped with LED
/SB         Specifies the display base (MB|GB|TB)
/Sync       Sets sync type for the IRRT volume.
           Valid with /IRRT (recovery volumes) only.
           It can be used to trigger manual volume synchronization on
given volume
           when is used standalone with recovery volume name.
/X          Removes metadata from all RAID disks.
           Use with /DS to delete metadata from specific RAID disks.
           If you specify one RAID disk from an existing volume,
           metadata will be removed from all disks belonging to that
volume.
=====
```



## 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 Console 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 RAID Configuration utility.

### 4.1 **Using the Intel Rapid Storage Technology Console**

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

Start→Programs→Intel Control Center ->Intel® Rapid Storage Technology→Intel Rapid Storage Console

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... (For EG: if you have only two volumes ...you can only see options to create RAID 0,RAID1 and Recovery(IRRT) ; 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 be 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 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

Run "raidcfg.exe" (DOS environment) or "raidcmd.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:\>raidcfg.exe /C:OEMRAID0 /DS:0 1 /SS:128 /L:0 /S:120  
C:\>raidcmd.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:\>raidcfg.exe /C:OEMRAID0
```



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

```
C:\>raidcfg.exe /?  
C:\>raidcmd.exe /h
```

**Note:**

Selecting the strip size is only applicable for RAID 0, RAID 5, RAID 10 levels, but not 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 Console from the following Start menu link:  
  
Start→All Programs→Intel® Rapid Storage Technology →Intel Rapid Storage Technology Console
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 "raidcfg.exe" (DOS environment) or "raidcmd.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:\>raidcfg.exe /D:OEMRAID0  
C:\>raidcmd.exe /D:OEMRAID0
```

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

```
C:\>raidcfg.exe /?  
  
C:\>raidcmd.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 striped and mirrored for RAID 10 across two or up to four SATA hard drives. 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.





6. Install the Intel Rapid Storage Technology software via the CD-ROM included with your motherboard or after downloading it from Intel's website on the Internet. This will add the Rapid Storage Technology console that can be used to manage the RAID configuration and to use the Intel Rapid Storage Technology migration feature to migrate a "RAID Ready" configuration to a RAID 0, RAID 1, RAID 5, RAID 10 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, 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. 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 Console and start menu links.
4. Open the Intel Rapid Storage Technology console 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 Console and start menu links.
4. Open the Intel Rapid Storage Technology Console from the Start Menu.
5. Follow steps 4 to 7 in section 6.3

## 6.5 Migrating from one RAID volume type to another volume type

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 Console application:  
Start Menu -> All Programs -> Intel Rapid Storage Technology -> Intel Rapid Storage Technology
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 Console and Start menu links.
4. Use the Intel Rapid Storage Technology Console 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 9.0* software from the CD-ROM included with your motherboard or after downloading it from the Internet. This will add the *Intel Rapid Storage Technology Console* 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 (raidcfg.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:\>raidcfg.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:\>raidcfg.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 start 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*

---

Intel® Rapid Recover Technology is supported on ICH8M-E and greater mobile platforms and ICH9R/DO and greater desktop platforms. This technology utilizes RAID 1 functionality to copy data from a designated Master drive to a designated Recovery drive. The size of the Master drive must be less than or equal to the size of the Recovery drive. 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 Console 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® Rapid Storage Technology Console

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



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 console'.
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 console.
4. Under 'Status' or 'Manage', in the storage system view, click the recovery volume to be rebuild. 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 Console from the Start Menu.
4. Under 'Status' or 'Manage', in the storage system view, click the recovery volume to be rebuild. 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 Console 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

```
Intel<R> Matrix Storage Manager option ROM v8.0.0.8888 ICH9R wRAID5
Copyright<C> 2003-07 Intel Corporation. All Rights Reserved.

[ MAIN MENU ]
1. Create RAID Volume          3. Reset Disks to Non-RAID
2. Delete RAID Volume         4. Recovery Volume Options
5. Exit

[ DISK/VOLUME INFORMATION ]

RAID Volumes:
ID   Name      Level          Strip  Size   Status      Bootable
0    IRR1      Recovery(OnReq) N/A    57.3GB NeedsUpdate Yes

Physical Disks:
Port Drive Model      Serial #      Size   Type/Status(Uo1 ID)
0    Maxtor 6Y060M0  Y2NGJ4FE     57.3GB Offline Disk(0)
1    Maxtor 6Y080M0  Y2R1Z18E     76.3GB Member Disk(0)
2    Maxtor 6Y200M0  Y60MQ3RE    189.9GB Non-RAID Disk
3    Maxtor 6Y080M0  SN0000      76.3GB Non-RAID Disk
4    Maxtor 6Y200M0  SN0001     189.9GB Non-RAID Disk
5    Maxtor 6Y080M0  SN0002      76.3GB Non-RAID Disk

[↑↓]-Select      [ESC]-Exit      [ENTER]-Select Menu
```

Then Select Option 2 Enable Only Recovery Disk.



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

---

The Intel Rapid Storage Technology driver must be installed before installing OS on a RAID volume or when in AHCI mode. Intel Rapid Storage Technology AHCI driver can be installed over Vista's native AHCI driver.

## 10.1 Pre-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 you can use Floppy, CD/DVD or USB.
  - Driver binary (iastor.sys),
  - INF files (iastor.inf and iaAhci.inf),
  - cat files (iaStor.cat and iaAhci.cat) and
  - TxtSetup OEM file (txtsetup.oem).

These should be in the root directory of the floppy. See Section 7.3 for instructions on making a driver installation disk.

### **For Windows Vista:**

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, CD/DVD Or USB) you created in step 3 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

[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 Rapid Storage Technology 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, 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:
6. Using Setup.exe – English version:
  - a. C:\>setup.exe -a -p c:\IAAdriver
7. Using iata\_enu.exe - English version (Alternate Method):
  - a. C:\>iata\_enu.exe -a -a -p c:\IAAdriver
8. Using iata\_cd - Multi-language version:
  - a. C:\>iata\_cd.exe -a -a -p c:\IAAdriver

\*\*Where C:\IAAdriver is an already existing folder on the c: drive
9. 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 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 floppy.
2. Now your floppy disk will have the valid Intel Rapid Storage Technology Driver which can be used for F6 method. Determining Software Component versions



# **11 Determining the version of the RAID driver**

---

There are two ways to do this. The first is to use the Intel Rapid Storage Technology. It reports the current version of the driver installed. The second is to locate the driver (iaStor.sys) itself and view its properties.

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

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. Now click on 'Intel® Rapid Storage Technology' link to expand the item. Under it you can view the driver version in the following format: 9.5.0.xxxx
5. This is the current version of the user interface utility installed on your system. The 9.5.0 portion is the product release number; the xxxx portion is the build number.

## **11.2 Using Intel Rapid Storage Technology File Properties**

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 Console



## **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 Rapid Storage Technology 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 Uninstallation

---

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 ICH9 I/O controller. Therefore, Intel does not provide a way to permanently remove the driver from the system. However, disabling the 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 Manager 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 Rapid Storage Technology 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

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. You should no longer see the RAID Option ROM status screen during boot, and you should no longer see the ICH9 I/O 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 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

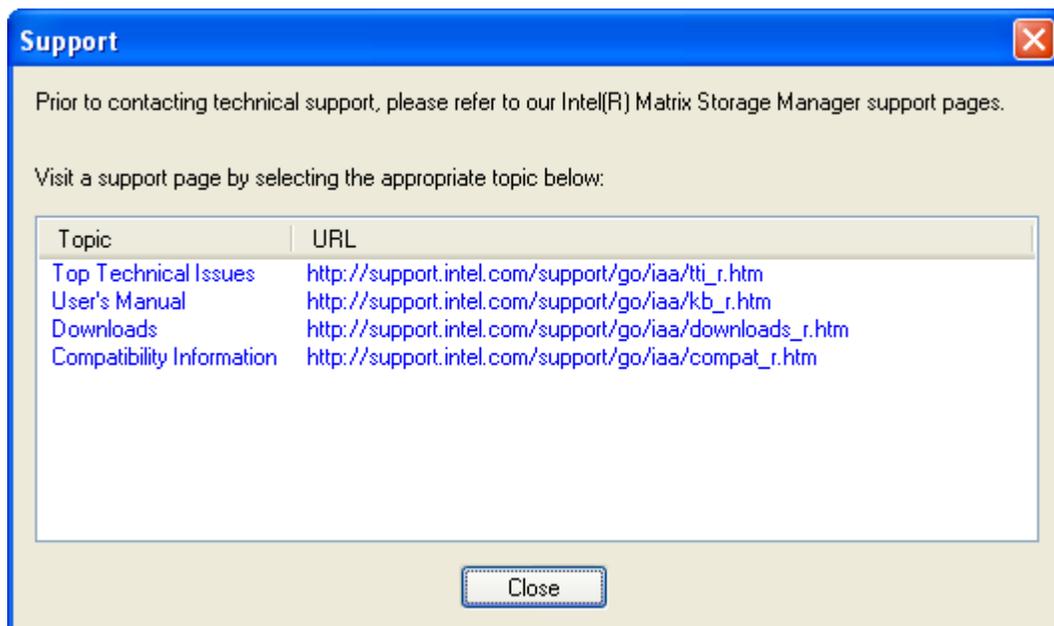
After installation of the Intel Rapid Storage Technology, the registry will contain keys to allow customization of several features. The registry key used to customize the Intel Rapid Storage Technology is at the following path:

HKEY\_LOCAL\_MACHINE/SOFTWARE/Intel/Intel Matrix Storage Manager

### 13.1 Tray Icon Configuration

### 13.2 Customize Support URLs in Rapid Storage Technology Console

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



These URLs can be changed to be OEM-specific by following the procedure listed below:

#### Method 1: Manual Registry

1. Run Regedit.exe from command prompt.



2. Select "HKEY\_LOCAL\_MACHINE\SOFTWARE\Intel\Intel Matrix Storage Console" key.
3. Add STRING values in the following form:  
ResourceN                    REG\_SZ    myurl  
ResourceNURL                REG\_SZ    <http://www.myurl.com>

Where N is any value from 0-49.

### Method 2: Using SetupCfg.reg

Edit the setcfg.reg file and append the following lines:

```
[HKEY_LOCAL_MACHINE\SOFTWARE\Intel\Intel Matrix Storage Console]
"Resource0URL"="http://www.xxx.com"
"Resource0"="Test xxx"
"Resource1"="Test yyy"
"Resource1URL"="http://www.yyy.com"
"Resource3"="Test zzz"
"Resource3URL"="http://www.zzz.com"
"Resource4"="Test aaa"
"Resource4URL"="http://www.aaa.com"
```

## 13.3 Disable the pop-up Window with BIOS incompatibility Warning

When a window warning is displayed in Intel Rapid Storage Technology, as shown in the figure below, follow the procedure given below if you intend to suppress the warning.



### Method 1:

Check the "Do not show this message again" box.

### Method 2: Registry entry

1. Run `regedit.exe` from the command prompt.
2. Select "HKEY\_LOCAL\_MACHINE\Software\Intel\Intel Matrix Storage Console" key.



3. Add a new string "PlugInRAID.pin\_CompatibilityMsg" with value to enable or disable:
  - "PlugInRAID.pin\_CompatibilityMsg"=dword:00000000
  - dword:00000000 - Disable
  - dword:00000001 - Enable
4. Exit the Regedit application.
5. Open Intel® Rapid Storage Technology from Start menu.
6. Confirm the above window is not displayed.

## 13.4 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)



# 14 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.

## 14.1 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.
IRRT 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 IRRT 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.



## **15 Port multiplier (PMP) support with Intel® Rapid Storage Technology**

---

Port multipliers are silicon-based devices that allow a single Serial ATA port to communicate with multiple disk drives. Intel® Rapid Storage Technology 9.5 supports port multipliers with PCHI/O controller to enable easy, cost effective storage scalability with SATA drives. Intel® Rapid Storage Technology supports only FIS based port multipliers plugged to either port 4 or port 5 only. Below are the limitations the users need to take into account when enabling port multipliers on their system

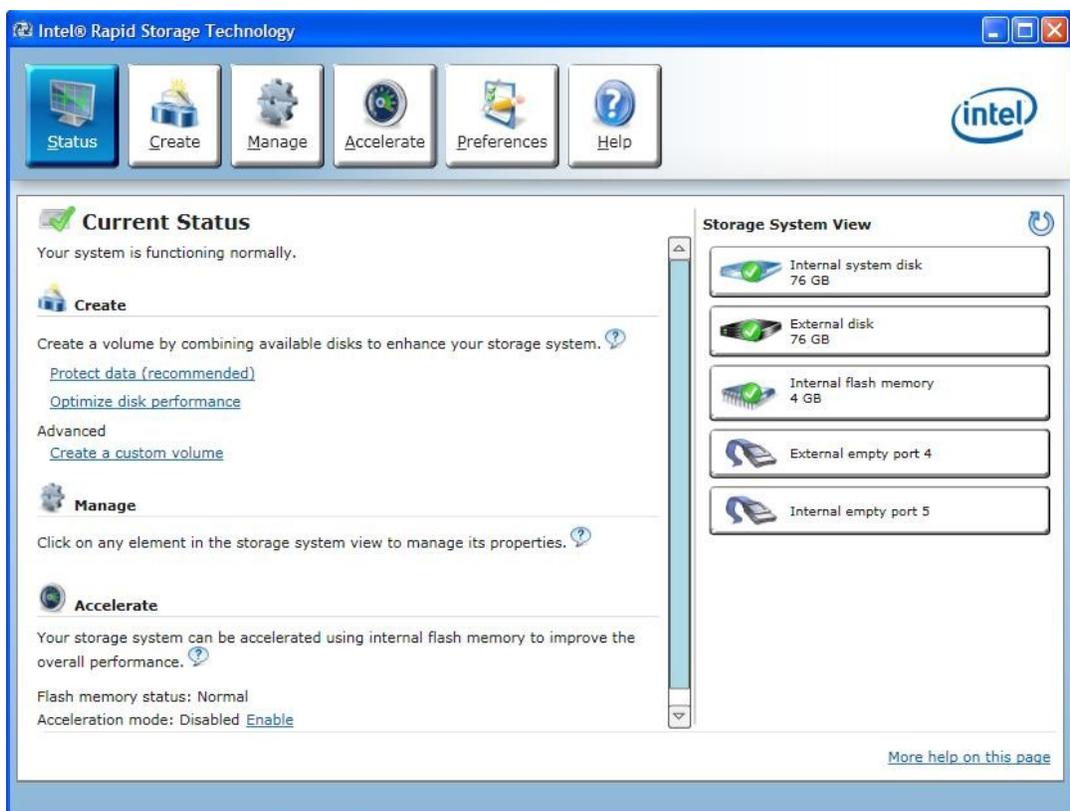
1. Port multipliers will only be supported on ports configured as eSATA ports (ESP bit set to '1'). Port multipliers that are attached to non eSATA ports will be treated as pass-through devices.
2. Only RAID data volumes that do not span across a port multiplier port and non-port multiplier port are supported on the port multiplier. Inhibits creating volumes that span across a port multiplier port and non-port multiplier port. Also doesn't support booting from a RAID volume or a pass-through drive on the port multiplier.
3. The OROM is PMP aware but does not support PMP accesses. It will ignore, not access and not display in its UI, any disks attached behind a PMP. But it will give visual indication in its UI that the port has a PMP attached.
4. Hot-plugs of PMP are handled without a manual HW scan. Hot-plugs of disks behind a PMP are handled only with a manual HW scan being ordered.
5. No LPM support on the port with port multiplier attached



# 16 Intel® Rapid Storage Technology UI

We live in a digital world. We digitally create, record, edit, share and save practically everything from the movies we watch, to the pictures we take, to the documents we store.

Is your PC optimized for storage performance and reliability?



## 16.1 Introduction

Intel® Rapid Storage Technology is a Windows\*-based application that provides improved performance and reliability for systems equipped with SATA disks for desktop, mobile, and server platforms. When using one or multiple SATA disks, you can take advantage of enhanced performance and lower power consumption. When using more than one disk, you can increase protection against data loss in the event of disk failure.

In this section:  
[Getting Started](#)



[Understanding the Application Notification Area](#)

### 16.1.1 Getting started

The Intel® Rapid Storage Technology software package provides high-performance SATA 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 Drives (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 Raid**  
This technology provides the ability to create RAID 0, RAID 1, RAID 5, and RAID 10 volumes. Data is distributed across two or more disks to provide data redundancy 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.
- **Flash memory configuration support**  
This feature allows you to improve the overall system performance by accelerating a disk, volume or applications and files.
- **Volume migration**  
This feature provides support for converting data into a high-performance RAID configuration.
- **Volume size increase**  
This feature allows you to increase the data storage capacity of a volume by utilizing 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.

#### 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.

- **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.

· **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.

· **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.

## 16.1.2 Understanding the application

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

In this section:

[Storage System Configuration](#)  
[Navigation](#)

### 16.1.2.1 Storage system configuration

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

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

· **Array**

An array is a collection of two or more SATA disks in a RAID configuration and is the highest element in the hierarchy of a storage system. Once a volume is created, the disks you used to create that volume form an array. Refer to the Creating Additional Volumes topic for details on how you can create two volumes across the same disks. An array can include one or two RAID volumes if the hardware allows it.

· **Volume**

A volume is the storage area on two or more disks whose type dictates the configuration of the data stored. If you created a volume for data protection, then your storage system may include a RAID 1 volume spanning two SATA disks, which mirrors data on each disk.

· **Disks**

A disk (i.e., hard disk or hard disk drive) physically stores data and allows read/write data access. If a disk is used to create a volume, it becomes an array disk because it has been grouped with other disks to form an array.

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



### 16.1.2.2 Navigation

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

<Insert> Status Icon

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.

<Insert> Create Icon

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.

<Insert> Manage Icon

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.

<Insert> Accelerate Icon

The 'Accelerate' area allows you to accelerate a disk, volume or specific applications and files using the internal flash memory, in order to improve the overall's system performance. Some of the benefits to such configuration are an increase in system responsiveness, multi-tasking enabling, and battery life extension.

<Insert> Note icon

Accelerate is only available if the the requirements listed in the Getting Started with Acceleration and Flash Memory Properties topics are met.

<Insert> Preferences Icon

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.



<Insert> Storage System View Icon

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.

### 16.1.3 Notification area

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

The notification area icon will automatically display in the notification area once Intel Rapid Storage Technology is installed.

#### Opening the application from the notification area

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

#### 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.

<Insert> Note icon

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.

## 16.2 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 you storage system.



## 16.2.1 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:

 <b>Normal</b>
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 an internal flash memory device is present and the requirements listed in the Accelerating the Storage System are met. This area typically reports the device health state and acceleration mode, and also provides the option to enable acceleration if it is disabled.
 <b>Warning</b>
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.
 The Accelerate subsection is only available if an internal flash memory device is present, at risk of failing, or has failed. Details and a recommended action are provided to help you fix the problem reported on the flash memory hardware.
 <b>Error</b>
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.



## 16.2.2 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.

<insert icon>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 disk is reported normal.	None
	An external disk is reported 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.	Back up your data and replace the disks as soon as possible. Refer to the Troubleshooting section for more information.
	An external disk is reported at risk.	Back up your data and replace the disks as soon as possible. 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.	<ul style="list-style-type: none"> <li>The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or,</li> <li>Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.</li> </ul>
	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.	<ul style="list-style-type: none"> <li>The recovery volume is in on request update mode. Change the volume update mode to continuous, if desired. Or,</li> <li>Your computer is running on battery and data updates to the recovery disk are not available. Reconnect your computer to the power supply.</li> </ul>
	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 disk is reported failed.	Refer to the Troubleshooting section for more information.
	An external disk is reported failed.	Refer to the Troubleshooting section for more information.



#### Volume states

Element	Description	Recommendation
	A volume is reported normal.	None
	A volume is reported degraded.	Refer to the Troubleshooting section for more information.
	A volume is reported failed.	Refer to the Troubleshooting section for more information.

#### Other storage system elements

Element	Description	Recommendation
	An internal flash memory device is present and functioning normally.	None
	An internal flash memory device is reported as at risk.	Refer to the Troubleshooting section for more information.
	An internal flash memory device is reported as failed.	Refer to the Troubleshooting section for more information.
	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

## 16.3 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.

<insert icon>Warning

Performing this action will permanently delete any existing data on the disks used to create a volume. Backup all valuable data before starting this process.

### 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, external or port multiplier ports. 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.

Volume type	Number of disks	Controller support
Recovery volume	2	ICH8M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3), PCHM.
RAID 0	2	ESB2, ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH8R, ICH8DH, ICH8DO, ICH8M-E, ICH8M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3), PCHM.
RAID 0	3 or 4	ESB2, ICH7R, ICH7DH, ICH8R, ICH8DH, ICH8DO, ICH8M, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3).
RAID 0	5 or 6	ESB2, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3).
RAID 1	2	ESB2, ICH7R, ICH7DH, ICH7MDH, ICH7M, ICH8R, ICH8DH, ICH8DO, ICH8M-E, ICH8M, ICH9R, ICH9DH, ICH9DO, ICH9M, ICH9M-E, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3), PCHM.
RAID 5	3 or 4	ESB2, ICH7R, ICH7DH, ICH8R, ICH8DH, ICH8DO, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3).
RAID 5	5 or 6	ESB2, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3).
RAID 10	4	ESB2, ICH7R, ICH7DH, ICH8R, ICH8DH, ICH8DO, ICH9R, ICH9DH, ICH9DO, ICH10R, ICH10D, ICH10DO, PCH (DT2/DT3).

### 16.3.1 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.



### 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'.

<b>Recovery volume: Flexible data protection</b>		<insert icon>
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, only changes to the master disk since the last update are copied to the recovery disk when you request it. The master and recovery disks must include 100% of the available disk space of an array, and no other volumes can be present on the system.		
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.	

<b>RAID 1: Real-time data protection</b>		<insert icon>
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.
-------------	---

<b>RAID 0: Optimized disk performance</b> Combines two to six SATA disks and breaks down data into units that are spread across the array disks.		<insert icon>
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. two-disk RAID 0 available in specific mobile configuration.	

<b>RAID 5: Efficient data hosting and protection</b> 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.		<insert icon>
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. Not available on mobile configurations.	

<b>RAID 10 : Balanced performance and data protection</b> 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.		<insert icon>
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. Not available on mobile configurations.

### **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.

### **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

<insert icon>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

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 you kept data from one of the disks.
6. If the selected disks did not have existing data or you chose not to keep existing data from one of the disks, you will need to partition the volume using Windows Disk Management\*.

<insert icon>Note

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

### Creating Additional Volumes

#### Creating multiple volumes on a single array

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. When you create the first volume, make sure to set the array allocation to be less than 100% in the Configure Volume step. The array may consist of two to six SATA disks, depending on the RAID type.

<inset icon>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.



## 16.4 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.

### 16.4.1 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.

#### 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 Gigabytes (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.

#### Adding a disk

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.

#### <inset icon>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.

#### Adding a Disk from Manage Array or Manage Volume



1. Under 'Status' or 'Manage', in the storage system view, click the array or 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'.
5. Restart your computer for changes to take effect, and then use Windows Disk Management\* to increase the capacity of the array and the volumes.

**<insert icon>Note**

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

**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. When you create the first volume, make sure to set the array allocation to be less than 100% in the Configure Volume step. The array may consist of two to six SATA disks, depending on the RAID type.

**<insert icon>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. 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.

### Increasing volume size

You can increase the size of a RAID volume by utilizing the remaining available space on the array. Hovering over the array name in the storage system view displays the amount of available space in MB.

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.

### 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.
4. Restart your computer for changes to take effect, and then use Windows Disk Management\* to utilize the available space.

### 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.
4. Restart your computer for changes to take effect, and then use Windows Disk Management\* to utilize the available space.

### <insert icon>Note

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

### Enabling disk data cache

Enabling the disk data cache for all SATA disks on the array allows you to allocate a portion of the system memory (RAM) 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.



Array properties report whether the data cache is enabled or disabled for all SATA disks in the array. Array disk properties report whether the data cache is enabled or disabled for the applicable disk, but does not provide the ability to change this setting.

**<insert icon>Warning**

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

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.

## 16.4.2 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.

### 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	<ul style="list-style-type: none"><li>· RAID 0 volume: indicates that one or both array disks are missing or have failed.</li><li>· RAID 1 volume: indicates that both array disks have failed.</li><li>· RAID 5 or 10 volume: two or more array disks are missing or have failed.</li></ul>
Inaccessible	Indicates that data on the accelerated volume cannot be accessed because the flash memory hardware it is associated with is missing, or the volume data does not match the flash memory data.
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	<ul style="list-style-type: none"> <li>· The recovery disk has failed, or</li> <li>· The master disk is missing or has failed and the volume is running off the recovery disk.</li> </ul>
Failed	Indicates that both array disks have failed.
Inaccessible	Indicates that data on the accelerated volume cannot be accessed because the flash memory hardware it is associated with is missing, or the volume data does not match the flash memory data.
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 coherent 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 migrated to a second disk. This state displays when a volume is created, the volume size is increased, or the type is changed.
Rebuilding	Indicates that data is being copied to a mirror or spare disk, and that data redundancy is being restored. A RAID 0 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 flash memory. Enhanced: indicates that the volume is accelerated for optimized data protection. Maximized: indicates that the volume is accelerated for optimized input/output performance. Transitioning: indicates that acceleration is transitioning from maximized to enhanced mode. This operation will start automatically if the flash memory hardware is reported as at risk of failing to avoid potential data loss. Disabling: indicates that acceleration on a volume previously in maximized mode is being disabled. As a result, the volume will no longer be associated with the flash memory and data will no longer be cached. Unavailable: indicates that the acceleration mode cannot be reported because the volume data may not match the data on the flash memory device. Refer to Troubleshooting Flash Memory States for more information.
Size	Reports the total capacity of the volume in gigabytes (GB and MB).
Data strip 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.

### 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'.

<insert icon>Note

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

### Rebuilding a volume

When a volume is reported degraded because of a failed disk, the failed disk must be replaced with a new one and the volume be rebuilt in order to maintain fault-tolerance. The option to rebuild is only available when a disk is connected, available and normal. If a spare disk is available, the rebuild process starts automatically when a disk fails for all volume types except for RAID 0 due to its configuration.

<insert icon>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'

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.
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'

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.

### 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.

<insert icon>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'.

<insert icon>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,

### Resetting volume to Normal

This action is only available when a RAID 0 volume is reported as failed, but both array disks are present and normal.

Under atypical conditions, when a RAID 0 volume is reported as failed, the ability to return the volume to a normal state will be available. In most cases, this will occur after a failed array disk was reset to normal and resetting the volume to normal will allow you to access and try recovering healthy volume data.

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.

### 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
2-disk RAID 1 (The RAID 1 volume must utilize 100% of the available space on the array in order to complete this action)	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

<insert icon>Note



Before starting, refer to the system requirements to determine which RAID types are supported by your computer and make sure the required number of SATA disks are connected. 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'.
5. The 'Manage' page refreshes and reports the new volume type.

<insert icon>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

### Increasing volume size

You can increase the size of a RAID volume by utilizing the remaining available space on the array. Hovering over the array name in the storage system view displays the amount of available space in MB.

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.

### 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.
4. Restart your computer for changes to take effect, and then use Windows Disk Management\* to utilize the available space.

### 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.
4. Restart your computer for changes to take effect, and then use Windows Disk Management\* to utilize the available space.

<insert icon>Warning

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

### Adding a disk

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.

<insert icon>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.

### **Adding a Disk from Manage Array or Manage Volume**

1. Under 'Status' or 'Manage', in the storage system view, click the array or 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'.
5. Restart your computer for changes to take effect, and then use Windows Disk Management\* to increase the capacity of the array and the volumes.

<insert icon>Warning

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

### **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.

<insert icon>Note

This action is only available if a recovery volume is present and in normal state

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.

### **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.

<insert icon>Note

This action is only available if a recovery volume is present, in a normal state, 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. 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'.

<insert icon>Note

You can follow the progress of the update by hovering over the notification area icon or under 'Status' or Manage Volume

### **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.

<insert icon>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.



### 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.

<insert icon>Note

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

<insert icon>Warning

Windows Explorer will not open until the disk has a partition on it because there are no files to be displayed

### 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 start correctly.

<insert icon>Warning

<insert icon>Warning

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.



### 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.

### 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 is first written to the cache memory, grouping multiple I/O requests into fewer requests, and writing from the cache to the volume at defined intervals. By default, the write-back cache is disabled.



<insert icon>Warning

### 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.

<insert icon>Note

### 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.

### 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.

### 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.
4. Once complete, click 'Verify' to start the verification process.

<insert icon>Note

<insert icon>Warning

### 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.

<insert icon>Note

### 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.

<insert icon>Note

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.

<insert icon>Warning

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.

## 16.4.3 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.

### 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, external or on a port multiplier.
Usage	Array disk: a disk that has been grouped with other disks to form an array containing RAID volumes.



	<p>Master disk: the disk that is the designated source drive in a recovery volume.</p> <p>Recovery disk: the disk that is the designated destination drive in a recovery volume.</p> <p>Spare: the disk has been designated as the destination drive for automatic rebuilds in the event of a failed or missing array disk. A spare disk cannot be used for a RAID 0 volume.</p> <p>Available: the disk is physically connected to the computer, and is unused or has unallocated space that can be used.</p> <p>Unknown: the disk is available but contains metadata that cannot be displayed in the operating system.</p>
Acceleration mode	<p>Reports the acceleration mode for the disk or volume associated with the flash memory.</p> <p>Enhanced: indicates that the disk is accelerated for optimized data protection.</p> <p>Maximized: indicates that the disk is accelerated for optimized input/output performance.</p> <p>Transitioning: indicates that acceleration is transitioning from maximized to enhanced mode. This operation will start automatically if the flash memory hardware is reported as at risk of failing to avoid potential data loss.</p> <p>Disabling: indicates that acceleration on a disk previously in maximized mode is being disabled. As a result, the disk will no longer be associated with the flash memory and data will no longer be cached.</p> <p>Unavailable: indicates that the acceleration mode cannot be reported because the disk data may not match the data on the flash memory device. Refer to Troubleshooting Flash Memory States for more information.</p>
Status	<p>Normal: the disk is present, functioning as expected, and unlocked.</p> <p>Locked: the disk is password-protected. Note: if a volume includes at least one locked disk, the volume will display as locked.</p> <p>At risk: an impending error condition was detected on the disk and it is now at risk of failure.</p> <p>Missing: the disk is not present or physically connected to the computer.</p> <p>Failed: the disk has exceeded its recoverable error threshold.</p> <p>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.</p> <p>Inaccessible: indicates that data on the accelerated disk cannot be accessed because the flash memory hardware it is associated with is missing, or the disk data does not match the flash memory data.</p>
Size	Reports the total capacity of the disk in megabytes (MB and GB).
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 semi-permanent code 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 for all array disks. This feature is controlled at the array level.
Native command queuing	Reports whether the disk supports this feature.
SATA transfer mode	Reports the transfer mode between the SATA controller and the SATA disk. The typical values for this parameter are: <ul style="list-style-type: none"><li>· Generation 1: 150 Mbytes/s</li><li>· Generation 2: 300 Mbytes/s</li></ul>
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).

### Activating a port LED

This action is only available if your system meets the hardware and configuration requirements to support this feature.

Activating the LED of a port on the SATA controller allows you to locate the port to which a specific disk is connected. An LED is an indicator light on a device connected to the computer that flashes when you select this option.

1. Under 'Status' or 'Manage', in the storage system view, click the disk for which you want to locate the port it is connected to. The volume properties are now displayed on the left.
2. Click 'Activate port LED'.
3. Once located, the corresponding port light will flash and timeout automatically.

### 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 set up using a third party software. Locked disk 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'.

<insert icon>Note

### Marking a disk as spare

This action is only available for desktops because a mobile computer is limited to one internal and one external disk, which are used to create a volume.

Marking a disk as spare allows you to designate an available SATA disk as the default destination for automatic rebuilds in the event of a failed or missing array disk that cannot be reconnected.



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'.

<insert icon>Note

<insert icon>Warning

### 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.

### 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

Early warning signs of failure with a disk that result from a slow degradation over time are detected. It may take a few alerts before the disk actually fails, but we recommend that you contact the manufacturer for more information to prevent potential data loss. In the meantime, you can reset the disk to normal by clearing the event in order to keep the overall storage system healthy and access all management options.

#### Failed

A SATA disk has exceeded its recoverable error threshold 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 read/write data access consistently fails, the disk will likely 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.

<insert icon>



### Accessing disk files

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

<insert icon>Note

<insert icon>Warning

### 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.

<insert icon>Note

### Connecting disks files

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.

### 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.



<insert icon>Note

### 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.

### Managing ports

A port is a connection point on your computer where you can physically connect devices, such as SATA disks or ATAPI devices, that transmit data into and out of the computer. A port allows a single connection, while a SATA port multiplier allows the connection of multiple SATA devices to a single SATA host port. This application currently supports six internal ports, where one can be used by a 15-port port multiplier.

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.

If a port is reported as empty in the storage system view, you can use that port to connect a new device, such as a SATA disk in order to increase the storage system capacity.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal, external or on a port multiplier.

### 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, external or on a port multiplier.
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 semi-permanent code found in the disk.
SATA transfer mode	Reports the transfer mode between the SATA controller and



	the SATA disk. The typical values for this parameter are: <ul style="list-style-type: none"> <li>· Generation 1: 150 Mbytes/s</li> <li>· Generation 2: 300 Mbytes/s</li> </ul>
--	--

## 16.5 Accelerating the Storage System

Using the Accelerate area allows you to configure the internal flash memory in order to improve the overall system performance. You can choose to accelerate a disk or a volume.

### Configuring flash memory

You can configure a non-volatile flash memory device, such as USB flash drives, to be used as an additional memory cache for a system or non-system disk or volume present on your system. Moving frequently accessed data over to the flash memory 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 flash memory instead of the SATA disk itself.

To configure the flash memory, you can accelerate a disk or volume in enhanced or maximized mode; by default, the system disk or volume is accelerated. You have the option to accelerate a different disk or volume present in the system if necessary.

<insert icon>Note

### Flash memory properties

The Flash Memory area is only available if the following requirements are met:

- Microsoft Windows\* Vista, Microsoft Windows\* Vista x64-bit Edition, or higher.
- Intel® PCHM SATA RAID controller is present.
- 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).
- The Non-Volatile Memory Host Controller Interface (NVMHCI) is available and the allocated size for caching has a minimum capacity of 3.5 GB.
- A flash memory device is connected.

### Limitations

- Only one disk or volume at a time can be accelerated per system.
- If two volumes are present on a single array and share array disks, neither volume can be accelerated.
- Acceleration on a disk must be disabled before it can be used to create a volume.

### Flash memory properties

Parameter	Value
Flash memory status	Reports the state of health of the flash memory device connected to the computer.
Acceleration mode	Reports the acceleration mode for the disk or volume associated with the flash memory. Disabled: indicates that no disk or volume is accelerated.



	<p>Enhanced: indicates that the disk or volume is accelerated for optimized data protection.</p> <p>Maximized: indicates that the disk or volume is accelerated for optimized input/output performance.</p> <p>Transitioning: indicates that acceleration is transitioning from maximized to enhanced mode. This operation will start automatically if the flash memory hardware is reported as at risk of failing to avoid potential data loss.</p> <p>Disabling: indicates that acceleration on a disk or volume, previously in maximized mode, is being disabled. As a result, the disk or volume will no longer be associated with the flash memory and data will no longer be cached.</p> <p>Unavailable: indicates that the acceleration mode cannot be reported because the disk or volume data may not match the data on the flash memory device. Refer to Troubleshooting Flash Memory States for more information.</p>
Accessible capacity	Reports the portion of the total flash memory capacity in gigabytes (GB) that can be used for accelerating system and non-system disks or volumes present in the system. In addition, part of the accessible capacity is reserved to accelerate applications and files. Refer to Accelerating Applications and Files for more information.
Serial number	Reports the manufacturer's serial number for the flash memory device.

### Enabling acceleration

System or non-system disks or volumes present in the system can be accelerated in enhanced or maximized mode:

- **Enhanced mode:** acceleration optimized for data protection.  
This mode uses the write-through cache method to write data to the flash memory and the disk simultaneously. In the event that the disk or volume associated with the flash memory becomes inaccessible, fails, or is disconnected from the flash memory hardware, there is no risk of data loss because data on the disk is always synchronized with the data on the flash memory.
- **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 disk or volume associated with the flash memory becomes inaccessible, fails, or is disconnected from the flash memory hardware, it is highly likely that data loss occurs. In most cases, data on the disk wasn't synchronized with the data on the flash 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 data on the flash memory.

<insert icon>Note

### Enabling acceleration on a system disk or volume

When acceleration is enabled from the Status area, the system disk or volume is accelerated by default. If you want to accelerate a non-system disk or volume, you will need to perform this action under 'Accelerate'.

1. Under 'Status', in the Accelerate subsection, click 'Enable' to accelerate the system disk or volume.
2. Select the acceleration mode that you want to enable, and then click 'OK'.



3. The page refreshes and reports the new acceleration mode. In the storage system view, hover over the accelerated disk or volume to verify that the acceleration mode is displayed.

### **Enabling acceleration on a non-system disk or volume**

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. In the Advanced section, click 'Customize acceleration settings'. Note: This option is only available if acceleration is disabled.
3. Select the disk or volume that you want to accelerate from the drop-down list.
4. Select the acceleration mode that you want to enable, and then click 'OK'.
5. The page refreshes and reports the new acceleration mode. Under 'Status', in the storage system view, hover over the accelerated disk or volume to verify that the acceleration mode is displayed.

### **Disabling acceleration**

You can choose to disable acceleration on a disk or volume if you want to use the accelerated disk to create a volume or if you want to enable acceleration on a different disk or volume.

Completing this action makes any cached data associated with the accelerated disk or volume immediately inaccessible.

<insert icon>Warning

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. In the Advanced section, click 'Disable', and then click 'Yes' to confirm.
3. The page refreshes and reports that acceleration is disabled.

### **Customizing acceleration**

This option is only available if acceleration on a disk or volume is disabled.

You can choose to enable acceleration on a different disk or volume other than the default selection by customizing acceleration settings. In order to perform this action, you must first disable the currently accelerated disk or volume if applicable.

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. In the Advanced section, click 'Customize acceleration settings'.
3. Select the disk or volume that you want to accelerate from the drop-down list.
4. Select the acceleration mode that you want to enable, and then click 'OK'.
5. The page refreshes and reports the new acceleration mode. Under 'Status', in the storage system view, hover over the accelerated disk or volume to verify that the acceleration mode is displayed.

### **Disassociating from flash memory**

This action is only available if the acceleration mode is displayed as unavailable. Refer to Troubleshooting Flash Memory States to review all available options before completing this action.

You may decide or have to remove the association between an accelerated disk or volume and the flash memory under one of the following circumstances:



- The accelerated disk or volume was removed and then hot-plugged back into the same computer.
- The accelerated disk or volume associated with the flash memory is missing.
- The flash memory associated with the accelerated disk or volume is missing.

Disassociating the flash memory from an accelerated disk or volume removes all acceleration information from both, and it allows you to enable acceleration on a different disk or volume, as long as the flash memory hardware is connected and functional. This action can be performed from the Manage or Accelerate area, except when the disk or volume is missing.

<insert icon>Warning

#### **From the Accelerate area**

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. Click 'Disassociate from flash memory'.
3. Click 'Yes' to confirm.
4. The page refreshes and indicates that acceleration is disabled.

#### **From the Manage area**

1. Under 'Status' or 'Manage', in the storage system view, click the accelerated disk or volume. The properties are now displayed on the left, and the status is reported as inaccessible.
2. Click 'Disassociate flash memory from volume' (or disk).
3. Click 'Yes' to confirm.
4. The page refreshes and indicates that acceleration is disabled.

<insert icon>Note

## **16.6 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 show, such as storage system warnings or errors, and be notified of any reported problems while the application is closed.

### **Showing and hiding system notifications**

1. Choose 'Show the notification area icon'.
2. Select the types of notifications you want to receive, and then click 'Apply Changes'.
3. To hide the notification area icon, deselect 'Show the notification area icon', and then click 'Apply Changes'.

<insert icon>Note

## **16.7 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.



## 16.7.1 Failed volumes

<b>RAID 0</b>	
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	<p>Follow this procedure to recover data:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and reconnect the missing disk.</li> <li>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.</li> <li>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.</li> <li>4. Under 'Status', verify that the volume and disks status display as 'Normal'. You can also review this information under 'Manage'.</li> </ol>
Failed array disk	<p>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.</p> <p>This procedure deletes the failed volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and replace the failed SATA disk with a new one that is of equal or greater capacity.</li> <li>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.</li> <li>3. Press Ctrl-I to access the main menu of the option ROM user interface.</li> <li>4. Select Delete RAID Volume from the main menu.</li> <li>5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.</li> <li>6. Press the 'Delete' key to delete the volume, then 'Y' to confirm.</li> <li>7. Create a new RAID 0 volume utilizing the new disk. If the failed disk was part of the system volume, you will also need to reinstall the operating system.</li> </ol>
<b>RAID 5</b>	
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	<p>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.</p> <p>This procedure deletes the failed volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity.</li> <li>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.</li> <li>3. Press Ctrl-I to access the main menu of the option ROM user interface.</li> <li>4. Select Delete RAID Volume from the main menu.</li> <li>5. From the Delete Volume menu, select the failed RAID volume,</li> </ol>



	<p>using the up and down arrow keys.</p> <ol style="list-style-type: none"> <li>6. Press the 'Delete' key to delete the volume, then 'Y' to confirm.</li> <li>7. Create a new RAID 5 volume utilizing the new disks. If the failed disk was part of the system volume, you will also need to reinstall the operating system.</li> </ol>
<p><b>RAID 10</b> 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.</p>	
<b>Cause</b>	<b>Solution</b>
Two adjacent array disks missing	<ol style="list-style-type: none"> <li>1. Power off your computer and reconnect the missing disks.</li> <li>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'.</li> </ol>
Three or four array disks missing	<p>In most cases, the volume cannot be recovered and any data on the volume is lost. This procedure deletes the failed volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and reconnect the missing disks.</li> <li>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.</li> <li>3. Press Ctrl-I to access the main menu of the option ROM user interface.</li> <li>4. Select Delete RAID Volume from the main menu.</li> <li>5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.</li> <li>6. Press the 'Delete' key to delete the volume, then 'Y' to confirm.</li> <li>7. Create a new RAID 10 volume utilizing the new disks.</li> <li>8. You will then need to reinstall the operating system on the new volume.</li> </ol>
Two or more array disks failed	<p>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.</p> <p>This procedure deletes the failed volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and replace the failed SATA disks with new ones that are of equal or greater capacity.</li> <li>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.</li> <li>3. Press Ctrl-I to access the main menu of the option ROM user interface.</li> <li>4. Select Delete RAID Volume from the main menu.</li> <li>5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.</li> <li>6. Press the 'Delete' key to delete the volume, then 'Y' to confirm.</li> <li>7. Create a new RAID 10 volume utilizing the new disks.</li> <li>8. You will then need to reinstall the operating system on the new volume.</li> </ol>

## 16.7.2 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	<p>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.</p> <p>If a SATA disk is available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. Under 'Status', click 'Rebuild to another disk'.</li> <li>2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'.</li> <li>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'.</li> <li>4. Once the operation successfully completed, the recovery disk and volume status will display as 'Normal'.</li> <li>5. The volume returns to the same update mode it was set to before the rebuild operation.</li> </ol> <p>&lt;insert icon&gt;Note</p>
Master disk missing	<p>If you can reconnect the missing disk, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and reconnect the missing disk.</li> <li>2. Turn on your computer and the system will automatically boot from the recovery disk. The rebuild operation will start automatically.</li> <li>3. The volume returns to the same update mode it was set to before the rebuild operation.</li> </ol> <p>If you cannot reconnect the missing disk and a SATA disk is available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. Under 'Status', click 'Rebuild to another disk'.</li> <li>2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'.</li> <li>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'.</li> <li>4. Once the operation successfully completed, the master disk and volume status will display as 'Normal'.</li> <li>5. The volume returns to the same update mode it was set to before the rebuild operation.</li> </ol> <p>&lt;insert icon&gt;Note</p>
Master disk failed	<p>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.</p> <p>If a SATA disk is available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. Under 'Status', click 'Rebuild to another disk'.</li> </ol>



	<ol style="list-style-type: none"> <li>2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'.</li> <li>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'.</li> <li>4. Once the operation successfully completed, the master disk and volume status will display as 'Normal'.</li> <li>5. The volume returns to the same update mode it was set to before the rebuild operation.</li> </ol> <p>&lt;insert icon&gt;Note</p>
--	--

<p><b>RAID 1</b> 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.</p>	
<p><b>RAID 5</b> 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.</p>	
<p><b>RAID 10</b> 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.</p>	
Cause	Solution
Missing array disk	<p>If you can reconnect the missing disk, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. Power off your computer and reconnect the missing disk.</li> <li>2. Turn on your computer and the rebuild operation will start automatically.</li> </ol> <p>If you cannot reconnect the missing disk and a SATA disk is available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> <li>1. If a SATA disk is available and normal, follow this procedure to rebuild the volume:</li> <li>2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'.</li> <li>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'.</li> <li>4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'.</li> </ol> <p>&lt;insert icon&gt;Note</p>
Failed array disk	<p>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.</p> <p>If a SATA disk is available and normal, follow this procedure to rebuild the volume:</p>



	<ol style="list-style-type: none"> <li>1. Under 'Status', click 'Rebuild to another disk'.</li> <li>2. Select the disk you want to use to rebuild the volume, and then click 'Rebuild'.</li> <li>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'.</li> <li>4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'.</li> </ol> <p>&lt;insert icon&gt;Note</p>
--	---

### 16.7.3 Other volume states

<b>Locked</b>	
Cause	Solution
At least one (but not all) disk included in the volume is locked with a password.	<p>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:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Unlock'.</li> <li>3. Enter the password, and then click 'Unlock'.</li> </ol> <p>Repeat this procedure for all locked disks included in the volume in order to unlock the volume.</p> <p>&lt;insert icon&gt;Note</p>
<b>Unknown</b>	
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.
<b>Power-saving mode (Recovery volumes only)</b>	
Cause	Solution
Your computer is running on battery and the volume is in continuous update mode. Data updates to the recovery disk or a data recovery operation are no occurring.	<p>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.</p> <p>&lt;insert icon&gt;Note</p>
<b>Data update needed (Recovery volumes only)</b>	
Cause	Solution
The data on the recovery disk is not synchronized with the data on the	<p>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:</p> <ol style="list-style-type: none"> <li>1. Under 'Status' or 'Manage', in the storage system view, click the</li> </ol>



master disk.	<p>recovery volume. The volume properties are now displayed on the left.</p> <ol style="list-style-type: none"> <li>2. Click 'Update data'.</li> <li>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.</li> <li>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'.</li> </ol>
<b>Running off recovery disk (Recovery volumes only)</b>	
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.	<p>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.</p> <p>&lt;insert icon&gt;Warning</p> <ol style="list-style-type: none"> <li>1. Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.</li> <li>2. Click 'Recover data', then 'OK' to confirm.</li> <li>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.</li> </ol>
<b>Recovery disk read-only (Recovery volumes only)</b>	
Cause	Solution
The recovery disk files have been accessed and display in Windows Explorer*.	<p>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:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume.</li> <li>3. The Windows Explorer window closes.</li> </ol> <p>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.</p>
<b>Master disk read-only (Recovery volumes only)</b>	
Cause	Solution
Your computer was booted from the recovery disk using the option ROM, and the volume is operating from that disk. The master	<p>Data mirroring and redundancy are lost, and your data is at risk in the event of a disk failure. We recommend that you hide the master disk files when finished with your review, and proceed with a data recovery to the master disk in order to maintain full redundancy. Follow this procedure to hide master disk files from Manage Disk:</p> <ol style="list-style-type: none"> <li>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.</li> </ol>



<p>disk files have been accessed and are displayed in Windows Explorer*.</p>	<ol style="list-style-type: none"> <li>2. Click 'Hide Files' from Manage Disk or 'Hide master disk files' from Manage Volume.</li> <li>3. The Windows Explorer window closes and the volume is displayed as running off recovery disk.</li> </ol> <p>Refer to the 'Running off recovery disk' procedure above to recover data to the master disk.</p>
--	---

### 16.7.4 Disk events

State	Cause	Solution
Locked	An internal or external disk is protected and locked with a password.	<p>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:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Unlock'.</li> <li>3. Enter the password, and then click 'Unlock'.</li> </ol>
At risk	An impending error condition was detected on an internal or external disk and is now at risk of failure.	<p>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:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.</li> </ol> <p>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.</p>
	An unexpected error was detected on a disk that has RAID configuration data (metadata) on it.	<p>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.</p> <p>Warning: Completing this action will permanently delete existing metadata. Back up any accessible data before continuing.</p>



		<ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Clear and reset disk', and then click 'Yes' to confirm.</li> <li>3. Once complete, the page refreshes with the disk returning to a normal state.</li> </ol>
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*.	<p>Hide the recovery or master disk files to return the disk status to offline and resume data updates in on request mode.</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Hide Files' from Manage Disk or 'Hide recovery disk files' from Manage Volume.</li> <li>3. The Windows Explorer window closes.</li> </ol>
Failed	An internal or external disk has exceeded its recoverable error threshold.	<p>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:</p> <ol style="list-style-type: none"> <li>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.</li> <li>2. Click 'Reset disk to normal'. The page refreshes instantly, returning to a normal state.</li> </ol>
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.



		<ol style="list-style-type: none"> <li>Under 'Status' or 'Manage', in the storage system view, click the recovery volume. The volume properties are now displayed on the left.</li> <li>Click 'Change mode', and then click 'Yes' to confirm.</li> <li>The page refreshes and the volume properties report the new update mode.</li> </ol>
	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.

S

## 16.7.5 Flash memory states

### Flash memory hardware issues

State	Cause	Solution
At risk	An impending error condition was detected on the flash memory hardware and is now at risk of failure.	<p>In this state, the flash memory hardware is still functional and data caching is most likely still occurring. If the accelerated disk or volume was in maximized mode, acceleration will automatically transition to enhanced mode to avoid data loss. As a result, the option to enable maximized mode will not be available until the issue on the flash memory hardware is fixed.</p> <p>We recommend that you contact the manufacturer for more information about how to return hardware to a healthy state and to prevent potential data loss.</p>
Failed	The flash memory hardware has exceeded its recoverable error threshold and is no longer operational.	In this state, cached data on the flash memory is lost and cannot be recovered. Acceleration is automatically being disabled because the disk or volume can no longer be associated with the flash memory. Replace the flash memory hardware, and then accelerate a disk or volume present on your computer. Refer to Enabling Acceleration for more information.

### Acceleration mode: Unavailable

This mode results from a potential lack of synchronicity between the disk or volume data and flash memory data resulting from a variety of scenarios. In all cases, details and recommended actions will be reported under 'Status' in the Manage and Accelerate subsections as well as under 'Manage' and 'Accelerate'.

#### Cause

The accelerated disk or volume was removed and then hot-plugged back into the same computer. The data on the disk or volume may not match the data in the flash memory,



depending on the acceleration mode.

#### Solution

In this situation, the disk or volume status will display as inaccessible because read/write activity on the disk or volume has been disabled until it is determined whether the data on the disk or volume and the flash memory matches. You can choose to return to the last remembered acceleration mode or disassociate the flash memory from the accelerated disk or volume.

#### **Return to acceleration in enhanced mode**

If the disk or volume was originally accelerated in enhanced mode, it is most likely that data on the disk or volume matches the data on the flash memory. Therefore, the system will automatically run from the accelerated disk or volume on reboot, and write data to the flash memory as files and applications are being used. Follow this procedure to return the acceleration to enhanced mode:

1. Restart your computer.
2. 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.
3. Under 'Status', in the Accelerate subsection, verify that the acceleration is reported in enhanced mode.

#### **Return to acceleration in maximized mode**

If the disk or volume was originally accelerated in maximized mode, it is likely that the data on the disk or volume does not match the data on the flash memory. For this reason, you need to decide which data you want to use to run applications and files.

Follow this procedure to return the acceleration to maximized mode:

1. Restart your computer.
2. Press Ctrl-I to access the main menu of the option ROM user interface.
3. From this location, choose to use the disk or volume data to rebuild the flash memory data, or to use the flash memory data to write to the accelerated disk or volume.
4. 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.
5. Under 'Status', in the Accelerate subsection, verify that the acceleration mode is reported as Maximized.

<insert icon>Note

#### **Disassociating the disk or volume from the flash memory**

Regardless of whether the data between the accelerated disk or volume and flash memory matches, you can choose to remove the association between the disk or volume and the flash memory. This operation can be completed from the Accelerate or Manage area, or using the option ROM user interface. This operation does not require you to restart your computer.

<insert icon>Warning

From the Accelerate area

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. Click 'Disassociate from flash memory'.
3. Click 'Yes' to confirm.
4. The page refreshes and indicates that acceleration is disabled.

From the Manage area

1. Under 'Status' or 'Manage', in the storage system view, click the accelerated disk or volume. The properties are now displayed on the left, and the status is reported as inaccessible.



2. Click 'Disassociate flash memory from volume' (or disk).
3. Click 'Yes' to confirm.
4. The page refreshes and indicates that acceleration is disabled.

From 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. Type 'r' to remove acceleration.
6. Press 'Y' to confirm

#### Cause

The accelerated disk or volume associated with the flash memory is missing.

#### Solution

##### **The disk or volume can be reconnected**

1. Power off your computer and reconnect the missing disk or volume.
2. Restart your computer.
3. Refer to the procedure described above for the applicable acceleration mode to resolve the problem, or disassociate the flash memory from the disk or volume if you no longer want acceleration enabled.
4. 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.
5. Under 'Status', in the Accelerate subsection, verify that acceleration is reported in enhanced or maximized mode. If you chose to disassociate the flash memory from the disk or volume, acceleration will be displayed as disabled.

##### **The disk or volume cannot be reconnected**

In this case, the only option you have is to disassociate the volume from the flash memory in order to remove acceleration information from the flash memory. Once this operation is complete, you will be able to accelerate a disk or volume that is connected to your computer. Follow this procedure to disassociate the flash memory from the missing disk or volume:

1. Click 'Accelerate', and then click 'Flash Memory' in the left-hand navigation if necessary.
2. Click 'Disassociate from flash memory'.
3. Click 'Yes' to confirm.
4. The page refreshes and indicates that acceleration is disabled.

<insert icon>Note

#### Cause

The flash memory associated with the accelerated disk or volume is missing.

#### Solution



In this case, the disk or volume status will be displayed as inaccessible because the volume cannot locate the flash memory data it is associated with. The acceleration mode is only displayed for the disk or volume, and no longer for the flash memory as the hardware is missing.

**The flash memory hardware can be reconnected**

1. Power off your computer and reconnect the missing disk or volume.
2. Restart your computer.
3. Refer to the procedure described above for the applicable acceleration mode to resolve the problem. You will need to access the option ROM user interface if the acceleration mode is maximized.
4. 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.
5. Under 'Status', in the Accelerate subsection, verify that acceleration is reported in enhanced or maximized mode. If you chose to disassociate the flash memory from the disk or volume, acceleration will be displayed as disabled.

**The flash memory hardware cannot be reconnected**

In this case, the only option you have is to disassociate the flash memory from the disk or volume in order to remove acceleration information. Follow this procedure to return the disk or volume to a normal state:

1. Under 'Status' or 'Manage', in the storage system view, click the accelerated disk or volume. The properties are now displayed on the left, and the status is reported as inaccessible.
2. Click 'Disassociate flash memory from volume' (or disk).
3. Click 'Yes' to confirm.
4. The page refreshes and the acceleration mode is no longer displayed.

**16.7.6 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: <ul style="list-style-type: none"> <li>• Missing components</li> <li>• Corrupted application</li> <li>• Application unable to connect to the service</li> <li>• Application fails to start.</li> </ul>	This option is only available if you are logged on as an administrator and if your computer is configured with at least 2 GB of system memory (RAM).
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. Follow the applicable procedure to resolve the problem.
Multiple users cannot run the application at the same time.	One or more users are attempting to open the application while an instance	Make sure only one instance of the application is running at a time.



	of the application is already running.	
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.



# 17 Braidwood Technology

Braidwood technology is Intel’s high performance, power efficient platform NVM solution. Braidwood uses Intel® High Speed NAND Flash Memory to increase system responsiveness and performance by caching hard disk drive data. Braidwood can also be configured as an Integrated SSD<sup>1</sup> or a combination of both.

## 17.1 Braidwood Platform

Braidwood technology is an integrated NVM platform solution. It is enabled through the use of Intel platform hardware, firmware and software. The platform consists of following hardware components: Integrated physical NAND controller where the NAND controller is integrated into the Platform Controller Hub (PCH), NAND module and ONFI 2.0 connector (for connector-based implementation). See Figure 1: Braidwood Platform. NAND module can be mother-board down (soldered) or connector-based (ONFI 2.0 defined). The integrated physical NAND controller in the Intel® 5 series chipset controls the physical connection to the NAND module and is exposed via AHCI Port 6 (reserved for future SSD partitioning of the NAND module) and via AHCI Port 7 (NVMHCI partition). The physical NAND controller is accessible only by Braidwood firmware.

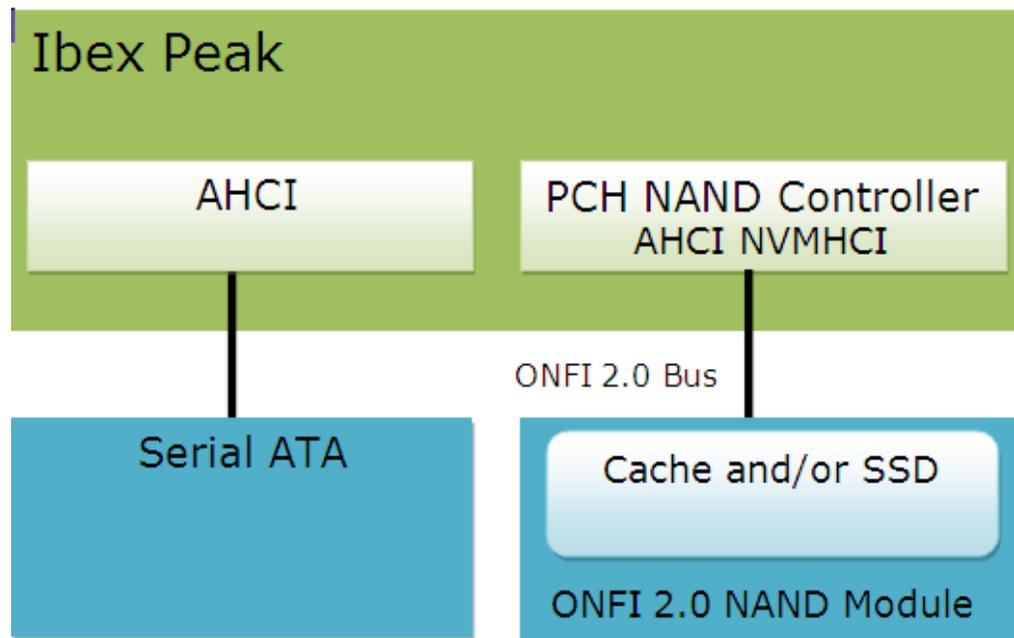


Figure 1: Braidwood Platform

<sup>1</sup> Braidwood Technology Integrated SSD feature will not be available until in a future update of the product



Braidwood firmware is stored in the SPI flash as a part of ME firmware image and downloaded into the Intel® 5 Series chipset at boot up. The firmware runs on the Virtual Engine in the chipset. The firmware is responsible for the NAND flash management including NAND initialization, bad block management, wear leveling, region configuration, NVMHCI support and ATA device emulation (Integrated SSD). Braidwood does not have any special ME image flashing requirements in a manufacturing environment. Current methods of flashing ME image in the manufacturing environment can still be used.

The Intel® Rapid Storage Technology driver provides RAID and AHCI support. The Intel® RST driver interfaces with the NAND via the AHCI NVMHCI (SATA Port 7) or the AHCI SSD (SATA Port 6). The AHCI NVMHCI port will be used to provide disk-cache (WBC/WTC) functionality. An OROM, in addition to providing RAID functionality, detects and handles logical or physical separation events. The OROM also handles pre-OS load I/O processing and power-failure recovery to keep the cache coherent with associated disk. Standard AHCI aware OS driver can be used for SSD access only. See Figure 2: Braidwood Architecture. Any OS with AHCI capabilities can take advantage of the Braidwood Integrated SSD (if configured) with no special knowledge of the NAND subsystem required. Any OS that supports NVMHCI can also take advantage of the NAND region dedicated to the host. Braidwood firmware performs all NAND management functionality such as wear-leveling, region configuration, bad block management, NVMHCI support and ATA device emulation (Integrated SSD).

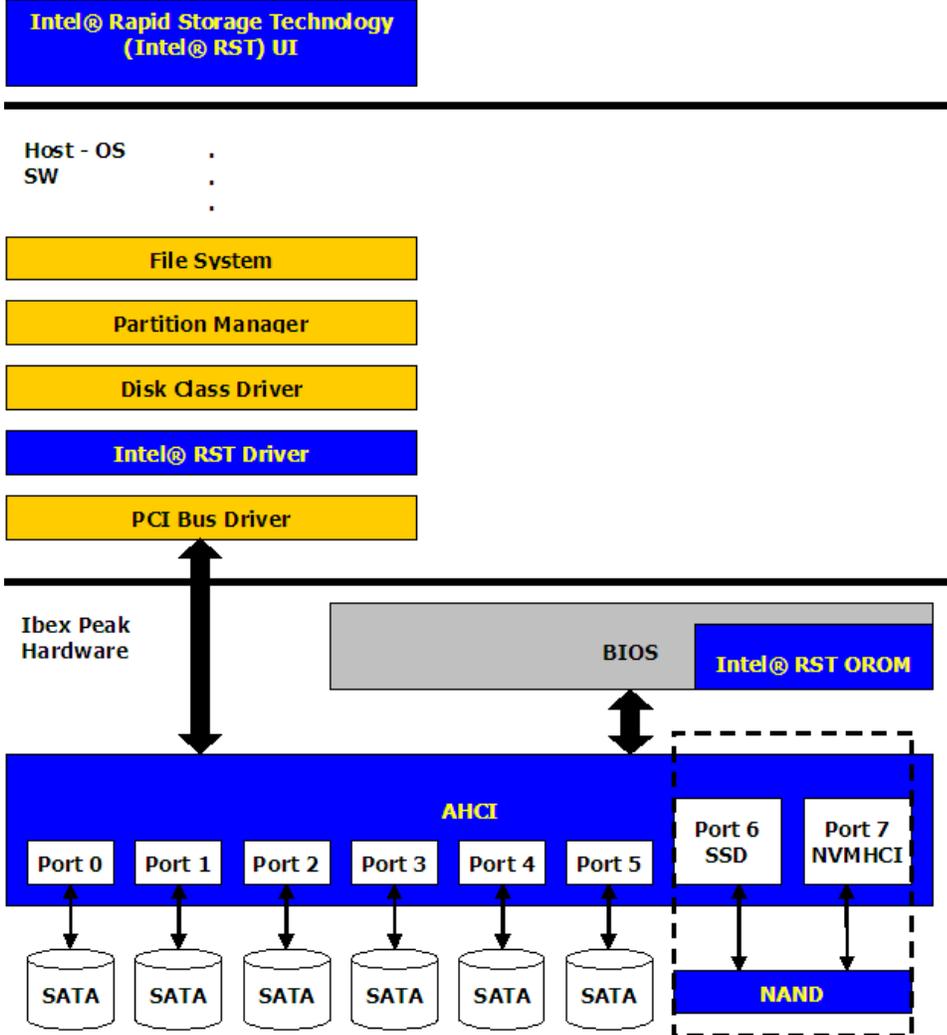


Figure 2: Braidwood Architecture

## 17.2 Benefits

Braidwood provides ease of use through simplified, unified interface and driver for NVM cache and Intel® RST. Braidwood offers significantly improved read/write performance through Intel® RST RAID Write Back Cache, higher capacity, lower power consumption (leading to longer battery life) and improved boot times.



## 17.3 Usage Models

The Braidwood NAND module can be configured so that the total usable<sup>2</sup> capacity of the NAND module is the cache region. The cache can be utilized as either write-back or write-through cache.

A future update of the Braidwood product will allow the NAND module to be configured as both cache and SSD regions. This will allow three usage models: total capacity cache only (currently available), cache + SSD (currently not available), and total capacity SSD only (currently not available).

### 17.3.1 Cache only (total capacity)

The Braidwood NAND module can be configured to create only the cache region. In this usage model, the total usable NAND memory capacity is dedicated for caching. This model uses caching for storage acceleration. This model delivers maximum performance for maximum capacity. Smaller caches can be used for lighter workloads.

## 17.4 Braidwood Implementation

This section covers the process to enable, configure, and to bring up Braidwood capabilities in the platform. This section will cover the use of the NANDUtil tool, how to bring up firmware with the correct setting that will enable the full set of Braidwood capabilities on the system, and also NAND configuration features that can be integrated into the BIOS that will provide the end user the ability to tailor the NAND region and for future upgradeability. It also provides an automated BIOS option for the OEM to configure the NAND module in the manufacturing line.

The following is a list of requirements needed to bring up Braidwood Technology:

1. Ibox Peak Braidwood chipset SKU
2. SATA controller mode:
  - a. RAID: for NVMHCI
3. Firmware with Braidwood SKU
4. BIOS that support all Braidwood technology requirements documented on the VE BIOS Writer's Guide
5. System with the ONFI 2.0 defined connector (for connector-based implementation)
6. Intel Fast Flash NAND module (solder down or connector based)

### 17.4.1 Factory Tools and Process

OEMs can use either the BIOS manufacturing option or the NANDUtil tool in their factory to configure the NAND module before it is shipped to end-users. The NAND

<sup>2</sup> The total usable capacity of the NAND module is less than the total physical capacity of the NAND module. This is due to the fact that a small portion of the NAND module capacity is used for housekeeping and maintenance purposes



module must be configured before it can be used. BIOS manufacturing option is an automated process that can configure any pristine or previously configured NAND module to any predefined configuration. NANDUtil is a DOS based tool that allows more in-depth testing of the NAND module. This tool will write and read back data written into each sector to verify that the NAND module can be accessed. More detailed information of the BIOS manufacturing option and the NANDUtil tool will be covered in the subsequent sections.

#### **17.4.1.1 High Level Process to Configure NAND Module Using the BIOS Manufacturing Option**

This section covers high level process to configure NAND module using the BIOS manufacturing option. More information on the BIOS flow will be covered in the Error: Reference source not found section.

1. Create Braidwood enabled firmware image, flash the image on the target system and set the necessary BIOS setup options. (For more details on Braidwood firmware bring up process, refer to 'FW Bringup Guide' in the latest Intel® Management Engine Firmware 6.0 kit on VIP). NAND module should be inserted during the FW bring up process.
2. Set the NAND configuration bit in BIOS through scripts.
3. Restart the system.
4. BIOS will configure the NAND module.
5. Restart the system.

#### **17.4.1.2 High Level Process to Configure NAND Module Using the NANDUtil Tool**

This section covers high level process to configure NAND module using the NANDUtil tool.

1. Create Braidwood enabled firmware image, flash the image on the target system and set the necessary BIOS setup options. (For more details on Braidwood firmware bring up process, refer to 'FW Bringup Guide' in the latest Intel® Management Engine Firmware 6.0 kit on VIP). NAND module should be inserted during the FW bring up process.
2. Boot the system to DOS environment.
3. Run NANDUtil tool to verify that VECI is running. (For more details on NANDUtil tool usage, please refer to Error: Reference source not found section).
4. If the NAND module is already configured, run NANDUtil tool destroy region command to destroy the existing region.
5. Create new region using NANDUtil tool create region command.
6. Restart the system.



### 17.4.1.3 NANDUtil Tool

NANDUtil is a DOS based tool provided by Intel to be used during development process of Braidwood technology. The NANDUtil tool provides a wide range of features that covers from configuration to testing, and compliancy of the Braidwood system.

The following is description of each capability provided by the NANDUtil tool:

- **Query:** This command is used for fetching the details of Flash and region(s) from NAND Flash devices on the NAND module.
- **Create:** This command creates a NVMHCI (cache) in the NAND module.
- **Destroy:** This command destroys the NVMHCI region on the NAND module.
- **Erase:** This command erases all data on the NVMHCI region in the NAND module.
- **Test:** This command initiates a memory test of the NAND module.
- **Status:** This command reports the version and status of the NAND FW
- **Compliance:** This command checks the platform to ensure it meets PCH requirements and will operate with the Braidwood NAND modules.

For more details or usage of the NANDUtil tool, refer to the 'Braidwood Tools User Guide' in the latest Intel® Management Engine Firmware 6.0 kit on VIP.

### 17.4.1.4 BIOS Options / Manufacturing BIOS Option

This section provides sample flows for NANDUtil features that can be implemented in the BIOS. These flows are recommended for the IBVs and OEMs to implement in their BIOS to allow the end user to re-configure the existing region in the NAND module or for future upgradeability. It also provides a manufacturing option that allows the OEMs to configure the NAND region during the manufacturing line process, allowing seamless integration of the Braidwood technology manufacturing requirement into their existing manufacturing line process.

The focus of these sample flows is the VECI message that BIOS needs to send for a specific feature. These are sample flows and should be treated as a guideline only. IBVs and OEMs may choose the best implementation to suit their unique needs.

Create:

This section shows a sample flow (see Figure 3: BIOS Option - Create Region(s)) for the create NVMHCI region in the NAND module. The create option should be located in the BIOS set up menu. This option will be grayed out if VECI failed to initialize or there is no NAND module connected in the system.



1. Check device status and characteristic
2. Verify NAND module has initialized correctly.
  - If NAND module has initialized correctly, display the NAND configuration option in the BIOS setup menu
  - If not, wait for 250 ms. If more than 15 seconds have elapsed and NAND module has not initialized, gray out the NAND configuration option and continue to boot.
3. Query existing regions in the NAND module
  - If there is no region, prompt user for option to create regions (step 4).
  - If both NVMeHCI and SSD regions are already present in the NAND module, display error messaging and exit NAND configuration option.
4. Prompt user for region creation options
5. Validate user request is valid. If not, loopback to step 4.
6. Send VECI create message.
7. Validate response. If there is no error, display successful messages.

## BIOS Option - Create

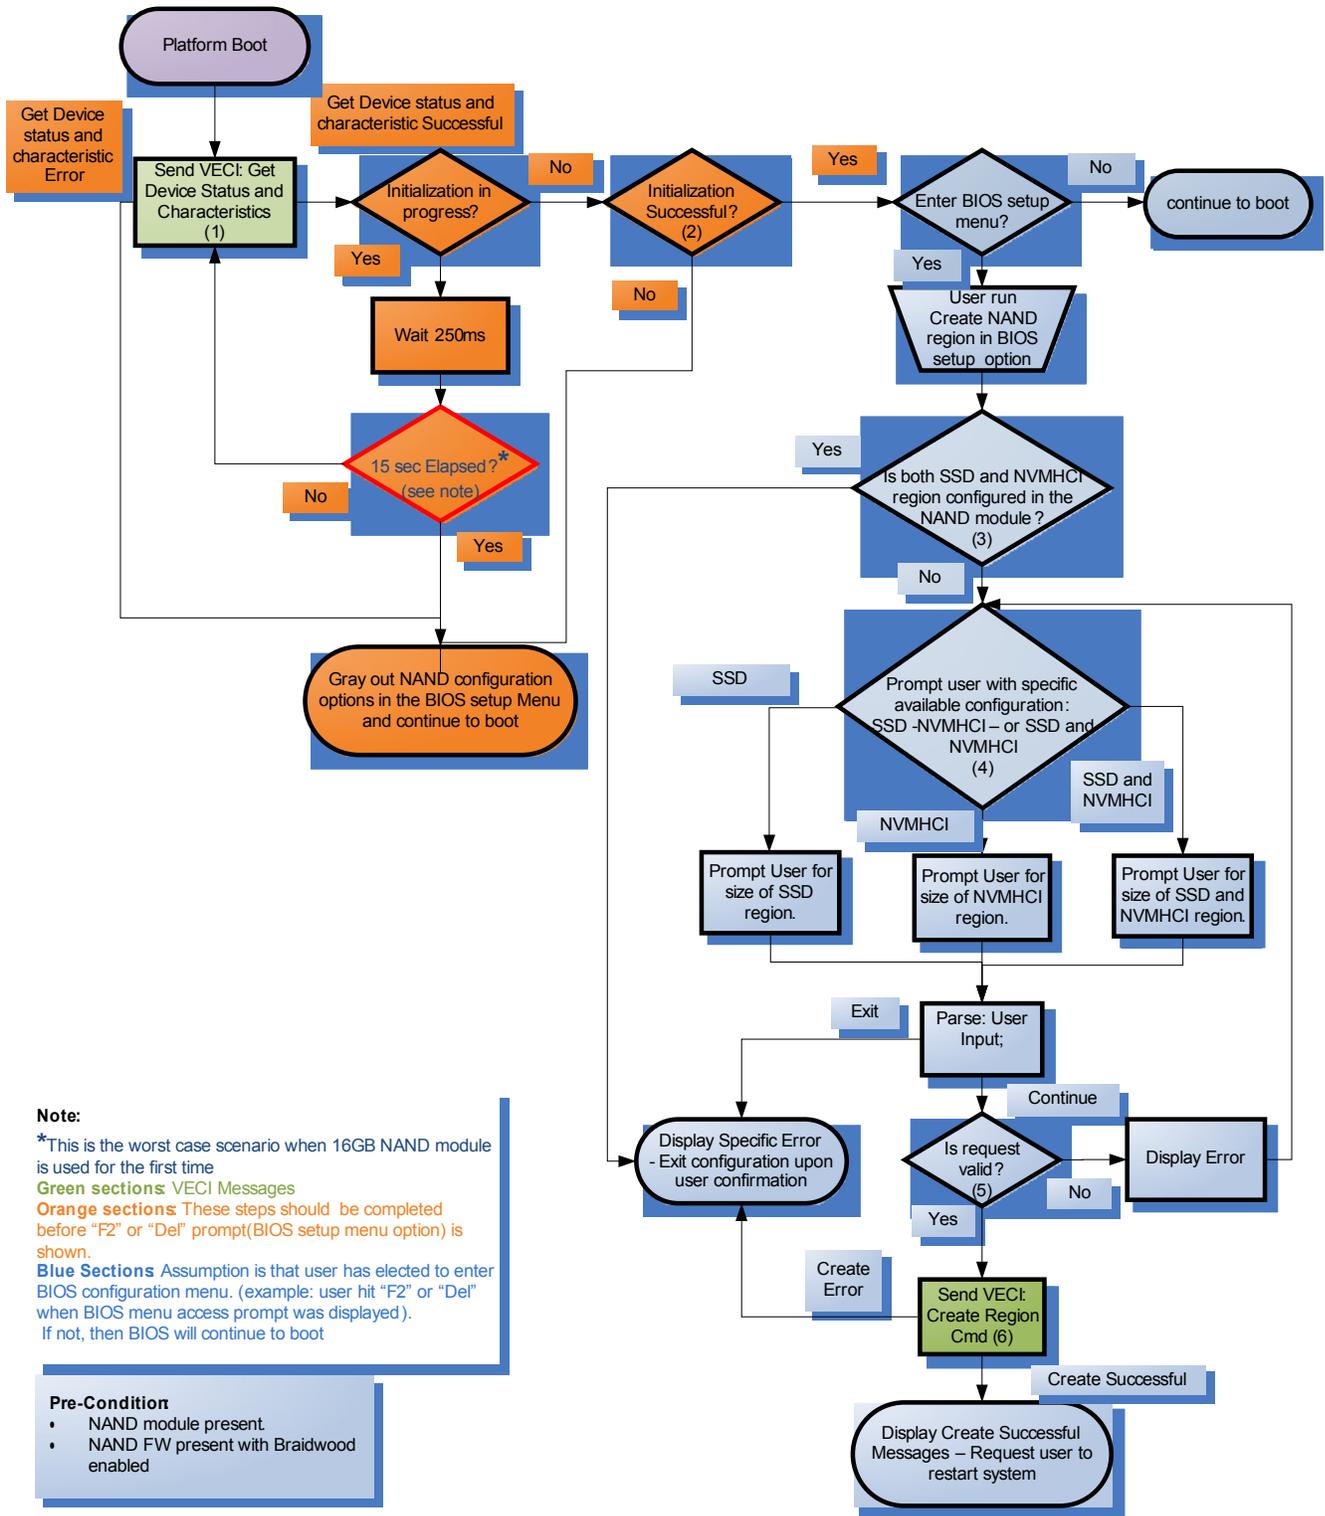




Figure 3: BIOS Option - Create Region(s)

Destroy:

This section shows a sample flow (see Figure 4: BIOS Option - Destroy Region(s)) for the destroy NVMHCI region in the NAND module. The destroy option is located in the BIOS set up menu. This option will be grayed out if VECI failed to initialize or there is no NAND module connected in the system.

1. Check device status and characteristic
2. Verify NAND module has initialized correctly.
  - If NAND module has initialized correctly, display the NAND configuration option in the BIOS setup menu
  - If not, wait for 250 ms. If more than 15 seconds have elapsed and NAND module has not initialized, gray out the NAND configuration option and continue to boot.
3. Query existing regions in the NAND module
  - If there is no region present in the NAND module, display error messaging and exit NAND configuration option.
  - If either NVMHCI or SSD is already present in the NAND module, prompt user for option to destroy regions (step 4).
4. Prompt user for destroy region warning.
5. Validate user request is valid. If not, exit NAND configuration option.
6. Send VECI destroy message.
7. Validate response. If there is no error, display "destroy successful" messages.

Note: all data in the NAND regions will be destroyed

## BIOS Option - Destroy

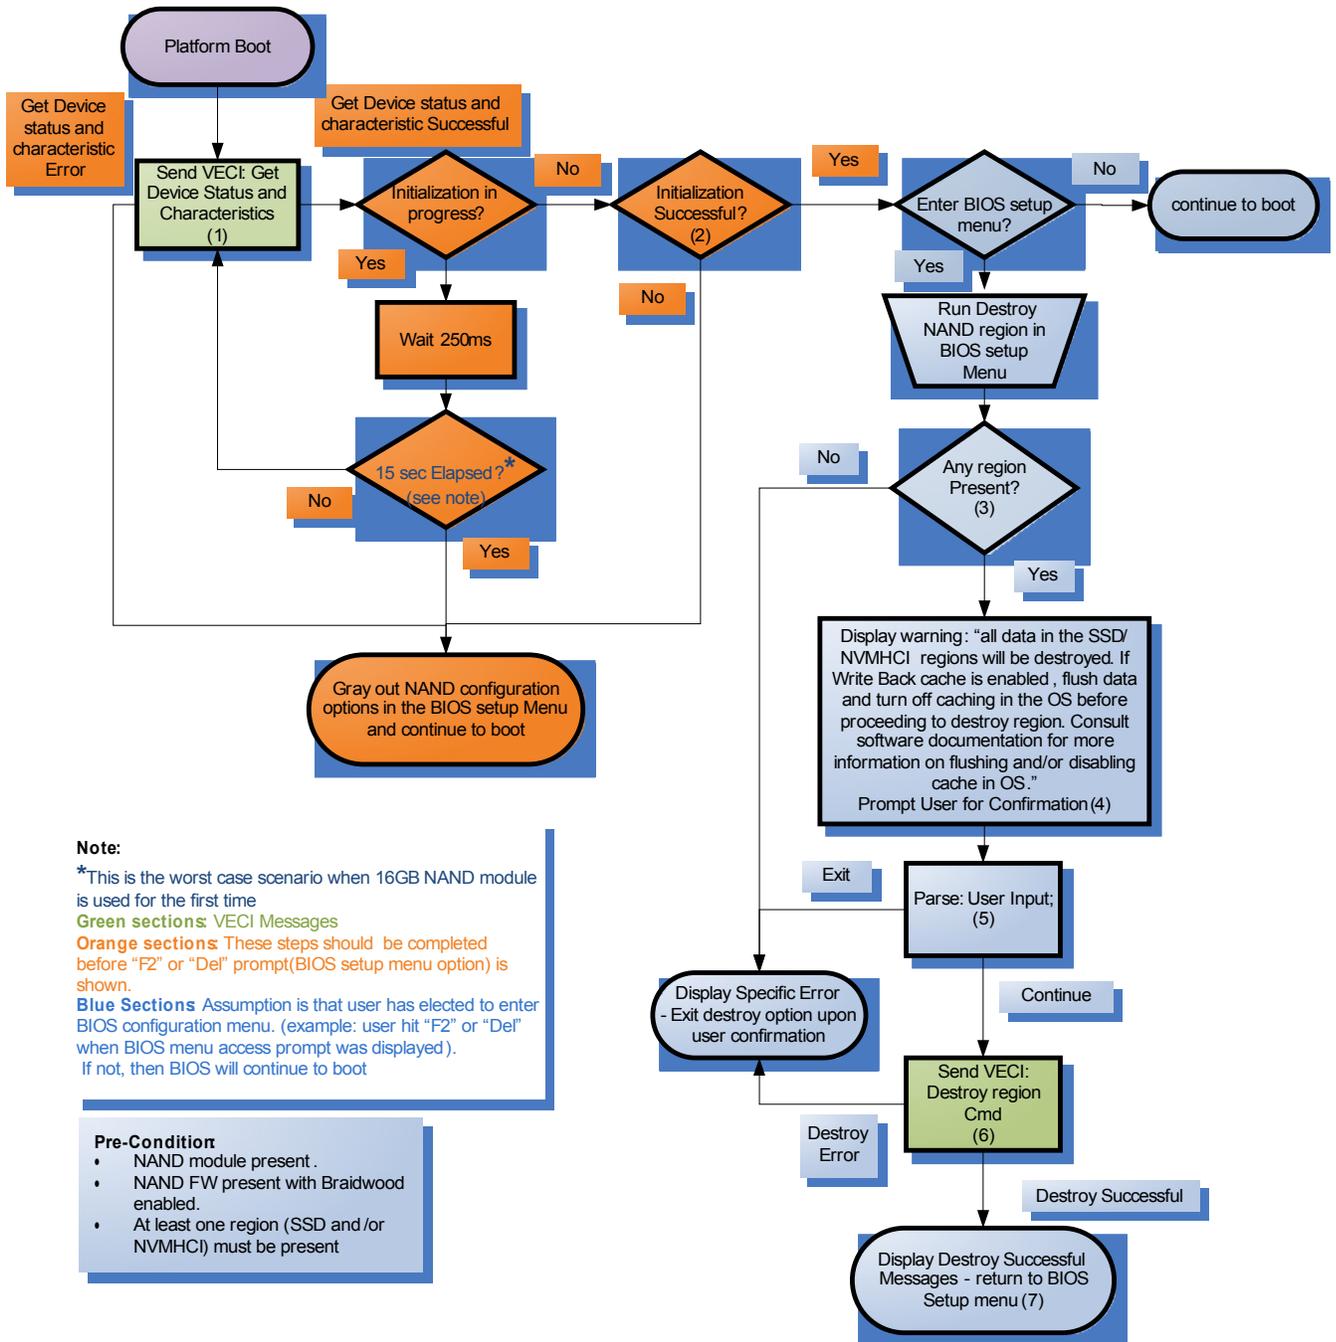


Figure 4: BIOS Option - Destroy Region(s)



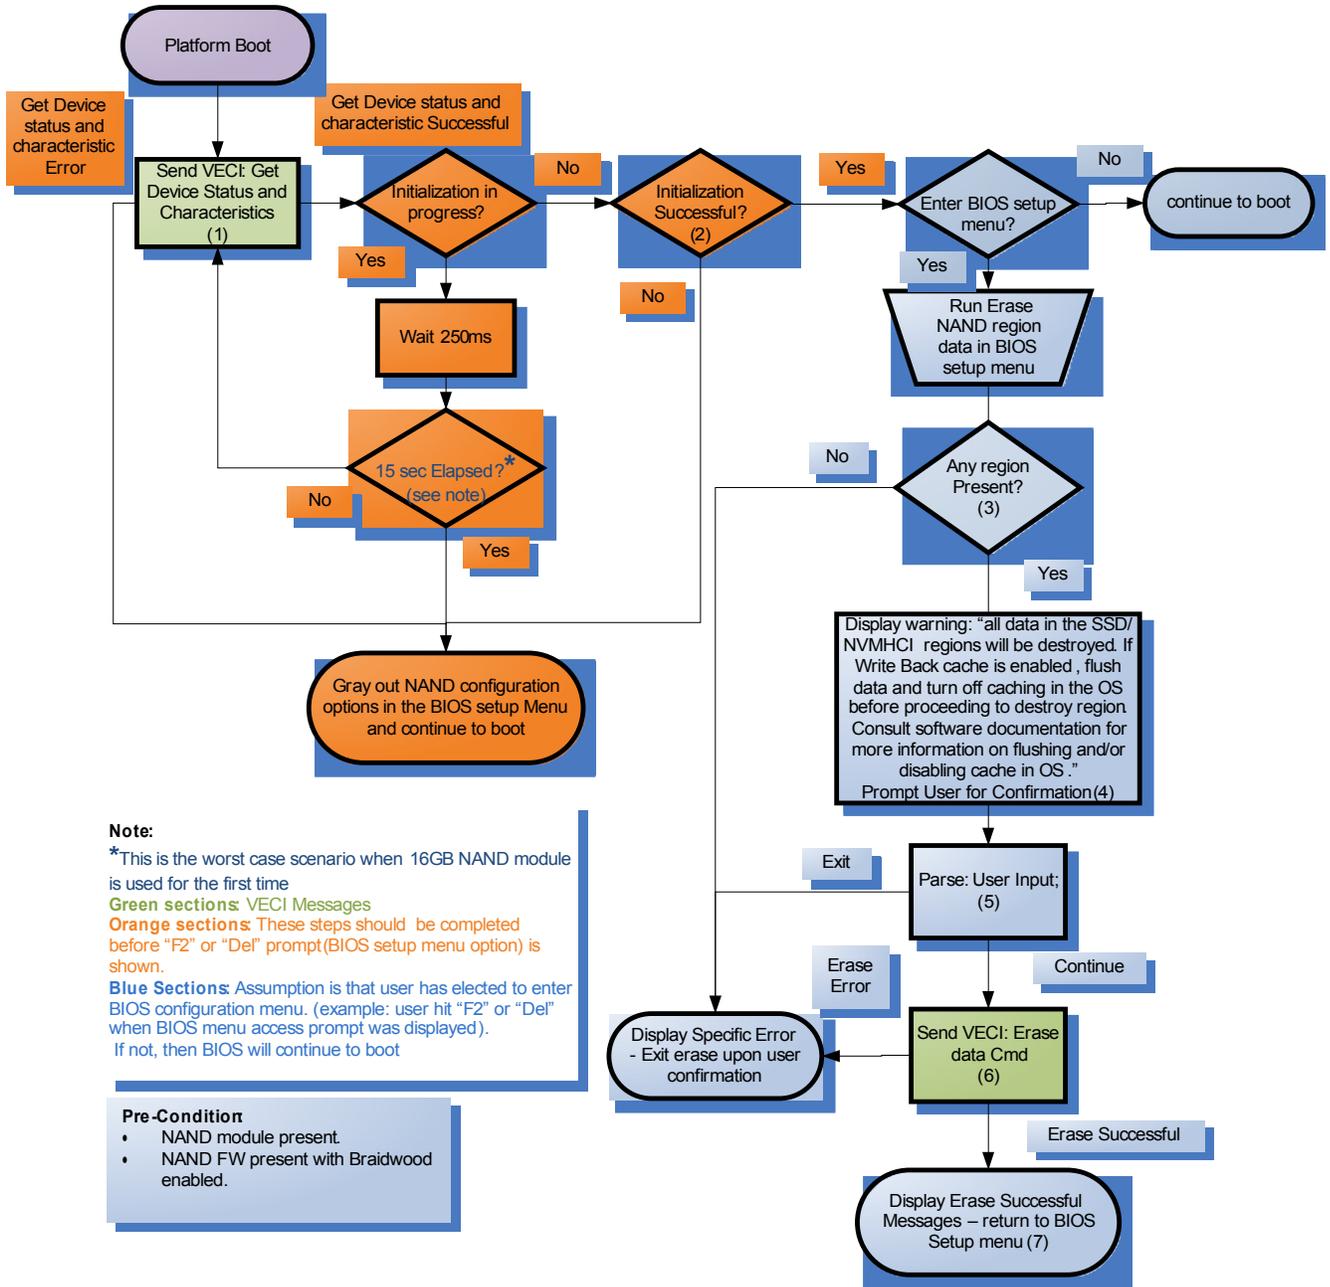
#### Erase:

This section shows a sample flow (see Figure 5: BIOS Option - Erase NAND) for the erase data in the NAND module. The erase option is located in the BIOS set up menu. This option will be grayed out if VECI failed to initialize or there is no NAND module connected in the system.

1. Check device status and characteristic
2. Verify NAND module has initialized correctly.
  - If NAND module has initialized correctly, display the NAND configuration option in the BIOS setup menu
  - If not, wait for 250 ms. If more than 15 seconds have elapsed and NAND module has not initialized, gray out the NAND configuration option and continue to boot.
3. Query existing regions in the NAND module
  - If there is no region present in the NAND module, display error messaging and exit NAND configuration option.
  - If NVMHCI region is already present in the NAND module, prompt user for option to erase region (step 4).
4. Prompt user for erase region warning.
5. Validate user request is valid. If not, exit NAND configuration option.
6. Send VECI erase message.
7. Validate response. If there is no error, display "erase successful" messages.

**Note:** all data in the NAND regions will be destroyed

## BIOS Option - Erase



**Figure 5: BIOS Option - Erase NAND**

BIOS Manufacturing Option:



This section shows a sample flow (see Figure 6: BIOS Manufacturing Option) for the BIOS manufacturing option. This sample flow provides the guidelines for IBVs and OEMs to implement the option in the BIOS. This option allows OEMs to configure the NAND module in the manufacturing line without adding an extra tool in their line. The BIOS is automated and can be used in the manufacturing line for pristine and non-pristine NAND modules. OEMs can reuse this BIOS option on their return/repair line process to bring the NAND module on the return system back to the manufacturing configuration state.

1. Platform Boot
2. NAND configuration is set using scripts or other methods
3. On the next system reboot
4. BIOS checks if NAND Configuration flag is set.
  - If set, go through NAND module configuration process (step 5)
  - If not, continue to boot

This step allows OEMs to control when the BIOS needs to go through this configuration process. This will prevent BIOS from going through the configuration process and configuring the NAND module multiple times during the manufacturing line. It also prevents this configuration process running in the BIOS after the system has been shipped.

5. BIOS checks if the system is in a manufacturing mode.
  - If it's in manufacturing flow, continue with the NAND module configuration process (step 6).
  - If not, exit the NAND module configuration process and continue to boot.

This step is optional and has the same function as the NAND configuration flag. It will prevent the BIOS from configuring the NAND module multiple times

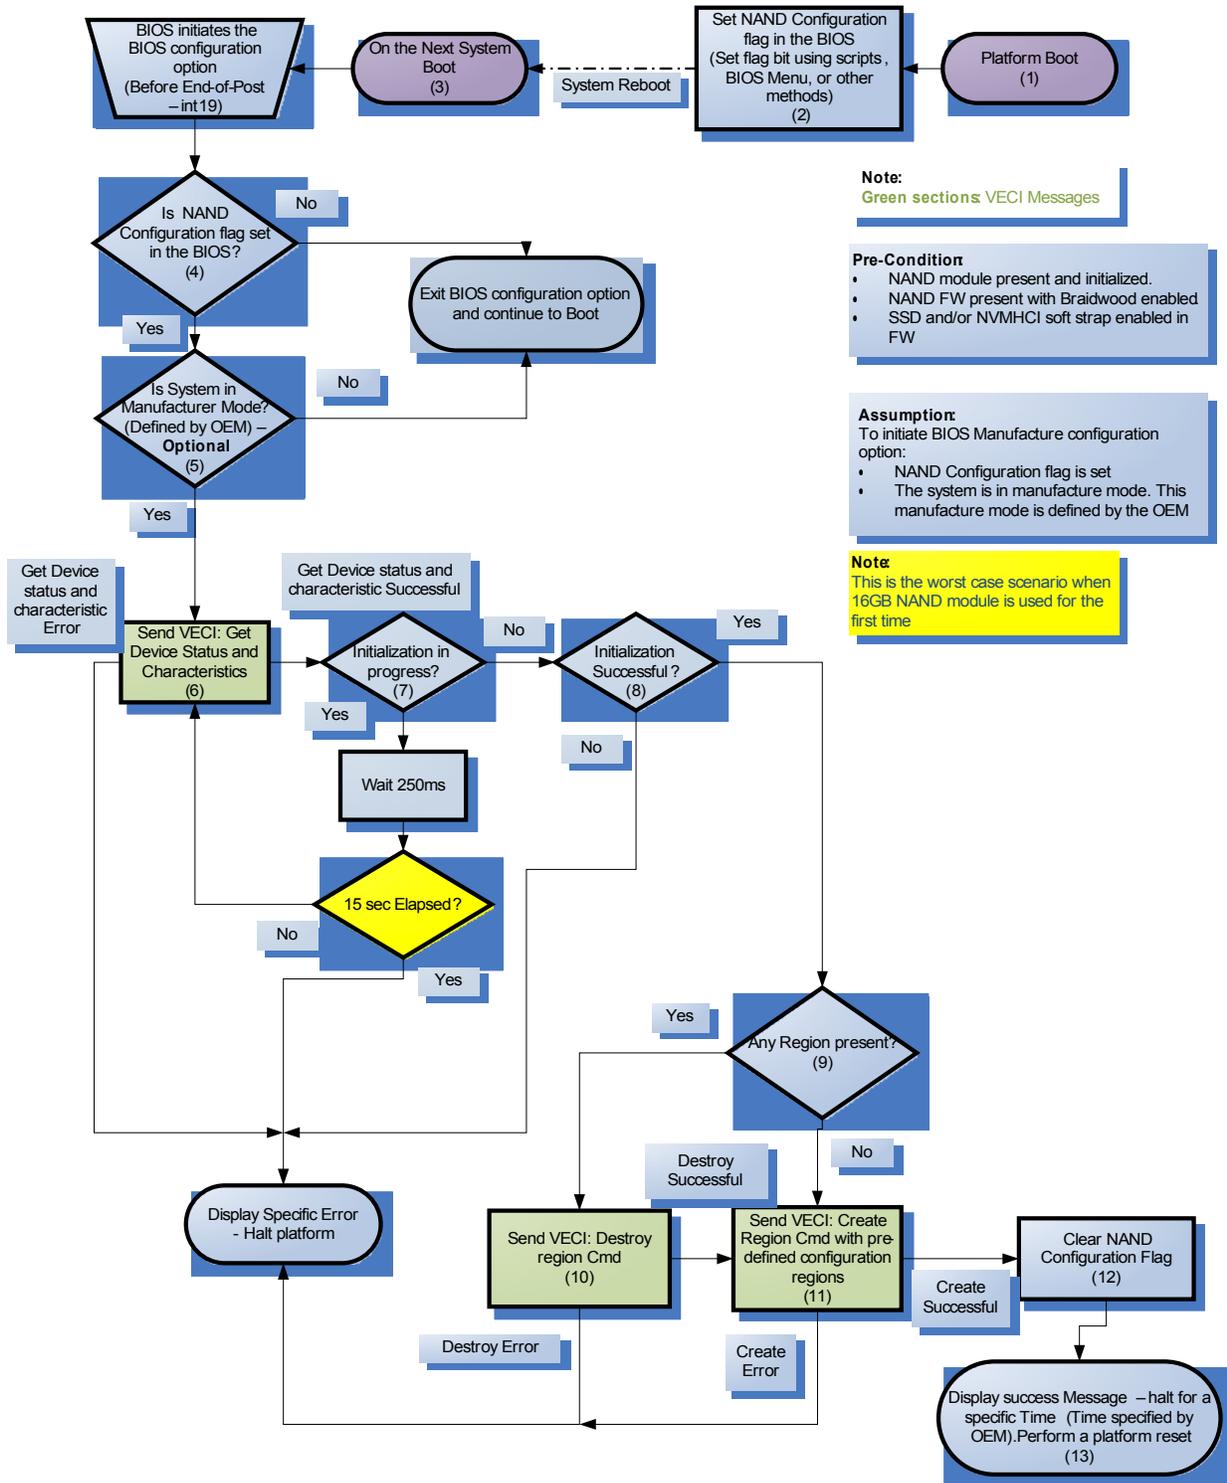
6. Check device status and characteristic
7. Verify if NAND module is still in the initialization process. Initialized correctly.
  - If NAND module has completed initialization process, continue with the NAND configuration process.
  - If not, wait for 250 ms and check again. If more than 15 seconds have elapsed and NAND module has not been initialized, exit the NAND configuration process and display error messages and halt system.
8. Verify if NAND module has initialized correctly.
  - If NAND module has initialized correctly, continue with the NAND configuration process.
  - If not, exit the NAND configuration process and display error messages and halt system.
9. Query existing regions in the NAND module
  - If there is no region present in the NAND module, create regions with predefined setting (step 11)
  - If there is region presents in the NAND module, destroy existing region(s) (Step 10).
10. Send VECI destroy message.
11. Send VECI Create message.



12. Validate response. If there is no error, clear NAND configuration flag.
  - *This will prevent BIOS from going through the NAND configuration process again unless the NAND configuration flag is set.*
13. Display successful messages and continue to boot.



## BIOS Manufacturing option



**Note:**  
Green sections: VECI Messages

**Pre-Condition**

- NAND module present and initialized.
- NAND FW present with Braidwood enabled.
- SSD and/or NVMeHCI soft strap enabled in FW

**Assumption**  
To initiate BIOS Manufacture configuration option:

- NAND Configuration flag is set
- The system is in manufacture mode. This manufacture mode is defined by the OEM

**Note**  
This is the worst case scenario when 16GB NAND module is used for the first time



Figure 6: BIOS Manufacturing Option

## 17.5 Caching Implementation

This section explains the requirements needed in order to setup the NAND module to use its caching functionality via the Intel® Rapid Storage Technology (Intel® RST) software suite.

### 17.5.1 Caching Requirements

The following computer system configurations must be met in order to use the caching functionality of the Braidwood module:

- A RAID-enabled Intel® 5 Series chipset SKU with Braidwood
- RAID support in the system BIOS (the user must be able to set the PCH chipset to RAID mode)
- The Intel® RST 9.5 driver and OROM installed
- The NAND module with the NVMHCI region configured

\*Note: although a RAID-enabled system set to RAID mode via the system BIOS is required for NV caching, the user is not required to create any RAID volumes. Non-RAID Pass-through disks can be used for caching.

For more details on RAID support requirement for caching, please refer to 'Intel Braidwood Technology and RAID Requirement' white paper (document # 421312).

### 17.5.2 Configuring the NAND module for Caching

The NANDUtil tool can be used to create an NVMHCI region in the NAND module which is required for caching. It is recommended that this NVMHCI region is of 3.2GB or larger capacity. For detailed information on NAND region configuration and usage of the NANDUtil tool refer to NANDUtil Tool section.

## 17.6 Caching Features and Configurations

This section gives details on the various features of the NV Cache and the user configurable features using the Intel® RST User Interface (Intel® RST UI).

### 17.6.1 Configuration Rules

The following rules must be followed for caching configuration:

- Only one RAID volume or pass-through disk is allowed to have NV Cache 'On' per computer system
- An Intel® Rapid Recovery Technology (IRRT) volume is allowed to use the NV Cache in 'Enhanced' mode only (the Maximized mode is not allowed)



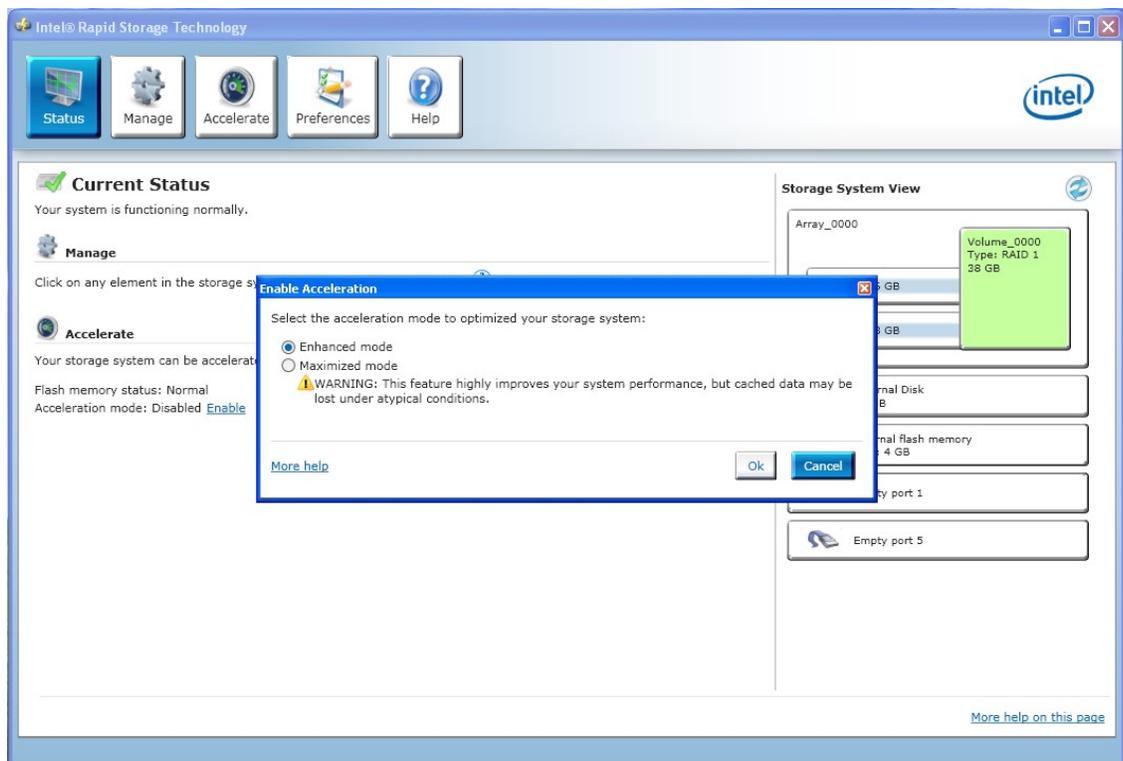
- Matrix RAID volumes (volumes that share their RAID disk array with other volumes) may not have NV Cache 'On'

## 17.6.2 Cache Modes

There are two NV Cache 'On' modes configurable by the user via the Intel® RST UI. The user can select the cache to be used in either Enhanced mode (also commonly referred to as Write Through Caching) or Maximized mode (also commonly referred to as Performance or Write Back Caching). These modes are discussed in greater detail in Enhanced Mode and Maximized Mode sections below. In both modes, the NV Cache provides boot acceleration, high read performance on cache hits, application pinning, and OROM I/O processing for power-fail, pre-driver load I/O, and separation handling.

To set the NV Cache to on, launch the Intel® RST UI and do the following:

1. From the default 'Status' page, go down to the Accelerate section
2. For "Acceleration mode", click on the [Enable](#) hyperlink
3. In the Enable Acceleration pop-up box select the desired acceleration mode, either Enhanced Mode or Maximized Mode (see Figure 7: Cache Modes)
4. Click [Ok] to complete the action; caching is enabled immediately





**Figure 7: Cache Modes**

### Enhanced Mode

Enhanced Mode utilizes the NV Cache as a write through cache where writes are simultaneously stored to both the NV Cache as well as the associated disk/volume. In this mode, both the NV Cache and associated disk/volume are in sync and can be removed from the system without risk of data loss. Performance benefits are realized during booting, for run-time reads, and for paging.

Characteristics of Enhanced Mode are:

1. Cache Coherency: The cached HDD/volume and the NV Cache are always in synch
  - All writes are issued to both HDD and the NV Cache
  - No flushes required on entry into OS power modes S3/S4/S5
2. Performance: lower performance of the two NV Cache modes
  - Lower write bandwidth, all writes slow down to HDD write speed
3. Power Savings: lower power savings in DC mode of the two NV Cache modes
  - Ability to keep HDD spun down is reduced due to periodic writes by the OS and/or applications
4. Separation Handling: This mode is fully separation safe
  - User can safely migrate HDD/volume separately from the NV Cache NAND module between systems

### Maximized Mode

Maximized Mode utilizes the NV Cache as a write back cache, where writes go to NV Cache and are not immediately stored to the associated disk/volume. The caching policy is discussed in greater detail below. This creates lower IO to the associated disk/volume, which can then lead to power savings inherit in less traffic to the disk/volume. Specific details are discussed below.

Maximized Mode is also referred to as Performance Mode, since this mode provides the most performance when considering both IO and hypothetical DC power savings. Performance benefits are realized during booting, run-time reads and writes, paging. In addition to performance gains this cache mode improves DC power savings.

Characteristics of Maximized Mode are:

1. Cache Coherency: The cached HDD and the NV Cache may not always be in synch
  - Writes are issued to HDD and/or to the NV Cache
  - Lazy writing opportunistically copies cache data to the HDD in the background when it's not spun down
  - Manual flush required to synchronize HDD with NV Cache
2. Performance: Higher performance of the two NV Cache modes
  - All no-streaming data written directly to the cache
  - Bounded-time to flush dirty data from cache to disk on power transitions
3. Power Savings: Best system power savings in DC mode
  - HDD is spun down more often, all writes captured in the cache



- Writes are batched and written to the HDD only periodic when spun down
4. Separation Handling: This mode is not fully separation safe
- Extra user intervention is required to synchronize “dirty” cache data to the HDD before migration between systems unless the HDD + NV Cache NAND module are moved as a single pair

## 17.7 OROM Implementation Supporting Braidwood

This section covers features and functionality in the OROM required to support Braidwood

### 17.7.1 Core OROM Features and Functionality

#### 17.7.1.1 Dual Boot

It is possible for the user to remove a disk that is cached and use it on a different system that doesn't know about the Intel® RST cache mechanism. It is also possible for the user to boot the system into an OS that doesn't have Intel® RST driver support and access the disk, bypassing the cache.

To determine dual boot, the OROM sets a flag if any INT13 occurred which would imply that some OS boot has occurred. If the driver has not reset the NVM flag, then the OROM can assume that some other non-NVM OS boot has happened.

It is also possible to roam a cached disk to a non-NVM supported OROM. The volume will be shown as “Incompatible”.

The effect that dual booting has on the accessibility of RAID volumes or pass-thru disks with NV Cache On is as follows. Dual booting is defined as booting into OS 'a' then re-booting into OS 'b' on the same system, regardless of whether OS 'b' resides on the same or a different partition, volume, disk or removable media from OS 'a'. See the table below for the behavior of dual boots.

Cache Clean Required	Driver on OS 'b'	Accessibility on OS 'b'	Accessibility upon re-boot to OS 'a'
Yes	RST 9.5 or later or Linux / Doesn't use Int13	Normal	Normal, but cache reset <sup>3</sup>
Don't care	RST 9.5 or later or Uses Int13	Normal	Normal
No	RST 9.5 or later or Linux / Doesn't use Int13	Incompatible disks / no access	Normal



Don't care	Non-IMSM / Non-Linux / Doesn't use Int13	No file system <sup>1</sup>	File system may be lost (unpredictable) <sup>1</sup>
Yes	Non-IMSM / Non-Linux / Doesn't use Int13	Normal <sup>2</sup>	If pass-thru then normal but cache reset, if R1 then unpredictable <sup>2</sup>

**Table 1: Dual boot behavior**

<sup>1</sup> If pass-thru or RAID-1 disk then there's an exception, see next row.

<sup>2</sup> Only if pass-thru or RAID-1 disk.

<sup>3</sup> "Cache reset" warning.

If OS 'a' resides on a separate volume or disk from OS 'b' and neither OS touches the files on the other OS's volume or disk then dual boot would succeed giving normal access to the volumes and/or disks.

### 17.7.1.2 Roaming/Separation

The only two ways for the user to safely roam a pass-thru disk or the disks of a RAID volume with NV Cache in Maximized mode to another platform are the following.

1. A 'Cache Clean' will have to be successfully completed in the OROM since the last boot.
2. Or the NVM part will have to be roamed in conjunction with the cached disks to another platform that is either running the RST 9.5 driver or using Int13 on an IMSM OROM version >= 9.5.

**Caution:** *When a system is on AC power and it is idle for a long period of time it may complete a cache clean the same way an explicit OROM cache clean does. In such cases roaming will be legal, but there are no guarantees about such a clean completing so it is not considered safe. See the table below for the results of various forms of disk roaming.*

**Warning:** *When the user illegally roams a volume or disk to another system with an RST 9.5 OROM and driver the following will happen.*

1. The volume or disk will appear as a failed volume on the new system. The user may access the failed volume by ordering a recover volume in the OROM or RST UI, but if the user does that then the NV Cache data is forgotten and lost from the volume or disk even if the disk is reunited with its NV module, **so data loss occurs.**
2. Or the user may choose not to "Recover" the volume or disk. In that case the volume or disk will remain failed, but it will be able to run normally with no data loss if it is reunited with its NV Module.

Clean	NVM Part roamed with	Driver on other system	Accessibility on other	Accessibility upon return
-------	----------------------	------------------------	------------------------	---------------------------



	disk(s)		system	
Yes	No	RST 9.5 or later or uses Int13	Normal	Normal, but cache reset <sup>3</sup>
Yes	Don't care	RST 9.5 or later or Linux or uses Int13 but <9.5 OROM	Normal	Normal, but cache reset <sup>3</sup>
Don't care	Yes	RST 9.5 or later or uses Int13	Normal	Normal
No	No	RST 9.5 or later or uses Int13	Failed volume / no access <sup>4</sup> , recovery means to get access but lose cached data <sup>5</sup>	If not recovered then normal, if recovered then normal access but cached data lost <sup>6</sup>
No	Don't care	RST 9.5 or later or Linux or uses Int13 but <9.5 OROM	Incompatible disks – no access	Normal
Don't care	Don't care	Non-IMSM / Non-Linux / Doesn't use Int13 and/or no IMSM OROM	No file system <sup>1</sup>	File system may be lost (unpredictable) <sup>1</sup>
Yes	Don't care	Non-IMSM / Non-Linux / Doesn't use Int13 and/or no IMSM OROM	Normal <sup>2</sup>	If pass-thru then normal but cache reset <sup>3</sup> , if R1 then unpredictable <sup>2</sup>

**Table 2: RAID Volume Disk roaming behavior**

<sup>1</sup> If pass-thru or RAID-1 disk then there's an exception, see next row.

<sup>2</sup> Only if pass-thru or RAID-1 disk.

<sup>3</sup> "Cache reset" warning.

<sup>4</sup> "Live dirty cache data missing" warning.

<sup>5</sup> "Live dirty cache data permanently discarded" warning.

<sup>6</sup> "Dangling cache data permanently discarded" warning.

When disks are moved to a new system without their NVM part the table above describes the resulting accessibility. In all such cases where the volume is accessible (not failed or incompatible) on the new system, the volume runs with NV cache Off. If the new system has the RST 9.5 or later driver then a "Missing live clean cache data" warning will be given. If those disks are re-united with their NVM part then they'll have a config mis-match with their NVM part.

If a system has a volume or pass-thru disk with NV Cache On and the disks of that volume or that pass-thru disk is / are removed but the NVM part is not removed then the following will occur on the system with the NVM part. There will be dangling NV Cache data. The RST UI will display information about the dangling NV Cache data



and give a “Dangling dirty NV Cache data” or “Dangling clean NV Cache data” warning as appropriate, but no volume or disk associated with the dangling NV Cache data will be created, accessible or visible. The user will have the ability to discard the dangling NV Cache data via the RST UI. If the user discards the dangling NV cache data then if the volume disks are returned the NVM part will have a configuration mismatch with the disks. The OROM will ignore dangling NV Cache data and not offer any handling thereof.

If the system is downgraded from supporting NVM to not supporting NVM regarding the system requirements then it is like roaming or separation. Loss of a system requirement is the same as an equivalent roaming case. For the purposes of mapping a loss of a system requirement to a roaming case consider the following. Losing the PCH platform or downgrading to pre-Vista is like losing the NVM part itself. If the RAID key or the BIOS SATA = RAID is lost then for any RAID volume or pass-thru disk with dirty data it is like running a RAID volume in AHCI-only mode. For a pass-thru disk with clean data it is like roaming to a pre-9.5 driver.

### 17.7.1.3 NV Cache to Disk Configuration Mismatch

The NV Cache configuration data is kept on both the NV module and the disks. The serial numbers of each are logged. Therefore the case where an associated NV module and disk have a configuration mismatch can be distinguished from the case where an NV module and/or disk are illegally removed and/or replaced. The following table describes the behavior of what happens if the disk(s) and NV module are both present and associated with each other but still do not agree on whether or not a particular volume or disk has NV Cache On.

This kind of mismatch can occur in any of the following cases:

1. The NVM module is erased
2. The NVM module is moved to another system or run on another OS, re-configured, then moved back or run back on this OS,
3. The disks are moved to another system, reconfigured, then moved back

Disk Config	NV Cache Config	Data dirty	Result
Volume/disk 'a' has NV Cache On	No info about volume/disk 'a'	Yes	Volume/disk 'a' runs with NV Cache Off. But possibly catastrophic data loss. <sup>1</sup>  NV module is available for NV Cache to be turned on.
Volume/disk 'a' has NV Cache On	No info about volume/disk 'a'.	No	Volume/disk 'a' runs with NV Cache On. No data loss. But cache reset. <sup>2</sup>  NV module is available for NV Cache to be turned on.



Volume/disk 'a' does not exist.	Volume/disk 'a' has NV Cache On.	Don't Care	Dangling cache data. NV Cache cannot be turned on for any volume or disks until the dangling cache data is discarded. <sup>4</sup>
Volume/disk 'a' exists but does not have NV Cache On	Volume/disk 'a' has NV Cache On	Don't care	Volume/disk 'a' runs with NV Cache On, but cache reset. <sup>2</sup>
Volume/disk 'a' has power-cycle count 'x'. It matches NV module's config in every other way.	Volume/disk 'a' has power cycle count <x-1. It matches disks' config in every other way.	No	Volume/disk 'a' runs with NV Cache On, no data loss, but cache reset <sup>2</sup>
Volume/disk 'a' has power-cycle count 'x'. It matches NV module's config in every other way.	Volume/disk 'a' has power cycle count <x-1. It matches disks' config in every other way.	Yes	Volume/disk 'a' runs with NV Cache On.  If this power cycle count mismatch is due to roaming to a non-RST driver then there could be a catastrophic cache data to disk data incoherency  If this power cycle count mismatch is due to just an extra reboot before it could be accounted for in MD then all data is intact.
Volume/disk 'a' has NV Cache On but is associated with another NV module	NV module is associated with volume or disk 'a' and has NV cache On	Don't care	Dangling cache data on the NV module discarded. <sup>3</sup>

**Table 3: Disk(s) to NV module configuration mismatch**

<sup>1</sup> "Live dirty cache data permanently discarded" warning.

<sup>2</sup> "Cache Reset" warning.

<sup>3</sup> "Dangling cache data discarded" warning.

<sup>4</sup> "Dangling NV Cache data" warning.

## 17.7.2 OROM UI Features and Functionality

There are features in the OROM UI that are there specifically to support NVM. These are:

- NVM status



- a user option to view accelerated Disk/Volume information
- 'Synchronize' and 'Remove' acceleration actions.

After every NVM UI action, the OROM UI will re-enumerate all devices and refresh the UI.

### 17.7.2.1 Main Window

#### 1. Accelerated "\*" Indicator

To indicate to the user that a disk and/or Volume is being cached using NVM aka "Accelerated", a green "\*" will be displayed before the port and/or ID number, respectively. Since Danbury and NVM Caching are mutually exclusive, there will be no display conflicts with using an asterisk in the same manner that was implemented for Danbury.

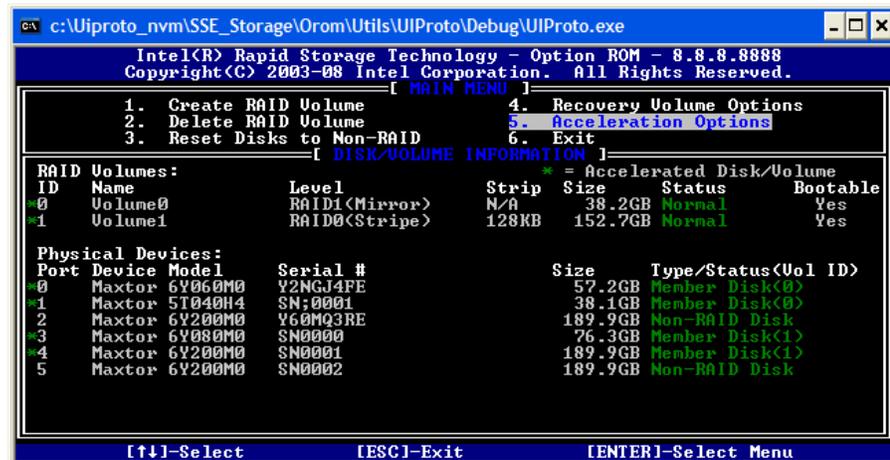


Figure 8 Acceleration Disk/Volume \* Indicator

#### 2. Disabled Status

The only status related to NVM Caching that will be displayed on the main window will be the status of "Disabled". Since the Status on the main window has been used for RAID volume status, it was determined that overloading this with the NVM status would be confusing to the user including having to identify which takes precedence, a Volume status vs a NVM status. Therefore, only a "Disabled" status will be shown in the main window because this status requires user interaction to rectify.

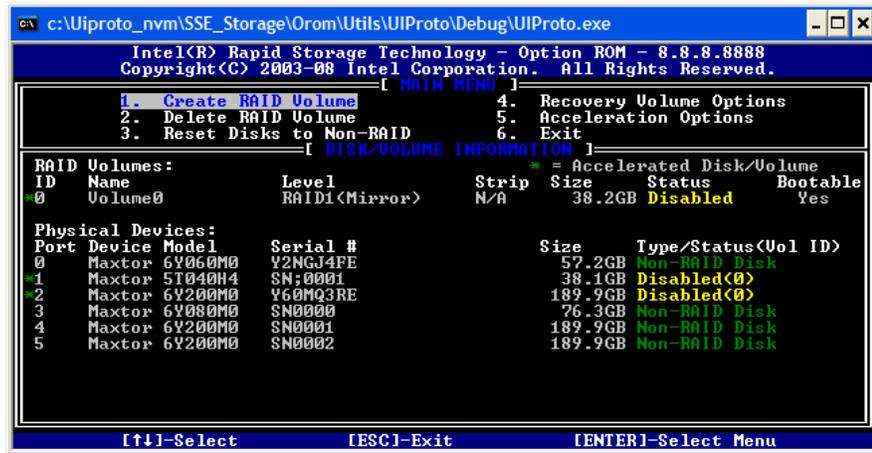


Figure 9: Disabled Acceleration Volume

### 17.7.2.2 Acceleration Options Window

An Acceleration Options Window will be added that allows the user to synchronize data, remove acceleration and it will provide more detailed information about the Accelerated Disk/Volumes on the system. A new option at the main menu will be added called "Acceleration Options" that takes the user to the Accelerations Options Window when selected.

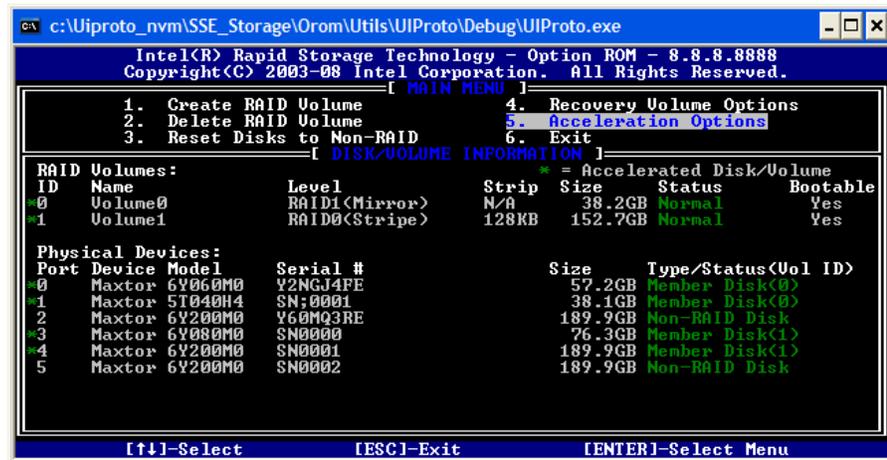


Figure 10: Addition of Main Menu Item "Acceleration Options"

If there are no cached objects on the system, then an error will be shown if the user tries to enter the window.

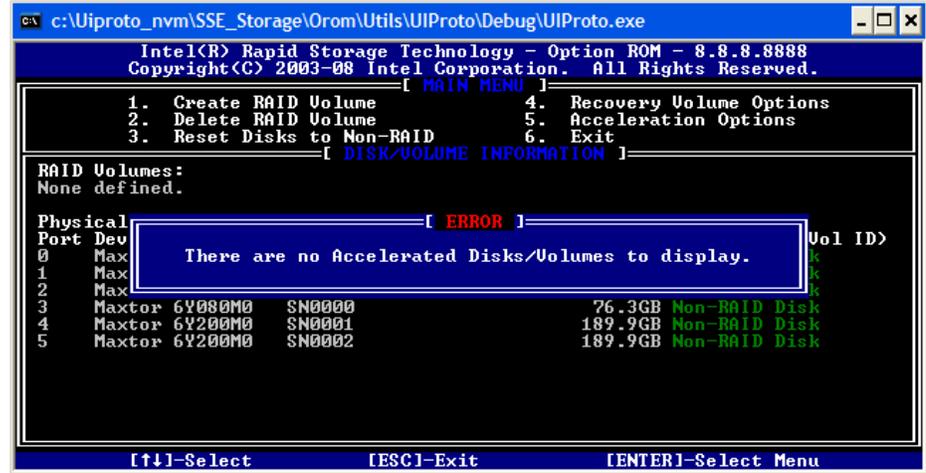
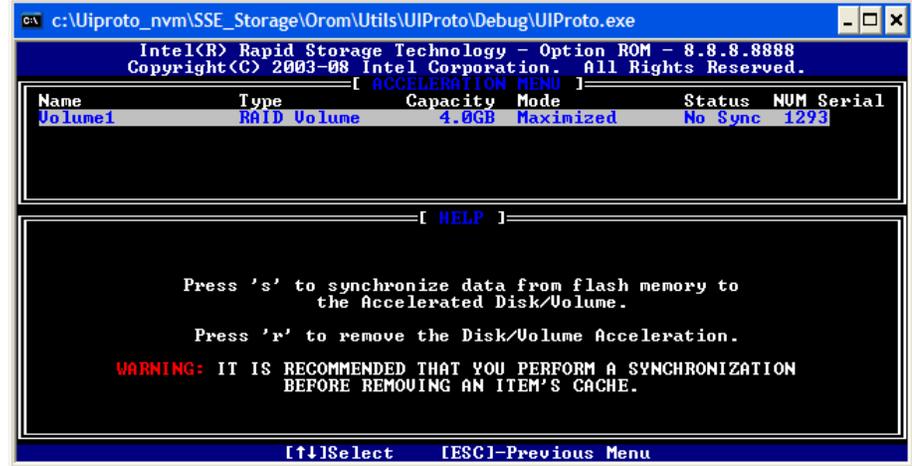


Figure 11: Error: no Accelerated Disks/Volumes to Display

#### 1. List of Accelerated Objects

The main window of the "Acceleration Options" window will display a list of Accelerated objects. The list can contain Volumes and non-RAID disks. For each Accelerated object, the following information will be displayed:

- Name** - "DISK PORT #" for SDR0, Volume Name for Volume
- Type** - non-Raid Disk or RAID Volume
- Capacity** - Size of cache in MB or GB if > 100MB.
- Mode** - Mode of cache (see Mode strings section)
- Status** - Status of Cache: In Sync or No Sync
- NVM Serial** - last 4 digits Serial number of NVMHCI part it is associated with



**Figure 12: Acceleration Options Window**

a. NVM Mode Strings

The following table contains the list of NVM Mode Strings that will be shown according to the NVM State of the Accelerated object.

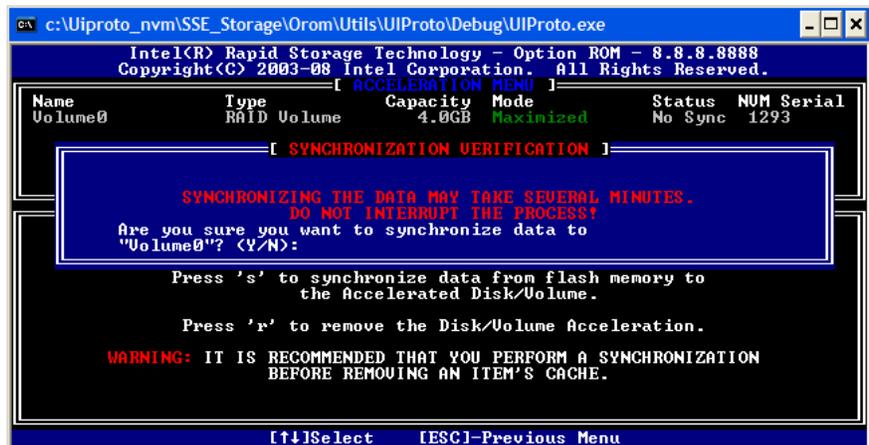
**Table 4 NVM Mode OROM UI String**

<b>CfgRaidDev::NvmCacheMode</b>	<b>Mode String</b>
NV_CACHE_MODE_DIS_PERF NV CACHE MODE DIS SAFE	Disabled
NV_CACHE_MODE_PERF_TO_OFF NV_CACHE_MODE_PERF_TO_SAFE NV_CACHE_MODE_OFF_TO_SAFE NV CACHE MODE OFF TO PERF	Transitioning
NV CACHE MODE SAFE	Enhanced
NV CACHE MODE PERF	Maximized
NV CACHE MODE IS FAILING	Failing
NV CACHE MODE HAS FAILED	Failed

2. Synchronize Data Action

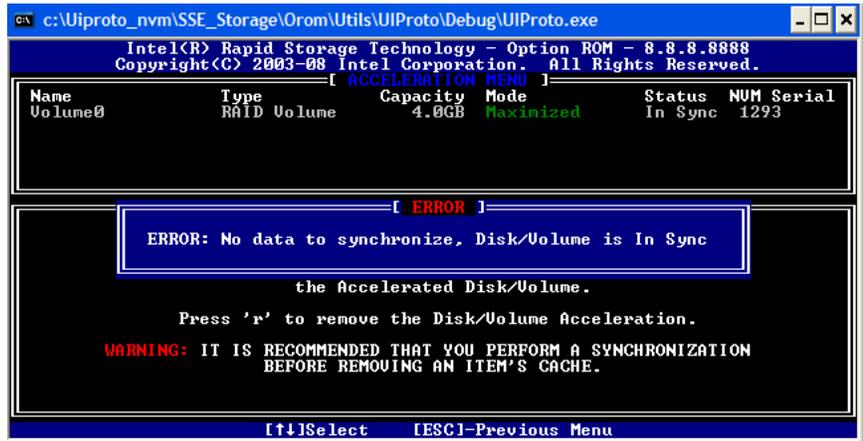


Only an Accelerate Disk/Volume with a “No Sync” status can be selected to be synced. In order to synchronize the data, the user must select the item in the list and hit the ‘S’ key. This will bring up a window asking the user to confirm that they wish to synchronize the data. If the user hit’s the ‘Y’ key, then a cache synchronization will begin on the selected item. Synchronization copies the data from the flash memory to the Accelerated Disk/Volume.



**Figure 13: Synchronize Data Verification Window**

The action will only be performed for a cached object of status “No Sync”. If the user presses the ‘S’ key on a selected item with a status other than “No Sync”, a message to the user will be shown that the action is not applicable for the selected item.



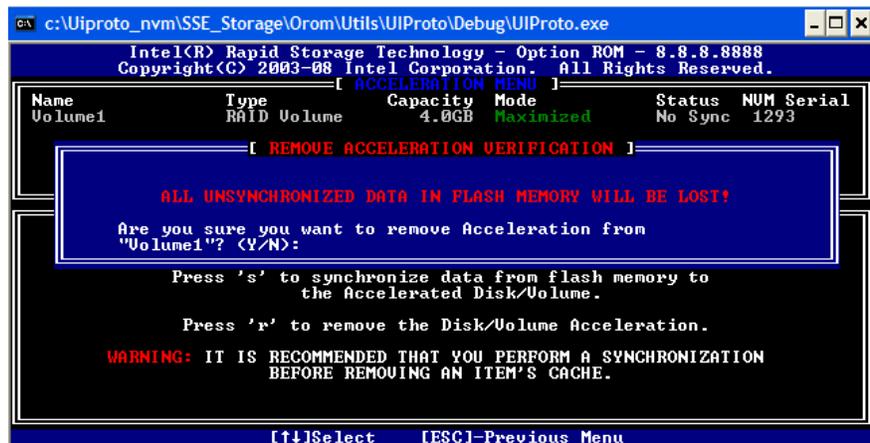
**Figure 14: Error: Sync action not available**

### 3. Remove Acceleration Action

In order to remove the acceleration, the user must select the item in the list and hit the 'R' key. This will bring up a window asking the user to confirm that they wish to remove the cache. If the user hit's the 'Y' key, the acceleration information will be removed from the Accelerated disk/volume and the information will be removed from the flash memory.

Removing the cache does not perform a synchronization. The user must synchronize the data using the synchronize action before removing the cache if they wish to save the data contained in the flash memory onto the disk/volume.

If the user deletes the cache for a single-disk RAID0, then the metadata will be completely removed from the disk/volume.



**Figure 15: Delete Cache Verification Window**

If the user removes the Acceleration for a non single-disk RAID0, all metadata will be removed from the disk (aka deleteVolume will be called).

If the user removes the Acceleration for a RAID Volume, the CfgRaidDev variables will be updated:

NvCacheModee = NV\_CACHE\_MODE\_OFF;

NvCacheFlags = 0x0;

NvmhciSerial = 0x0;

NvmUiState = NvmOff;

#### 4. Synchronize Data Progress

Since the time required to synchronize the data from the flash memory to the Accelerated Disk/Volume can be significant (a few minutes), the user will be presented a percentage indicating the progress of the synchronization. A user will not be able to perform any other actions in the OROM UI until the synchronization is complete.

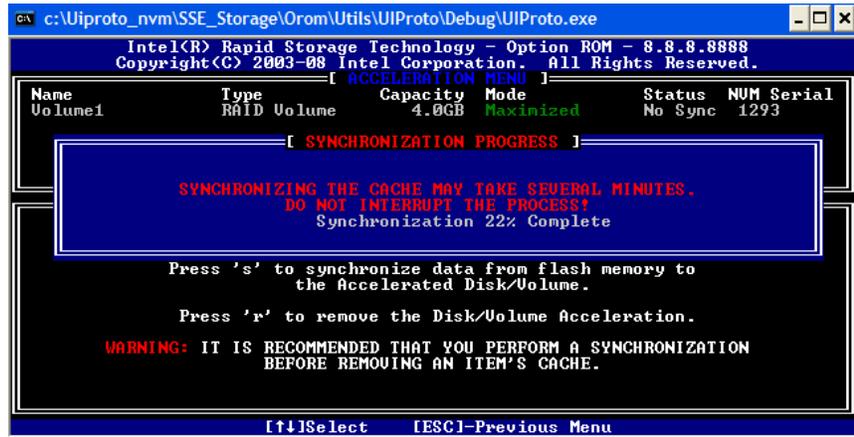


Figure 16: Progress Percentage

### 17.7.2.3 User Required NVM Prompts

There are specific scenarios that will require a user to make a decision concerning their cached disk/volume before it is usable. During enumeration, the disk/volumes will be disabled, and the user must enter the OROM UI to enable these cached disk/volumes. Otherwise, they will have a status shown as Disabled and will not be accessible.

#### 1. Roamed Object Prompt

If the user connects a disk/volume that is not being cached on the system it is connected to, then the user will be queried to remove the cache from the disk/volume or to keep the disk/volume disabled.

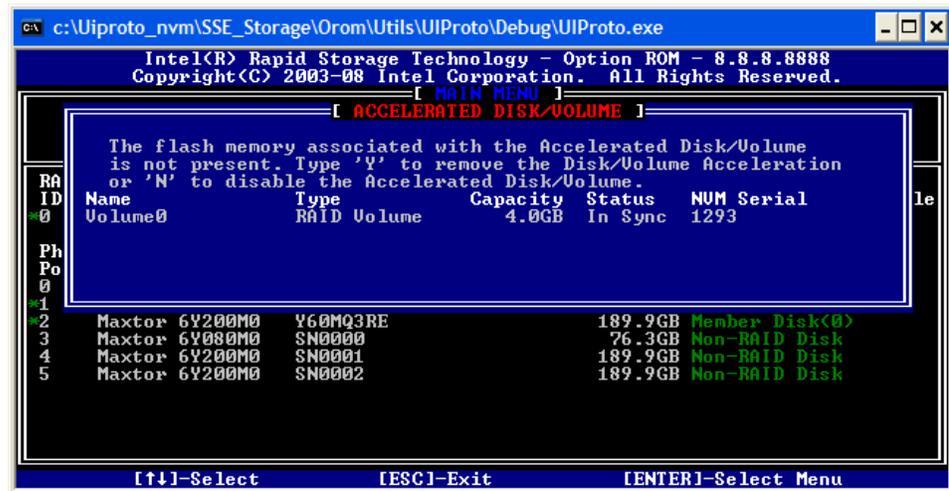
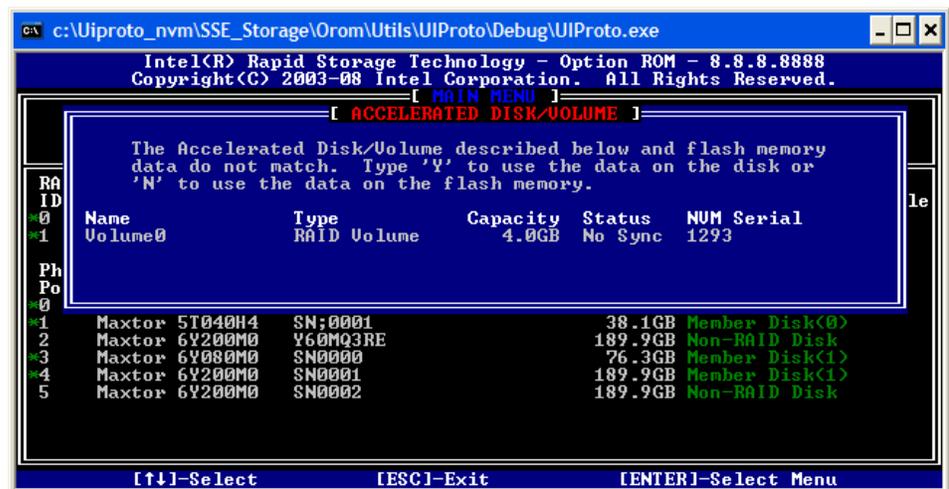


Figure 17: User Prompt NVM not present for Accelerated Disk/Volume

2. Power CycleCount Mismatch Prompt

During the enumeration, the OROM will compare the power cycle count of the NVM's information about the cached disk/volume to the information of the cached disk/volume's metadata. If there is a difference greater than 1, then the user will be queried to choose to use the cached disk/volume's contents or to use the NVM cache data contents.



3. Figure 18: User Prompt Power Cycle Count Mismatch Last Touched by OROM Prompt

During the enumeration, the OROM will check the lastInitByOROM flag in the NVM cache. If it is still set (indicating the driver didn't update the NVM cache) then the user will be prompted to choose to use the cached disk/volume's contents or to use the NVM cache data contents.

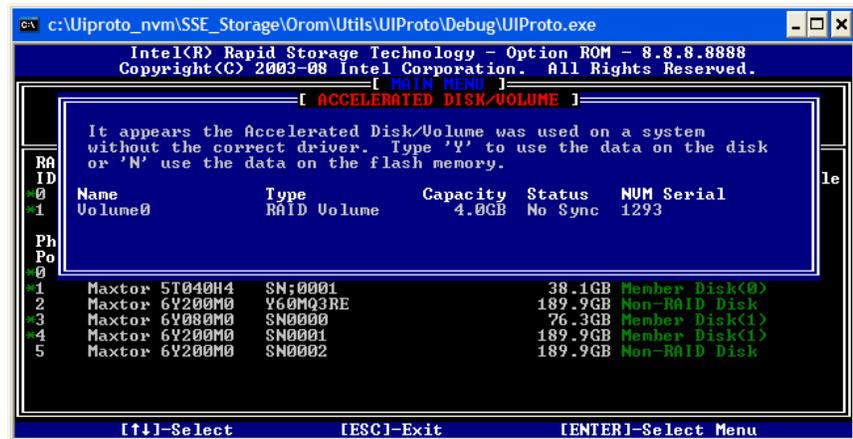


Figure 19: User Prompt Last Touched by OROM

#### 17.7.2.4 Delete Volume Menu

An Accelerated disk/volume will not be shown in the Delete Volume list. Only Accelerated and non-Accelerated Volumes will be shown in the Delete Volume list. If an Accelerated disk/volume is deleted from the Delete Volume Menu, it will not synchronize the data before deleting the Accelerated disk/volume. The user should synchronize the data for RAID 1 and IRRT volumes using the Acceleration Options menu first.

#### 17.7.2.5 Reset Disks to Non-RAID Menu

An Accelerated disk/volume (SDR0) will not be shown in the disk list for the Reset Disks to Non-RAID menu. This is because it is shown as a non-RAID Disk in the main window. Instead, the user must use the Accelerated Options to remove the acceleration if they wish to remove all metadata on the disk/volume and be able to roam it to an older OROM/driver.



## 18 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



<b>Term</b>	<b>Definition</b>
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 striped across the hard drives to provide good read/write performance and data redundancy. The parity is striped in a rotating sequence (aka Striping 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 Striping 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