

Intel® Virtual RAID on CPU (VROC) for Windows*

Technical Product Specification

April 2019
Revision 1.8



Revision History

Revision	Description	Date
001	Initial release.	May 2017
002	Updates to support Intel RSTe 5.1 PV Release	May 2017
003	Updates to support Intel RSTe 5.2 PV Release	June 2017
004	Updates to support Intel RSTe 5.3 PV Release	October 2017
005	Updates to include support for Intel® Xeon® Processor D-2100 Product Family	January 2018
006	Updates to support Intel RSTe 5.5 PV Release	August 2018
007	Updated to reflect the name change to Intel VROC	December 2018
008	6.1 Release	April 2019

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

No computer system can provide absolute security. Requires an enabled Intel® processor, enabled chipset, firmware and/or software optimized to use the technologies. Consult your system manufacturer and/or software vendor for more information.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. Check with your system manufacturer or retailer or learn more at intel.com.

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

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

All products, computer systems, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

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

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

Copyright © 2019 Intel Corporation. All rights reserved.



Contents

Revision History	2
1 Introduction	5
1.1 Intended Use	5
1.2 Intended Audience	5
1.3 Terminology.....	6
2 Product Overview	9
2.1 Product Release Numbering Scheme	9
2.2 Supported Operating Systems per Platform.....	10
2.3 Supported Platforms/Chipsets/SKUs.....	11
2.4 New Features Introduced in Intel VROC 6.1 Release.....	12
2.5 New Features Introduced in Intel VROC 6.0 Release.....	13
2.6 Features Introduced in Intel RSTe 5.5 Release.....	14
2.7 Features Introduced in Intel RSTe 5.4 Release.....	15
3 Intel VROC 6.1 Package Components.....	16
3.1 Intel VROC 6.1 CLI.....	16
3.2 Intel VROC 6.1 OS Installer	16
3.3 Intel VROC 6.1 Environment	16
3.4 Intel VROC (SATA RAID) Environment.....	19
3.5 Intel VROC (NonVMD NVMe RAID) Environment.....	23
4 Intel® VROC Key Features.....	24
4.1 Intel® Volume Management Device (Intel® VMD).....	24
4.2 Intel VROC (VMD NVMe RAID) Pre-OS Features and Functionality	28
4.3 Intel VROC 6.1 Monitoring and Management	32
4.4 Intel VROC Features and Functionality	32
5 Intel VROC 6.1 Common Features.....	42
5.1 Intel VROC 6.1 RAID Features and Functionality.....	42
5.2 Intel VROC 6.1 Pre-OS Features and Functionality	65
5.3 Intel VROC 6.1 Monitoring and Management	72
5.4 Intel VROC 6.1 Power Management	109
5.5 Intel VROC 6.1 Miscellaneous Features and Functionality	109
5.6 Intel VROC Package Updates.....	115
6 Intel® VROC (SATA RAID) Key Features.....	117
6.1 Intel VROC (SATA RAID) Pre-OS Features and Functionality.....	117
6.2 Intel VROC (SATA RAID) Miscellaneous Features and Functionality	123
7 Intel VROC (NonVMD NVMe RAID) Key Features	126
7.1 PCIe Intel NVMe Device Usage Model	126
7.2 Intel VROC (NonVMD NVMe RAID) Boot Support	126



7.3	Intel VROC (NonVMD NVMe RAID) Installation	126
7.4	Intel VROC (NonVMD NVMe RAID) Pass-through IOCTL Support.....	126
7.5	Intel NVMe SSDs Support List	127
7.6	Intel VROC (NonVMD NVMe RAID) RAID Volume Support with Multiple Controllers	127
7.7	Intel VROC (NonVMD NVMe RAID) Feature Limitations	127
Appendix A: Related Documentation		128
Appendix B: Storage System Events.....		129
Appendix C: External Hardware Capability		133
Appendix D: Intel C620 Series Chipset Legacy OROM Boot Option		136
Appendix E: VROC (SATA RAID) Port Bitmap Implementation		138
	Legacy OROM.....	138
	UEFI Driver.....	139
	EFI_DEVICE_PATH_PROTOCOL	139
Appendix F: Troubleshooting		140
	Failed Volumes.....	140
	Degraded Volumes.....	142
	Other Volume States	143
	Disk Events.....	145
	Software Errors	147



1 Introduction

This is the Technical Product Specifications for the Intel® Virtual RAID on CPU (Intel® VROC) 6.1 family of products. The intent of this document is to present the functional requirements (or features) and technical details that make up these products. The features outlined will include those for the Pre Operating System (Pre-OS) components, the Windows* based drivers and tools, the Intel Accelerated Storage Manager (ASM) web-based remote management tool and the Graphical User Interface (GUI).

With the release of Intel VROC 6.1, the name of the Intel Rapid Storage Technology enterprise (RSTe) family of products has been changed to the following:

- Intel VROC (VMD NVMe RAID) – Previously referred to as Intel VROC, this is the product to support NVMe SSDs attached to an Intel Volume Management Device (VMD) controller.
- Intel VROC (SATA RAID) – Previously referred to as Intel RSTe, this is the product that supports SATA drives attached to the platform's Intel Platform Control Hub (PCH) (SATA and/or sSATA) controller configured in RAID mode.
- Intel VROC (NonVMD NVMe RAID) – Previously referred to as Intel RSTe NVMe, this is the product that supports DATA RAID on Intel NVMe SSDs attached to the platform CPU. This product ONLY applies to supported platforms with Intel CPUs that do not contain Intel VMD technology.

NOTE: This product cannot be installed or used on platforms that have CPUs that contain the Intel VMD technology, whether enabled or disabled.

1.1 Intended Use

This document is intended to provide high level information on the technical features of the Intel VROC 6.1 family of products for Intel VROC (VMD NVMe RAID), Intel VROC (SATA RAID) and Intel VROC (NonVMD NVMe RAID).

1.2 Intended Audience

The intended audience of this document is OEMs and ODMs that require detailed information about the features and technical specifications of the Intel VROC 6.1 family of products. It contains information pertaining to the Intel VROC 6.1 family of products include:

- Intel VROC (VMD NVMe RAID)
- Intel VROC (SATA RAID)
- Intel VROC (NonVMD NVMe RAID)
- Intel VROC PreOS packages
- Intel VROC GUI and other management tools



1.3 Terminology

Term	Definition
AER	Advanced Error Reporting
AHCI	Advanced Host Controller Interface
AIC	Add In Card
API	Application Programming Interface
ATA	Advanced Technology Attachment
ATAPI	Advanced Technology Attachment Packet Interface
ASM	Intel® Accelerated Storage Manager
BIOS	Basic Input / Output System
Chipset	Term used to define a collection of integrated components required to make a PC function.
CLI	Command Line Interface
CSMI	Common Storage Management Interface
Disk Coercion	Applying the coercion scaling factor to the disk to reduce the amount of available capacity on the disk when integrated within a system.
DOS	Disk Operating System
DSJ	Dirty Stripe Journaling
Disk's Write Cache	A memory device within a drive, which is allocated for the temporary storage of data before that data is copied to its permanent storage location.
EBDA	Extended BIOS Data Area
EN	Entry Server
EP	Efficient Performance
FUA	Flush Unit Access
GB	Giga-byte
GUI	Graphical User Interface
HDD	Hard Disk Drive
HII	Human Interface Infrastructure
Hot Plug	A term used to describe the removal or insertion of a SATA drive when the system is powered on.



Term	Definition
Hot Spare Disk	A disk flagged so as to be automatically used to rebuild a failed or degraded RAID volume without user interaction.
I/O	Input / Output
JD	Journaling Drive
LPM	Link Power Management
Matrix RAID	Multiple RAID volumes residing within the same Array
MB	Mega-bytes
NAI	Notification Area Icon
NCQ	Native Command Queuing
NTFS	NT File System
NVMe	Non Volatile Memory PCI Express
ODD	Optical Disk Devices
ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
OROM	Option ROM
OS	Operating System
PCH	Platform Control Hub
Pre-OS	Pre Operating System Environment (Legacy OROM and/or UEFI)
Private Protocol	The vendor specific protocol to extract device information from a device managed by the vendor's device driver.
Port	The point at which a SATA drive physically connects to the SATA controller.
PPL	Partial Parity Logging
PRD	Product Requirements Document
RAID	Redundant Array of Independent Disks
RHEL	Redhat Enterprise Linux
RSTe	Rapid Storage Technology enterprise
RWH	RAID Write Hole
SATA	Serial Advanced Technology Attachment
sSATA	Secondary Serial ATA



Term	Definition
SGPIO	Serial General Purpose I/O
SMART	Self-Monitoring, Analysis and Reporting Technology: an open standard for developing drives and software systems that automatically monitors a drive's health and reports potential problems.
SP	Scalable Processor
SSD	Solid State Device – non-volatile memory
UI	User Interface
UEFI	Unified Extensible Firmware Interface
UWD	Universal Windows Driver
VMD	Intel Volume Management Device
VMD Domain	A grouping of drives behind a single Intel VMD
Volume Roaming	Moving of RAID volumes and all its array members between two enabled platforms to maintaining data availability or different supported operating systems.
VROC	Virtual RAID on CPU

§



2 Product Overview

The Intel VROC 6.1 family of products provide enterprise RAID solutions for both NVMe SSD and SATA devices for enterprise servers and workstations. The product family includes the following three products.

1. Intel VROC (VMD NVMe RAID) – This product provides an enterprise RAID solution on platforms that support the Intel VMD technology.
2. Intel VROC (SATA RAID) – This product provides an enterprise RAID solution for SATA devices connected to the SATA/sATA Intel Platform Control Hub (PCH) configured for RAID mode.
3. Intel VROC (NonVMD NVMe RAID) – This product provides an enterprise RAID solution for Intel NVMe SSDs attached to PCIe slots managed by the Platform CPU. Intel VROC (NonVMD NVMe RAID) is not intended for, nor supports:
 - a. Non-Intel NVMe SSDs.
 - b. Platforms that have Intel CPUs that contain Intel VMD technology (weather enabled or disabled).

Note: Intel VROC 6.1 is a high level blanket product reference for Intel VROC (VMD NVMe RAID), Intel VROC (SATA RAID) and Intel VROC (NonVMD NVMe RAID) .

2.1 Product Release Numbering Scheme

The product release **version** is divided into 4 sections or numbers (**AA.BB.CC.DDDD**, e.g. **5.0.0.1001**).

Number / Section	Description
AA: Major Release Number	This section represents the major release version of the product. It is usually associated with a major change in features and/or new platform/chipset launches.
BB: Minor (Maintenance) Release Number	This section represents the minor release version of the product. If this number is non-zero , then the release is a minor release of the AA major release version. This can represent a maintenance release with several bug fixes or it can align with a platform refresh.
CC: Hot Fix Release Number	This section represents customer specific hot fixes. If this number is non-zero , then the release addresses a customer specific hot fix request to resolve their specific issue. These releases are not fully validated and are only intended for a specific customer.
DDDD: Release Build Number	This section represents the build number of release AA.BB.CC , Note: for production releases, the build number generally begins with the number '1' (e.g. AA.BB.CC.1001). In the cases of Intel VROC 6.1 NVMe, this number will begin with '2'



2.2 Supported Operating Systems per Platform

Platform	Windows* 10 (RS2, RS3, RS4, RS5)	Windows* Server 2016 Enterprise	Windows* Server 2019 Enterprise
Intel® C610/C220/C230 series chipset based platforms	Y	Y	Y
Intel® C240 series chipset based platforms	Y	Y	Y
Intel® Xeon® Scalable Processor family based platforms	Y		
Intel® Xeon® Scalable Processor family Workstation based platforms	Y		
Intel® Xeon® Scalable Processor family Server based platforms		Y	Y
Intel® Xeon® Processor D-2100 Product Family based platform		Y	Y

*Only 64bit OS is supported



2.3 Supported Platforms/Chipsets/SKUs

The Intel VROC 6.1 product package was designed to work with, tested and validated on Intel Customer Reference Boards (CRBs) outlined in this section.

2.3.1 Supported Platforms for Intel VROC (VMD NVMe RAID)

CPU	Platform	VMD Device ID	# of VMD
Intel® Xeon® Scalable Processor family – W	Intel® Xeon® Scalable Processor family workstation ⁺	201D	3 per CPU
Intel® Xeon® Scalable Processor family – SP	Intel® Xeon® Scalable Processor family server and workstation ⁺	201D	3 per CPU
Intel® Xeon® Processor D-2100 Product Family	Intel® Xeon® Processor D-2100 Product Family based platform ⁺	201D	3 per CPU
Intel® X299 High End Desktop	Intel® Xeon® Scalable Processor family server and workstation ⁺	201D	3 per CPU

⁺ Unless Otherwise Specified in the Release Notes

Note: Intel VROC support on the Intel X299 High End Desktop platforms is restricted to Intel NVMe SSDs only. This is a platform limitation. As a result, when Intel VROC is installed onto an Intel X299 High End Desktop platform, the customer will only see Intel NVMe SSDs plugged into the platform.

Note: Intel VROC (NonVMD NVMe RAID) is not supported on platforms that support Intel VROC.



2.3.2 Supported Chipset SKU for Intel VROC (SATA RAID)

Chipset	Platform	RAID controller Device ID	Intel VROC (NonVMD NVMe RAID) Virtual Device ID	# of ports
Intel® C610 series chipset	Platforms containing the Intel® C610 series chipset ⁺	2826 (SATA) 2827 (sSATA)	2F9c	6 SATA 4 sSATA
Intel® C610 series chipset	Platforms Refreshes containing the Intel® C610 series chipset ⁺	2826 (SATA) 2827 (sSATA)	6F9C	8 SATA 6 sSATA
Intel® C220 series chipset	Platforms containing the Intel® C220 series chipset ⁺	2826 (SATA)	N/A	6 SATA
Intel® C230 series chipset	Platforms containing the Intel® C230 series chipset Platform ⁺	2826 (SATA)	A135 (Integrated Sensor Hub enabled)	8 SATA
Intel® C240 series chipset	Platforms containing the Intel® C240 series chipset ⁺	2826 (SATA)	A37C (Integrated Sensor Hub enabled)	8 SATA
Intel® 620 series chipset	Platforms containing the Intel® 620 series chipset ⁺	2826 (SATA) 2827 (sSATA)	N/A	8 SATA 6 sSATA
Intel® C422 series chipset	Platforms containing the Intel® C422 series chipset ⁺	2826 (SATA)	N/A	8 SATA

⁺ Unless otherwise specified in the Release Notes

2.4 New Features Introduced in Intel VROC 6.1 Release

2.4.1 Windows* 7-64bit Support

The Intel VROC 6.1 release package includes targeted support for Windows* 7-64bit. Within this package, is included Intel RSTe 5.6 specifically for Windows* 7-64bit. This was added back into the product packaging to support the older platforms.

On platforms that are installing Windows* 7-64bit, the Intel RSTe 5.6 driver from the F6 directory should be used. Once the OS is successfully installed, running the Intel VROC 6.1 installation application will install the Intel RSTe 5.6 GUI and middleware and update/install the driver.

This configuration is the exception to the backwards compatibility of the Intel VROC PreOS. Meaning, that after the platform BIOS has been updated to include the Intel VROC 6.1 PreOS, the Intel RSTe 5.6 package (for Windows 7-64bit) will be supported. Table 1 shows the Intel VROC PreOS versions to Intel RSTe 5.6 OS version support.

Table 1: Intel VROC Compatibility Matrix

	Intel VROC PreOS 5.5	Intel VROC PreOS 6.0	Intel VROC PreOS 6.1
Intel RSTe 5.6 Win7-64bit	Supported	Not Supported	Supported



2.4.2 Microsoft* .NET and Intel ASM No Longer Included

To address functional and security updates, this version of the Intel® Virtual RAID on CPU (Intel® VROC) 6.1 PV Release Package has removed the Microsoft .NET Framework as well as the Intel Acceleration Storage Manager (ASM). Users should update to the latest version.

For the customer's convenience, the Intel VROC product installation application was designed to automatically install the Microsoft .NET Framework and provide an option to install the Intel Acceleration Storage Manager (ASM).

The .NET Framework was included because the Intel VROC user interface relies on the .NET Framework to operate properly. To ensure that customers are able to get the latest version available, Intel is no longer including .NET Framework in the Intel VROC production package. This is not needed because the supported Windows operating systems already include a version of .NET Framework. If the latest version of the .NET Framework is not installed, it can be obtained/downloaded either via a Web update or offline directly from <https://dotnet.microsoft.com/>.

In addition to removal of the .NET Framework installation, this release also removes the Intel ASM component. The Intel ASM installer has some dependencies on 3rd party libraries and Intel would like to reduce or eliminate these dependencies. Until this is accomplished, the Intel ASM component is being removed. Please contact your Intel FAE for future release details.

For more information please refer to Technical Advisory Reference Number 610700.

2.5 New Features Introduced in Intel VROC 6.0 Release

The following table outlines those new features introduced or enhanced with this release of the Intel VROC 6.0 family of products.

New Feature in Intel VROC 6.0	Description
UWD Support	The Intel VROC 6.0 release package includes a GUI that is UWD "H" compliant to support new platforms releasing with Microsoft* Windows* 10 RS5.



2.6 Features Introduced in Intel RSTe 5.5 Release

The following table outlines those new features introduced or enhancements made with the release of the Intel RSTe 5.5 release package.

Feature in Intel RSTe 5.5	Description
UEFI LED Management	The Intel RSTe 5.5 release package added support for LED Management via the PreOS UEFI environment. This includes Locate, Rebuild and Failure LED indicators. This is for both Intel VROC UEFI and Intel VROC (SATA RAID) UEFI PreOS images.
Improved Error and Information Logging	The Intel RSTe 5.5 release package includes improved error logging information in the Windows Event Logs.
Improved Error handling	The Intel RSTe 5.5 release package includes improved error handling when communication errors are encountered.
Support of Older Platforms	The Intel RSTe 5.5 release package includes the merging of the Intel VROC (SATA RAID) 4.X and Intel RSTe 5.X baselines. Intel RSTe 5.5 is now used to support all platforms with the exception of platforms with the Intel C600 or C200 series chipset. Included is the support for the Intel RSTe NVMe product as well.
Changing Controller Default Values	The Intel RSTe 5.5 release package introduced the ability to set/change the controller's default settings for: <ul style="list-style-type: none">• Rebuild on Hot Insert• Read Patrol
Warning Message when Creating RAID Volumes with different types of drives	The Intel RSTe 5.5 release package introduced warning messages when a RAID volume is created when: <ul style="list-style-type: none">• There is a drive mix of SSDs and HDDs• Drive size differences are greater than 10%
Merging 4.X and 5.X Baselines	The Intel RSTe 5.5 release package combines the Intel RSTe 4.X and 5.X baselines into one product baseline. That means that the Intel RSTe 5.5 release package is intended to support the older Intel RSTe supported platforms. Please see the section on Supported Platforms for more information.



2.7 Features Introduced in Intel RSTe 5.4 Release

This section outlines those new features introduced with the release of the document.

Feature in Intel RSTe 5.4	Description
UWD Support for Windows 10 RS3	Intel RSTe 5.4 meets UWD Declarative (D) and Componentized (C) requirements to support the Microsoft* modern Windows* Driver plan
Enhanced LED Management	The LED management feature was enhanced to provide greater flexibility and the ability to customize aspects of the feature. Please see the section titled LED Management for full details.
48 NVMe SSDs	Intel RSTe 5.4 increased the maximum number of NVMe SSDs supported to 48 NVMe SSDs.
Support for additional non-Intel NVMe SSDs	Intel VROC has added non-Intel NVMe SSDs to the support list.



3 Intel VROC 6.1 Package Components

This section outlines the individual components that are included in the Intel VROC 6.1 PV package that is delivered to customers. The specific details on each component are in the following sections of this document. The Intel VROC 6.1 package is comprised of and Intel VROC 6.1 Command Line Interface (CLI) tool, and Intel VROC 6.1 installation utility, Intel VROC 6.1 “F6” drivers and the Intel VROC 6.1 PreOS packages. These package components comprise and support the Intel VROC (VMD NVMe RAID) environment, the Intel VROC (SATA RAID) environment, the Intel VROC (NonVMD NVMe RAID) environment and the Intel ASM installer.

3.1 Intel VROC 6.1 CLI

The VROC 6.1 CLI tool (IntelVROCCli.exe) is a Microsoft® Windows® executable that can be run either from a Microsoft® WinPE environment or from a fully installed Windows operating system. This is a tool that is executed from a command window and can be used as part of a manufacturing scripting environment. This tool supports Intel VROC (VMD NVMe RAID) environment, the Intel VROC (SATA RAID) environment, the Intel VROC (NonVMD NVMe RAID).

Note: The CLI tool is tied to the driver version included in the release package. The CLI tool used needs to match that of the driver version installed on the platform being used.

3.2 Intel VROC 6.1 OS Installer

The Intel VROC 6.1 OS installer (SetupVROC.exe) is a Microsoft® Windows® executable that runs from a fully installed Windows® operating system. The installer will install the latest device drivers, the Intel VROC 6.1 Graphical User Interface (GUI) and supporting components. An option is also available to install the Intel ASM package. This single installer is used to Intel VROC (VMD NVMe RAID), the Intel VROC (SATA RAID), the Intel VROC (NonVMD NVMe RAID) and Intel ASM.

On platforms with Intel CPU's that contain Intel VMD technology (weather enabled or not), the installer will only install Intel VROC (VMD NVMe RAID), the Intel VROC (SATA RAID) and (if desired) Intel ASM products.

On supported platforms with Intel CPU's that do not contain Intel VMD technology, the installer will install Intel VROC (SATA RAID) and/or Intel VROC (NonVMD NVMe RAID).

On platform running Windows® 7-64bit, the Intel RSTe 5.6 package will be installed.

3.3 Intel VROC 6.1 Environment

The Intel VROC 6.1 production package is broken into a Pre-OS component, an OS installer, an F6 (or “Load driver”) component and a Command Line Interface (CLI) component (referenced above).



3.3.1 Intel VROC (VMD NVMe RAID) Pre-OS Components

The Intel VROC Pre-OS components are described in the following table.

VMDVROC_1.efi/ VMDVROC_2.efi	<p>These two UEFI device driver files are to be included in the platform BIOS. Both of these files are required and work together to properly support the Intel VMD controller when enabled. When these files are properly incorporated into the platform BIOS, the BIOS setup will have the ability to manage NVMe drives (create/delete RAID Volumes) connected to the Intel VMD controller. Other features include:</p> <ul style="list-style-type: none"> • Seeing and reporting (to the BIOS) any NVMe SSD attached to the Intel VMD • Installing an OS onto a drive (or RAID Volume) managed by the Intel VMD • Recognizing of Intel VROC RAID Upgrade Hardware Keys • Proper configuration of the system based on the Intel VROC RAID Upgrade Hardware Key seen
HWKeyCheckVROC.efi	<p>This utility can be used by OEMs to test Intel VROC RAID Upgrade Hardware Key configuration via a UEFI Shell environment. This helps to make sure the system can properly see the Intel VROC RAID Upgrade Hardware Key plugged into the system.</p> <ul style="list-style-type: none"> • This tool must match the major release version of the UEFI Intel VROC (VMD NVMe RAID) driver. To find this release version, boot to the efi shell and type "drivers -sfo" and note the major release version which is the first number (capital X). Look for "Intel(R) VROC with VMD Technology AA.BB.CC.DDDD". Compare this value to what is reported in the Customer Release Notes associated with the latest release. Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of Intel VROC (VMD NVMe RAID) UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: HWKeyCheckVROC.efi). • The output from this efi tool can also be saved as a .txt file to attach to an IPS sighting filed against Intel VROC (ex: HWKeyCheckVROC.efi > HWKeyCheckVROC_Output.txt)
RCfgVROC.efi	<p>This utility can be used by the OEMs to manage/test Intel VROC capabilities from a UEFI Shell environment. This utility must be copied to and executed from a USB key.</p> <ul style="list-style-type: none"> • This tool must match the major release version of the UEFI VROC driver. Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of Intel VROC (VMD NVMe RAID) UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: RCfgVROC.efi) • The output from this efi tool can also be saved as a .txt file to attach to an IPS sighting filed against Intel VROC (ex: RCfgVROC.efi /I > RCfgVROC_Output.txt - will save all of the Drive/Volume/Array information in the .txt file)



	Note: When creating a RAID 5 Volume with RWH enabled and when that volume must be usable in either Windows* or Linux*, the RWH feature must be enabled with PPL. Intel VROC for Linux does not support RWH closure using a Journaling Drive.
RCmpVROC.efi	<p>This is a debug utility to help verify that the two efi drivers have been properly configured/incorporated into the BIOS. When reporting an issue to Intel, the output from this file will most likely be asked for by the Intel representative.</p> <ul style="list-style-type: none">• This tool must match the major release version of the UEFI VROC driver. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: RCmpVROCRS.efi).• The output from this efi tool can also be saved as a .txt file to attach to an IPS sighting filed against Intel VROC (ex: RCmpVROC.efi /I > RCmpVROC_Output.txt txt - will save all of the Drive/Volume/Array information in the .txt file
LedToolVMDVROC.efi	<p>This is UEFI base tool that can be used to test LED connectivity to determine if the hardware is setup correctly. The tool provides the ability to simulate a Locate, Fault or Rebuild signal to the enclosure.</p> <ul style="list-style-type: none">• This tool must match the major release version of the UEFI Intel VROC (VMD NVME RAID) driver. Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of the Intel VROC (VMD NVME RAID) UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: LedToolVMDVROC.efi /D1 1 /TIMEOUT 2). This will generate a Locate signal for 2 seconds.• LedToolVMDVROC.efi [/?] [/HELP] [/L] [/Dn state-number] [/TIMEOUT time] /? , /HELP Displays Help Screen. Other options are ignored. /L Lists all detected disks. /Dn Sends the specified state [0-3] to device selected by 'n'. Valid state number is from 0 to 3 and they are interpreted as following: 0 = LED_OFF, 1 = LOCATE, 2 = FAULT, 3 = REBUILD /TIMEOUT After sending the state, waits the specified number of seconds and then turns all LEDs off. Example: "LedtoolVMDVROC.efi /D1 2 /D2 1 /TIMEOUT 2"• This will set state FAULT on device 1 and state LOCATE on device 2. After 2 seconds delay, messages will be reset



3.3.2 Intel VROC F6 (or “Load driver”) Drivers

The Intel VROC 6.1 package includes two sets of individual device drivers for the Intel VMD controller.

The first set is intended to support the manual update or installation of the Intel VROC (VMD NVMe RAID) device driver from within an operational Microsoft® Windows® OS environment. This installation process utilizes the device manager driver update process and is intended to support those configurations that may not have an Intel VROC GUI installed. These device drivers are located under the F6-drivers directory in the release package and are:

- iaVROC.free.win764bit.5.BB.CC.DDDD - for installing Intel RSTe 5.6 (VMD NVMe RAID) on Windows® 7 64bit operating systems.
- iaVROC.free.win8.64bit.6.BB.CC.DDDD - for installing Intel VROC (VMD NVMe RAID) on all other supported Windows® operating systems.

The second set of drivers is intended to support the installation of a supported Microsoft® Windows® OS. Microsoft® Windows® does not contain Intel VROC drivers “inbox”. As a result, installing a supported Microsoft® Windows® OS requires the use of the F6 or “Load driver” option. The Intel VROC F6 Drivers are located in the VROC_6.BB.CC.DDDD_F6-drivers directory, at the same directory level as those previously mentioned.

Note: Microsoft® Windows® OSs will support the use of any of the for mentioned device drivers.

3.4 Intel VROC (SATA RAID) Environment

The Intel VROC (SATA RAID) production package is also broken into a Pre-OS component, an OS installer, an F6 (or “Load driver”) component and a Command Line Interface (CLI) component (referenced above).

3.4.1 Intel VROC (SATA RAID) Pre-OS Components

The VROC (SATA RAID) Pre-OS components are comprised of Legacy Option ROM (OROM) support along with UEFI support. The Intel VROC (SATA RAID) Pre-OS components are:

3.4.1.1 EFI Support for the SATA/sSATA Controller

SataDriver.efi/ sSATADriver.efi	<p>These are the UEFI device drivers to support the SATA/sSATA Controllers when configured for RAID mode. These files are to be included in the platform BIOS. When properly incorporated, it will provide the ability to manage SATA/sSATA drives (create/delete RAID Volumes) connected to the SATA Controller. Other features included:</p> <ul style="list-style-type: none"> • Seeing and reporting (to the BIOS) any SATA drive attached to the SATA/sSATA Controller when configured in RAID mode. • Installing an OS onto a drive (or RAID Volume) managed by the SATA/sSATA Controller when configured in RAID mode.
--	---



LedToolSata.efi/ LedToolsSata.efi	<p>These are UEFI base tools that can be used to test SGPIO LED connections to determine if the hardware is setup correctly. These tools need to be copied to a USB drive and executed from a UEFI Shell.</p> <ul style="list-style-type: none">• The tool must match the major release version of the Intel VROC (SATA RAID) UEFI driver. To find this release version, boot to the efi shell and type "drivers" and note the major release version which is the first number (capital X). Look for "Intel® VROC with VMD Technology X.x.x.xxxx Drive..." Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of VROC 6.1 UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: LedToolSata.efi).• Help for SATA/sSATA LedTool: Intel(R) SGPIO Led Testing Utility for Patsburg AHCI Controller LedTool.efi [/?] [/HELP] [/D:state_number] [/D0-5:state_number] [/TIMEOUT:time] /? Same as /HELP. Displays Help Screen. Other options ignored. /D Send the specified state to all ports. /D0-9 option will be ignored. /D0-5 Send the specified state to selected port. /TIMEOUT After sending the states waits the specified number of seconds and then turns off all LEDs and performs a resets the Messages sending logic. Note: After a reset, the first sent state will turn on all other LEDs. Valid state_number is from 0 to 5 and they are interpreted as following: 0 = NO_DISK, 1 = NO_ACTIVITY, 2 = ACTIVITY, 3 = LOCATE, 4 = FAIL, 5 = REBUILD. Example use: "Ledtool.efi /D2 0 /D0 2 /TIMEOUT 2" this will set state NO_DISK on port 2 and state ACTIVITY on port 0. After 2 seconds delay Messages will be reset.
RCfgSata.efi/ RCfgsSATA.efi	<p>These utilities can be used by the OEMs to manage/test Intel VROC (SATA RAID) capabilities from a UEFI Shell environment. The utilities must be copied to and executed from a USB key in a UEFI Shell.</p> <ul style="list-style-type: none">• The tool must match the major release version of the Intel VROC (SATA RAID) UEFI driver. Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of Intel VROC (SATA RAID) UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type "map -r" to find the FS#. Type FS#. Then run the tool (ex: RCFgSata.efi).• The output from this efi tool can also be saved as a .txt file to attach to an IPS sighting filed against Intel VROC (SATA RAID) (ex: RCFgSata.efi /I > RCFgSata_Output.txt - will save all of the Drive/Volume/Array information in the .txt file)



RCmpSata.efi/ RCmpsSata.efi	<p>These are debug utilities to help verify that the SATA/sSATA efi drivers have been properly configured/incorporated into the BIOS. When reporting an issue to Intel, the output from this file will most likely be asked for by the Intel representative. This utility must be copied to and executed from a USB key in a UEFI Shell.</p> <ul style="list-style-type: none"> The tool must match the major release version of the Intel VROC (SATA RAID) UEFI driver. Using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of Intel VROC (SATA RAID) UEFI driver found in the previous step. Attach the device to the system under test. Navigate to the file system of the fat32 formatted USB disk. Type “map –r” to find the FS#. Type FS#. Then run the tool (ex: RCmpSata.efi). The output from this efi tool can also be saved as a .txt file to attach to an IPS sighting filed against Intel VROC (SATA RAID) (ex: RCmpSata.efi > RCmpSata_Output.txt)
RClrSata.efi/ RClrsSata.efi	<p>This is a utility provided that will clear off any RAID metadata from a SATA drive attached to the SATA controller. This will basically wipe out any data that is on the drive that this tool is executed on. So great care must be taken. This utility must be copied to and executed from a USB key in a UEFI Shell.</p>

3.4.1.2 Legacy OROM Support for the SATA/sSATA Controller

SataOrom.bin/ sSataOrom.bin	<p>These are the Legacy OROM images to support the SATA/sSATA Controllers when configured for RAID mode. These files need to be included in the platform BIOS. When properly incorporated, it will provide the ability to manage SATA drives (create/delete RAID Volumes) connected to the SATA Controller. Other features included:</p> <ul style="list-style-type: none"> Seeing and reporting (to the BIOS) any SATA drive attached to the SATA Controller when configured in RAID mode. Installing an OS onto a drive (or RAID Volume) managed by the SATA Controller when configured in RAID mode.
RCfgSata.exe/ RcfsSata.exe	<p>This utility can be used by the OEMs to manage/test Intel VROC (SATA RAID) capabilities from a DOS Command environment. This utility must be copied to (and executed from) a DOS bootable USB key with the target environment booted from this DOS bootable USB key.</p>
RCmpSata.exe/ RCmpsSata.exe	<p>This is a debug utility to help verify that the SATA Legacy OROM image has been properly configured/incorporated into the BIOS. When reporting an issue to Intel, the output from this file will most likely be asked for by the Intel representative. This utility must be copied to (and executed from) a DOS bootable USB key with the target environment booted from this DOS bootable USB key.</p>



3.4.2 Intel VROC (SATA RAID) F6 (or “Load driver”) Drivers

The Intel VROC (SATA RAID) package includes two sets of individual device drivers for the PCH controller in RAID mode.

The first set is intended to support the manual update or installation of the PCH device driver from within an operational Microsoft® Windows® OS environment. This installation process utilizes the device manager driver update process and is intended to support those configurations that may not have an Intel VROC 6.1 GUI installed. These device drivers are located under the F6-drivers directory in the release package and are:

- iaStorE.free.win7.64bit.5.BB.CC.DDDD - for installing Intel RSTe 5.6 (SATA RAID) on Windows® 7 64bit operating systems
- iaStorE.free.win8.64bit.6.BB.CC.DDDD - for installing Intel VROC (SATA RAID) on all other supported Windows® operating systems.

The second set of drivers is intended to support the installation of a supported Microsoft Windows OS. Microsoft Windows. Although Microsoft Windows does contain support for the SATA controller in RAID mode, Intel recommends against using this driver. There is no inbox driver support for the sSATA controller in RAID mode. The Intel SATA/sSATA F6 Drivers are located in the VROC_6.BB.CC.DDDD_F6-drivers directory, at the same directory level as those previously mentioned.

Note: Microsoft Windows OSs will support the use of any of the for mentioned device drivers.

3.4.2.1 Intel VROC (SATA RAID) F6 (or “Load driver”) Limitation

Starting with Microsoft® Windows® 8, the Microsoft® inbox RAID driver (Rapid Storage Technology or Intel RST) began to include support for the platform Device ID 0x2826, which is used by Intel VROC (SATA RAID). The intent of this functionality is to allow customers who do not have access to the latest F6 drivers and need to re-install their Operating System. The OS installation is to be on a single drive only. Once the OS is installed and the latest Intel VROC (SATA RAID) driver can be installed, the desired RAID volumes can be generated. One of the limitations to the Intel RST “inbox” RAID driver is that it only supports 6 SATA ports. The Intel® Xeon® Scalable Processor family platform has 8 SATA ports. This leads to compatibility issues. For those scenarios where a RAID volume is created in the Intel VROC (SATA RAID) Pre-OS and that RAID volume contains drives on SATA ports 7 and/or 8, installing one of the supported Microsoft® operating systems (except Microsoft® Windows® 7) on to the Intel VROC (SATA RAID) RAID volume may fail. This issue can occur even if the intent is to use the Intel VROC (SATA RAID) F6 driver. This is due to the fact that the Microsoft® installation process will search for a suitable “inbox” driver and if one is found, will install it before the option to load the F6 driver is presented to the user. The “inbox” Intel RST RAID driver will see the RAID volume and also see that one or two drives are missing (those drives that are on ports 7 and/or 8). It will then mark the RAID volume as either “Failed” or “Degraded”. When the Microsoft® installer present the option to “Load Driver” the Intel VROC (SATA RAID) RAID volume has already been marked. The RAID volume status is not visible to the user. When the Intel VROC (SATA RAID) F6 driver is loaded, it will see that the RAID volume is either “Failed” or “Degraded” and will continue on as such. This status of the RAID volume is passed onto the Microsoft® installer and the user will only find out when attempting to continue on with the installation. The following are the two possibilities:

1. For a RAID volume that is marked as “Degraded”, the installation process will continue on as normal. The user may not see the issue until the OS is installed and the Intel VROC (SATA RAID) GUI has also been installed. The status of the RAID volume is presented to the user in the Intel VROC (SATA RAID) GUI. At that point, the user can take the appropriate action.
2. For a RAID volume that is marked as “Failed”, the user will see that the Microsoft® installation will not continue. A message will be presented to the user reporting that the process cannot continue.



There are two workarounds that can be implemented. They are as follows:

1. Intel recommends avoiding using ports 7 and/or 8 as part of the OS installation.
2. Using a USB drive with the Intel VROC (SATA RAID) IntelVROCCLI.exe utility copied to it:
 - a. After the installation process encounters the above mentioned failure (the installation process does not continue, press f10 to invoke a CMD window.
 - b. If you have not already done so, please insert the USB drive into the system. Navigate to your USB drive with the IntelVROCCLI.exe utility.
 - c. Run command: IntelVROCCLI.exe --manage --normal-volume <volumeName>
This will reset the volume to a normal state
 - d. Close the CMD window.
 - e. In the Windows* disk selection window, reload the Intel VROC (SATA RAID) f6 driver.
 - f. Once completed, Windows* should allow installation on the RAID volume.

This limitation does not occur on nor does it apply to the sSATA controller. This is because there is no Microsoft* "inbox" RAID driver for the sSATA controller.

3.5 Intel VROC (NonVMD NVMe RAID) Environment

The Intel VROC (NonVMD NVMe RAID) production package is only includes the F6 drivers to facilitate the installation of an OS using the Intel VROC (NonVMD NVMe RAID) driver.

Note: Intel strongly recommends using the Microsoft Windows NVMe inbox driver to install any of the supported Microsoft Windows operating systems.

3.5.1 Intel VROC (NonVMD NVMe RAID) F6 (or "Load driver") Drivers

There is a set of F6 drivers (iaRNVMe.free.win7.64bit.5.AA.BB.CC.DDDD) for installing Windows* 7 operating systems. There is also a set a F6 drivers (iaRNVMe.free.win8.64bit.5.AA.BB.CC.DDDD) for installing all other supported Windows* operating systems.

- iaRNVMe.free.win7.64bit.5.BB.CC.DDDD --for installing Intel RSTe 5.6 (NonVMD NVMe RAID) on Windows* 7 64bit operating systems
- iaRNVMe.free.win7.64bit.5.BB.CC.DDDD - for installing Intel VROC (NonVMD NVMe RAID) on all other supported Windows* operating systems.

e included with each Kit.



4 Intel® VROC Key Features

The Intel VROC 6.1 product package is comprised of several components that provide a complete platform solution.

The following is a brief list of the features that are supported on Intel VROC (VMD NVMe RAID).

Features	<ul style="list-style-type: none">• Surprise Hot-Plug• LED Management (VMD Method)• Error Management (VMD First)• NVMe RAID Boot• Enhanced GUI• UEFI HII• RAID Volume Roaming• EFI Public UEFI NVMe Pass-Through Command Support• Windows* Driver Pass-through IOCTL Support• Intel VROC Private UEFI Device Information Protocol• Intel VROC Private UEFI Volume Information Protocol	<ul style="list-style-type: none">• 4K native drives• NVMe Deallocate/TRIM• EFI Private protocol to extract device information on non-RAID NVMe drives• Intel VROC 90 Day Trial Period• Intel VROC Pass-Thru Mode With RAID 0 with Intel P3608• RAID 0/1/10 with Intel VROC Standard Upgrade Key• RAID 0/1/5/10 with Intel VROC Premium Upgrade Key• RAID 0/1/5/10 for Intel NVMe SSDs with Intel VROC Intel-SSD-Only Upgrade Key (Non-Intel NVMe SSDs supported in Pass-thru)• RAID Write Hole Closure (RAID 5 w/ Intel VROC Premium Upgrade Key)
Utilities	<ul style="list-style-type: none">• Install/Uninstall Utility• Configuration and Management Utilities	

4.1 Intel® Volume Management Device (Intel® VMD)

With the introduction of the Intel® Xeon® Scalable Processor family, one of the key features included is the Intel VMD. The Intel VMD is an integrated PCIe endpoint within the CPU root complex. The class code for the Intel VMD device is a RAID controller. Intel VMD driver support is provided with Intel VROC package and includes the following:

- Multiple Intel VMD Controllers
- LED Management (VMD Method of LED Management)
- Surprise Hot Plug
- Error Handling

Each Intel® Xeon® Scalable Processor family CPU provides up to 48 PCIe lanes which is subdivided into 3 domains. Each VMD domain manages x16 lanes. Intel VMD can be turned on/off on x4 lane granularity and supports either NVMe SSD device or PCIe switch device.

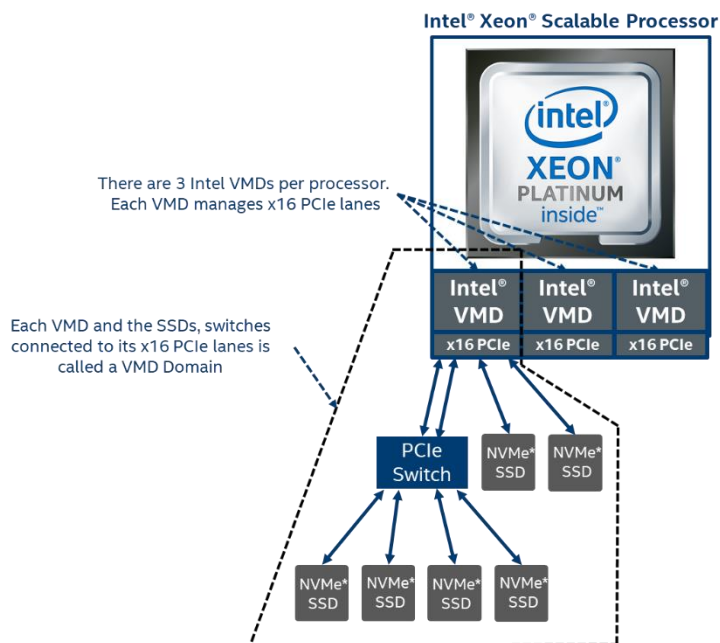
Intel VMD is also designed to support the following types of NVMe Solid State Drives:

- PCI express Add in Cards (AIC)
- M.2 form factor drives
- U.2 form factor drives



4.1.1 Multiple Intel VMD Controllers

Each Intel® Xeon® Scalable Processor family CPU contains 3 Intel VMD controllers.



4.1.2 Intel VMD Method of LED Management

Backplane LED management is a critical part of enterprise RAID management. Intel VMD provides a mechanism to support LED management on Intel® Xeon® Scalable Processor family platforms. The mechanism utilizes repurposed two of the PCIe Slot Control Registers (Power Indicator Control and Attention Indicator Control of the Slot Control register) to support IBPI blink patterns defined in the SFF-8489 standard.

For details on how to implement the Intel VMD Method, refer to the *LED Management for PCI Express SSDs with Intel RAID* document (IBL Doc ID: 334005-004US).

For details on how to utilize and/or customize LED management behavior in this product, please see the section on [LED Management](#) in this document.

Note: LED management outside of a supported backplane environment is not supported.

4.1.3 Intel VMD Error Management

Intel VMD technology handles Advanced Error Reporting (AER) Handling and functions as a crash dump and hibernate target within Windows® OS, therefore providing protection to the rest of the PCIe bus, preventing whole system shut down. Upon receiving an interrupt at MSI-X 0 for each root port, the Intel® VMD uses a sophisticated algorithm to determine the precise device to apply error handling based severity. The actions taken are specific to the error on the device, and preserve the remaining eco system.



Intel VMD controller will handle the following errors. If non-fatal error is encountered, the Intel VROC (VMD NVMe RAID) driver will handle the error and standard recovery efforts are taken. If a fatal error is encountered, the Intel VROC (VMD NVMe RAID) driver will treat the error as a “device surprise hot removed”. But system should continue to operate without hang, reboot, or crash.

These messages will be seen in the System Event Viewer environment. The Event ID that will signify one of these conditions has occurred is 11. Please gather this information and provide it to your Intel VROC representative for analysis.

Advanced Uncorrectable Error	Severity Register
Data Link Protocol Error	Fatal (1b).
Surprise Down Error	Fatal (1b).
Flow Control Protocol Error	Fatal (1b).
Receiver Overflow	Fatal (1b).
Malformed TLP	Fatal (1b).
ECRC Error	Non-Fatal (0b).
Unsupported Request Error	Non-Fatal (0b).
Poisoned TLP	Non-Fatal (0b).
Completion Timeout	Non-Fatal (0b).
Completer Abort	Non-Fatal (0b).
Unexpected Completion	Non-Fatal (0b).
ECRC Error	Non-Fatal (0b).
Unsupported Request Error	Non-Fatal (0b).
ACS Violation	Non-Fatal (0b).
MC Blocked TLP	Non-Fatal (0b).
Atomic Op Egress Blocked	Non-Fatal (0b).
TLP Prefix Blocked	Non-Fatal (0b).

4.1.4 Intel VMD Enabling

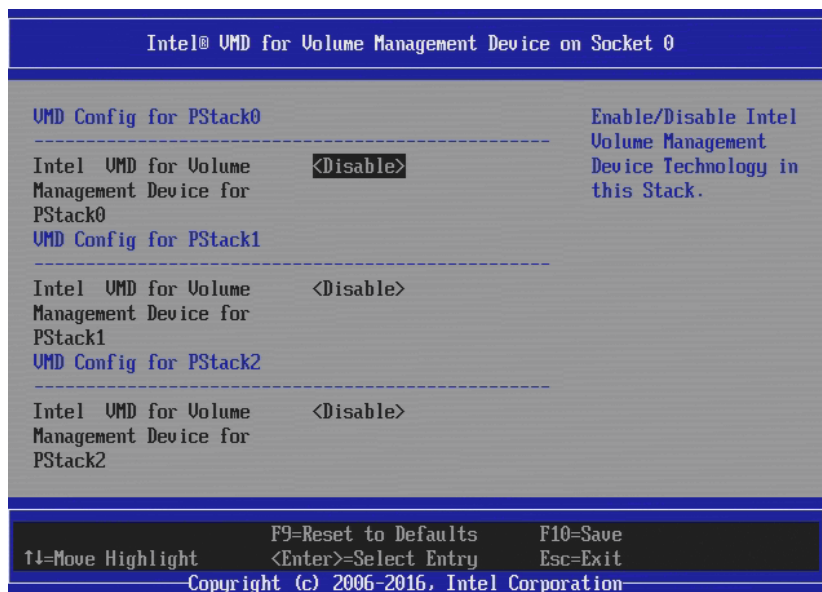
These instructions will cover how to enable Intel VMD in an Intel Customer Reference Board (CRB).

Note: Refer to the instructions provided by the user's platform BIOS vendor because those instructions will most likely be different from these instructions.

1. Immediately following POST, select the option that will allow the user to access the BIOS setup menu. This example uses F2.
2. For the Intel CRB reference BIOS, The user will want to use the arrow keys to move the cursor to **the EDKII Menu** (it will become highlighted) and press <Enter>.
3. Using the arrow keys, move the cursor to **Socket Configuration** and press <Enter>.
4. Using the arrow keys, move the cursor to **Intel® VMD technology** and press <Enter>.
5. Using the arrow keys, move the cursor to **Intel® VMD for Volume Management Device on SocketX** (where X represents the number of the socket the user is to modify).



6. Using the arrow keys, move the cursor to select the desired PStack# to **enable** or **disable** Intel VMD according which hardware configuration is being used. See below:



7. Repeat Step 6 for each PStackX to be **enabled** or **disabled**.

See the user's Platform BIOS manufacturer documentation for a complete list of options that can be configured.

4.1.5 Intel VMD Surprise Hot-Plug of NVMe Devices

On Intel® Xeon® Scalable Processor family platforms with supported Hot Plug backplanes, when Intel VMD is enabled surprise insertion and removal of NVMe drives is expected to work just like SATA drives surprise insertion and removal. Drives will automatically appear or disappear without requiring manual scan in device manager.

Note: Surprise hot-plug is only supported for U.2 form factor NVMe drives because other form factor (AIC, M.2) drives themselves doesn't support surprise hot-plug.

When Intel VMD is enabled, surprise insertion and removal of NVMe drives is isolated from the PCIe bus and managed by Intel VMD. The Intel VROC (VMD NVMe RAID) driver takes appropriate action on NVMe SSD device changes based on SFF-8639 interface (U.2 Form Factor only with this release). This includes surprise insertion/removal of 2.5" form-factor devices in the OS runtime environment.



4.2 Intel VROC (VMD NVMe RAID) Pre-OS Features and Functionality

With the introduction of Intel VMD, and with that, Intel VROC (VMD NVMe RAID), one of the key features is the ability to install an OS to and boot from a RAID volume. A key component in this feature are the Intel VROC (VMD NVMe RAID) UEFI drivers. The Intel VROC (VMD NVMe RAID) UEFI drivers are a set of binary images that are compiled into the platform BIOS and provide a method by which the BIOS environment will be able to do the following:

1. Initialize and enumerate NVMe devices managed by the Intel VMD domains (when enabled) and report that information to the BIOS.
2. Scan for the Intel VROC RAID Upgrade Key and properly configure the Intel VROC (VMD NVMe RAID) software to properly manage the SKU the system is configured for.
3. Provide an Intel VROC (VMD NVMe RAID) management user interface to the BIOS to manage RAID volumes NVMe devices connected to the Intel VMD controller.
4. Interface to and boot from NVMe devices managed by the Intel VMD.
5. Interface to and boot from Intel VROC (VMD NVMe RAID) volumes.

4.2.1 Intel® VROC (VMD NVMe RAID) for UEFI

For information on how to incorporate the Intel VROC (VMD NVMe RAID) UEFI images into the platform BIOS, please see section [Intel® VROC 6.1 Unified Extensible Firmware Interface \(UEFI\) Drivers](#).

Note: There is no Legacy Option ROM (OROM) support for Intel VROC.

4.2.2 Intel VROC (VMD NVMe RAID) UEFI Human Interface Infrastructure (HII)

The Intel VROC 6.1 product supports RAID management via the Intel VROC UEFI HII. The Intel VROC HII is part of the UEFI Protocol and provides a way for customers to manage RAID Volumes behind the Intel VMD controller in the BIOS environment.

With Intel VROC 6.1, the Intel VROC (VMD NVMe RAID) HII will report the following for each device discovered:

- The Slot Number to which the NVMe drive is connected
- Bus/Device/Function information to which the NVMe drive is connected
- The CPU of the Intel VMD that the NVMe drive is connected
- The Intel VMD Controller Number to which the NVMe drive is connected
- The Root Port Offset to which the NVMe drive is connected
- The NVMe drive vendor identification
- The NVMe drive serial number
- The NVMe drive size

Note: Intel VROC does not support nor provide a Legacy Option ROM image as part of the Intel VROC Pre-OS package.



4.2.3 Intel VROC (VMD NVMe RAID) UEFI Maximum Number of Drives

The Intel VROC (VMD NVMe RAID) UEFI driver supports a maximum number of 48 NVMe drives across all Intel VMD controllers within the supported platform.

4.2.4 Using the Intel VROC (VMD NVMe RAID) HII to Create a RAID Volume

The following are instructions on how to create a RAID volume using the Intel VROC HII interface. This procedure should only be used for a newly-built system or a reinstall of the operating system. The user should use the Intel VROC (SATA RAID) GUI in the Windows* operating system for creating RAID volumes after the operating system is up and running.

Note: Please consult the user's platform documentation for instructions on how to enter into the Intel VROC (VMD NVMe RAID) HII interface.

The following assumptions have been made:

- It is known how to enter into the appropriate platform BIOS level setup menus
- The Intel VMD functionality has been enabled
- The appropriate Intel VROC RAID Upgrade Key has been installed
- The appropriate number of NVMe SSDs have been plugged into the enabled Intel VMD controller.

To create a RAID volume, do the following:

- Enter into BIOS configuration setup menu to access the Intel VROC (VMD NVMe RAID) HII user interface.
- Navigate to and select "Intel® Virtual RAID on CPU"
- Navigate to and select "Create RAID Volume"
- Type in a volume name and press the <Enter> key, or press the <Enter> key to accept the default name.
- Select the RAID level by pressing <Enter> to access the Menu and using the arrow keys to scroll through the available values, then press the <Enter> key to select the desired RAID type.
 - This step is optional, and applies to RAIDs that have volumes connected to more than one Intel VMD controller. To enable the RAID to be spanned over multiple controllers, highlight the < >, and press <Enter>.

This will open a small menu where the default is set as blank, indicating that this feature is not enabled. To enable, use the down arrow key, and press <Enter>.

- Using the arrow keys, highlight the drives one by one by selecting the < > bracket on the line next to that drive's port number. Press <Enter> to open the selection menu which will be set initially to blank. Use the arrow key to highlight the X and press <Enter> to enable as part of the RAID.
- Repeat previous step for each drive required in this RAID.
- Unless the user has selected RAID 1, select the strip size by using the arrow keys to scroll through the available values, then press the <Enter> key.
- Select the volume capacity and press the <Enter> key. The default value indicates the maximum volume capacity using the selected disks. This value is calculated in bytes. For demonstration 700GB is 716800. (700 X 1024 = 716800)



- Use the down arrow key to select Create Volume and press <Enter>.

Note: At this point in the process if there are any significant discrepancies between the drives selected to be RAID members, a warning message will be displayed if one of the following conditions is encountered:

- There is a combination of SSDs and HDDs used
 - There are at least two drives that have a size difference of more than 10%.
- The user will be presented with a screen reporting on the status of the Create process.

If the user's RAID created successfully, it will be listed under the heading of RAID Volumes.

Other disks or portions of volumes that were not included within the RAID will be listed as part of the selections under Non-RAID Physical Disks. These may be used to create additional RAID volumes.

- To exit the user interface, press <Esc>. Press <Esc>, and the user will be presented with the following message: Changes have not saved. Save Changes and exit? Press 'Y' to save and exit, 'N' to discard and exit. 'ESC' to cancel. Press Y to save and confirm (if necessary) to exit.

The user's RAID configuration has now been saved.

- To save and reboot to begin OS installation, press <Esc> to return to the Main Menu. Highlight Reset, and press <Enter> to reboot the system back to the boot menu. The user may be required to accept the changes made to the BIOS setup before exiting. Also, the steps to exit out of the BIOS setup may differ depending on the platform BIOS.

4.2.5 NVM Express Pass-Thru Protocol Support

The Intel VROC (VMD NVMe RAID) UEFI driver will provide support for the following public UEFI NVMe Pass through protocol (EFI_NVMM_EXPRESS_PASS_THRU_PROTOCOL) commands, as defined by the UEFI public specification:

- **Identify** - The Identify command returns a data buffer that describes the NVMe controller. To read identification parameters from NVMe device applications will send identify command. Intel® VROC (VMD NVMe RAID) will pass the ADMIN_IDENTIFY command to appropriate NVMe device.
- **Set Features** - The Set Features command specifies the attributes of the Feature indicated. Intel® VROC (VMD NVMe RAID) will pass the ADMIN_SET_FEATURES command. Intel® VROC (VMD NVMe RAID) passes this command to NVMe device in pass through mode and RAID mode.
- **Get Features** - The Get Features command retrieves the attributes of the Feature specified. Intel® VROC (VMD NVMe RAID) will pass the ADMIN_GET_FEATURES command. Intel® VROC (VMD NVMe RAID) passes this command to NVMe device in pass through mode and RAID mode.
- **Get Log Page** - To read values of SMART and general health information the application sends ADMIN_GET_LOG_PAGE command. Intel® VROC (VMD NVMe RAID) will pass the ADMIN_GET_LOG_PAGE command to appropriate NVMe device. Self-Monitoring, Analysis and Reporting Technology (S.M.A.R.T.) is an open standard used by hard-drives and hosts to monitor drive health and report potential problems. The SMART and health information are collected over the life of the NVMe controller and is retained across power cycles. By default, SMART monitoring is always enabled on NVMe products.
- **Format NVM** - The Format NVM command is used to low level format the NVM media. This is used when the host wants to change the LBA data size and/or metadata size. Intel® VROC (VMD NVMe RAID) passes this command to NVMe device in pass through mode and RAID mode.



- **Firmware Commit/Activate** - Intel® VROC (VMD NVME RAID) will pass the ADMIN_FIRMWARE_ACTIVATE command. The Firmware Activate command is used to verify that a valid firmware image has been downloaded and to commit that revision to a specific firmware slot. The host may select the firmware image to activate on the next controller reset as part of this command.
- **Firmware Download/Update** - Intel® VROC (VMD NVME RAID) will pass the ADMIN_FIRMWARE_IMAGE_DOWNLOAD command to appropriate NVMe device. The ADMIN_FIRMWARE_IMAGE_DOWNLOAD command is used to download firmware image to the controller. The new firmware image will not start to run right after ADMIN_FIRMWARE_IMAGE_DOWNLOAD command. To select which firmware version will be executed after NVMe device reset ADMIN_FIRMWARE_ACTIVATE command must be used.
- **Device Self-Test (DST)** - The Device Self-test command is used to start the device self-test operation or abort a device self-test operation. Intel® VROC (VMD NVME RAID) will pass the ADMIN_DEVICE_SELF_TEST command. Intel® VROC (VMD NVME RAID) passes this command to NVMe device in pass through mode and RAID mode.

4.2.6 Intel VROC (VMD NVME RAID) IOCTL Support

The Intel VROC (VMD NVME RAID) package supports passing certain NVMe admin commands to a particular NVMe device attached to an Intel VMD domain. The Intel VROC (VMD NVME RAID) IOCTLs provide support for obtaining information about the disks in the system and to be able to use that information to send NVMe private IOCTLs to a specific disk.

For specific information about the Intel VROC (VMD NVME RAID) IOCTL support, please see the 1.4 version Intel VROC (VMD NVME RAID) IOCTL document included in this package.

4.2.7 Intel VROC (VMD NVME RAID) Private Device/Member Info Protocol

The Intel VROC (VMD NVME RAID) package will provide a customer or custom interface to provide information on NVMe devices in pass-thru mode (non-RAID) as well as those NVMe devices that are part of an Intel VROC (VMD NVME RAID) RAID volume when those devices are managed by Intel VMD. This interface (or protocol) is installed on each handle of the Intel VMD managed pass-thru NVMe drives. It is important that the consumer, utilizing this information, not alter any of the information directly as this could lead to unpredictable behavior. If the information is to be retained, a copy of the information should be made and the protocol should be closed. This is because the underlying information can change over time. The protocol information can be found in Intel® Virtual RAID on CPU (Intel® VROC (VMD NVME RAID)) Private UEFI Device Info Protocol which is included in the Kit.



4.2.7.1 Reference Documents

Document	Description	Document Number/Location
Intel® Virtual RAID on CPU (Intel® VROC (VMD NVME RAID)) Private UEFI Device Info Protocol Intel®_VROC (VMD NVME RAID)_UEFI_DEVICE_INFO_PROTOCOL_1_5	This Intel VROC (VMD NVME RAID) protocol can be used to retrieve information for non-RAID and RAID member NVMe devices connected to Intel VMD domains.	Included with the Kit
Intel® Virtual RAID on CPU (Intel® VROC (VMD NVME RAID)) Private UEFI Volume Info Protocol Intel®_VROC (VMD NVME RAID)_UEFI_VOLUME_INFO_PROTOCOL_REV_1_0	Intel VROC (VMD NVME RAID) volume info protocol can be used to retrieve information for all RAID volumes created thru Intel Virtual RAID on CPU (Intel® VROC) UEFI mode only, and Intel VROC (VMD NVME RAID). Changes can also be seen using the HII and Intel VROC (VMD NVME RAID) / Intel VROC (VMD NVME RAID) UEFI tools.	Included with the Kit
Intel® Virtual RAID on CPU (Intel® VROC (VMD NVME RAID)) IOCTLs Intel®VROC_IOCTLs_1.04	This document presents all NVMe admin commands supported by Intel® Virtual RAID on CPU (Intel® VROC (VMD NVME RAID)) and describes how to send them to a particular NVMe device.	Included with the Kit

4.3 Intel VROC 6.1 Monitoring and Management

This section is intended to outline those specific features in the user interfaces that are specific to Intel VROC.

4.3.1 Intel VROC 6.1 Graphical User Interface (GUI)

The Intel VROC 6.1 GUI will present all VMD domains in the system under a single header marked “Intel(R) VROC”. The individual NVMe SSDs will be displayed and associated with the Intel VMD domain controller it is attached to. For example “NVMe SSD on Controller 0” would be an NVMe SSD plugged into Intel VMD 0.

For further information on the Intel VROC 6.1 GUI, see section 5.1, Intel VROC 6.1 RAID Features and Functionality.

4.4 Intel VROC Features and Functionality

This section is intended to outline those features that are specific to Intel VROC 6.1.

4.4.1 Intel VROC RAID Modes of Operation Upgrade Hardware Key for RAID Support

The Intel VROC 6.1 package will support four operational SKUs (or modes) of the platform. They are the Intel VROC Pass-thru SKU, Intel VROC Standard SKU, Intel VROC Premium SKU and Intel VROC Intel-SSD-Only SKU.



SKU	Description	HW Key Required	Key Features
Intel VROC Pass-thru	<p>The expected default platform configuration will be the Intel VROC Pass-thru SKU. This SKU is a platform that supports Intel VROC (VMD NVME RAID) but does not contain an Intel VROC RAID Upgrade Key. In this configuration, the Intel VROC 6.1 product will enumerate all NVMe drives (including non-Intel drives) and will expose them to the platform as pass through drives. Intel VROC (VMD NVMe RAID) functionality is disabled with the exception of support for RAID 0 on any x8 (XX08) Intel SSD as long as it has launched and is still active. E.g. P3608/P3708/P4608/etc.</p> <p>Note: In this mode, Intel VROC (VMD NVME RAID) does not block, or restrict the creation of RAID 0 volumes to only Intel P3608/4608/etc. NVMe SSDs. User will be able to create RAID 0 volumes on other NVMe SSD managed by the Intel VMD. However, issues encountered with non XX08 drives are not supported.</p> <p>The message displayed in this mode is: Intel® VROC (VMD NVME RAID) (in pass-thru mode)</p>	Not needed	<ul style="list-style-type: none"> All Intel NVMe SSDs supported in Pass-thru mode only (no RAID) <ul style="list-style-type: none"> RAID 0 management only supported on any x8 (XX08) Intel SSD as long as it has launched and is still active Non-Intel NVMe SSDs supported in Pass-thru mode only (no RAID) (See NVMe Drives for list of supported drives) LED Management Hot Plug Support
Intel VROC Standard	<p>The Intel VROC Standard SKU configuration is enabled when the Intel VROC RAID Standard Upgrade Key is installed in the platform. This SKU has the same behavior as the Intel VROC Pass-thru SKU with the addition of support for RAID 0, 1 and 10. When the Intel VROC RAID Standard Upgrade Key is installed, and Intel VMD is enabled, the Intel VROC (VMD NVME RAID) UEFI HII will enable RAID 0, 1 and 10 management on all supported NVMe drives. This functionality applies only to those PCIe slots that have Intel VMD enabled.</p> <p>The message displayed in this mode is: Intel® VROC (Standard)</p>	Standard Key	<ul style="list-style-type: none"> Pass-thru SKU features RAID 0, 1, 10 (Bootable and Data)
Intel VROC Premium	<p>The Intel VROC Premium SKU configuration is enabled when the Intel VROC RAID Premium Upgrade Key is installed in the platform. This SKU has the same behavior as the Intel VROC Standard SKU with the addition of support for RAID 5 and RAID 5 Write Hole Closure. When the Intel VROC RAID Premium Upgrade Key is installed, and Intel VMD is enabled, the Intel VROC (VMD NVME RAID) UEFI HII will enable RAID 0, 1, 5 and 10 management on all supported NVMe drives. This functionality applies only to those PCIe slots that have Intel VMD enabled.</p> <p>The message displayed in this mode is: Intel® VROC (Premium)</p>	Premium Key	<ul style="list-style-type: none"> Standard SKU features RAID 5 (Bootable and Data) RAID 5 Write Hole Closure



Intel VROC Intel-SSD-Only	<p>The Intel VROC Intel-SSD-Only SKU configuration is enabled when the Intel VROC RAID Intel-SSD-Only Upgrade Key is installed in the platform. This SKU has the same behavior as the Intel VROC Premium SKU but will only support only Intel NVMe SSDs. When the Intel VROC RAID Intel-SSD-Only Upgrade Key is installed, and Intel VMD is enabled, the Intel VROC (VMD NVMe RAID) UEFI HII will enable RAID 0, 1, 5 and 10 management only on Intel NVMe drives. This functionality applies only to those PCIe slots that have Intel VMD enabled. Non-Intel NVMe SSDs (on the supported list in NVMe Drives) will be supported in pass-thru mode only.</p> <p>The message displayed in this mode is: Intel® VROC (in pass-thru mode)</p>	Intel-SSD-Only Key	<ul style="list-style-type: none">• Full Intel Premium Key feature Support only for Intel NVMe SSDs• Non-Intel NVMe SSDs supported in Pass-thru mode only (no RAID) (See NVMe Drives for list of supported drives)
---------------------------	--	--------------------	---

Please consult the user's platform guides to identify which PCIe slots are managed by the CPU.

To enable either the Intel VROC Standard SKU, the Intel VROC Premium SKU or Intel VROC Intel-SSD-Only SKU, an Intel VROC RAID Upgrade Key must be purchased and installed in the platform.

This Intel VROC RAID Upgrade Hardware Keys are available for purchase for use on platforms based on Intel® Xeon® Scalable Processor family platforms with SKX microarchitecture. The Intel VROC RAID Upgrade Hardware Key is a small PCB board that has a security EEPROM that, when read by the Intel VROC (VMD NVMe RAID) UEFI driver, can enable different Intel VROC (VMD NVMe RAID) software stack features when Intel VMD is enabled. Please refer to the user's platform documentation for the location of the Intel VROC RAID Upgrade Hardware Key placement.

Note: It is assumed that all Intel VROC (VMD NVMe RAID) supported platforms contain support (a physical hardware header) to support the Intel VROC RAID Upgrade Key. Please reference Activation Key Product Architecture Specification Revision 1.4 (or later) for additional information.

4.4.2 Intel VROC (VMD NVMe RAID) Volumes on a Single Intel VMD Domain

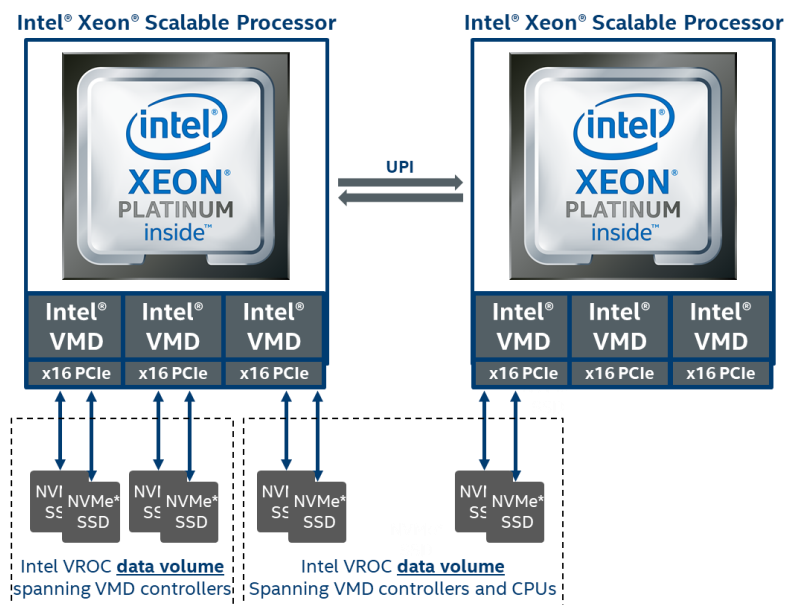
Intel VROC (VMD NVMe RAID) allows both DATA RAID and Boot System RAID volumes on NVMe disks connected to a single Intel VMD domain. These RAID volumes are reported to the OS as single disk drives connected to the Intel VMD domain where the RAID volume member disks are connected.

4.4.3 Spanning Intel VROC (VMD NVMe RAID) RAID Data Volumes across Intel VMD Domains

Intel VROC (VMD NVMe RAID) allows DATA RAID volumes on NVMe SSDs that span across multiple Intel VMD domains. DATA RAID volume members can reside on different Intel VMD domains and can be used for storing data.

Note: Although supported, spanning across CPUs is generally not recommended. This configuration may incur performance penalties.

Note: Intel VROC (VMD NVMe RAID) does not support installing an OS onto an Intel VROC (VMD NVMe RAID) volume that spans across Intel VMD controllers/domains.



4.4.3.1 Intel VROC (VMD NVMe RAID) on Pass-through Boot Drives

The Intel VROC 6.1 will allow installing to and booting from an OS on a pass-through NVMe SSD drive (not part of a RAID Volume) that is behind an enabled Intel VMD controller.

4.4.3.2 Intel® VROC (VMD NVMe RAID) Spanned RAID Volumes across Intel VMD controllers

Intel VROC (VMD NVMe RAID) supports the ability to create an Intel VROC (VMD NVMe RAID) Volume that will span Intel VMD domains. These volumes can be created at any point before or after the user's system is successfully running Windows, but may not be used as the boot volume. The following instructions will cover creation of spanned volumes using the BIOS interface.

The following assumptions will be made about the configuration and setup for configuration of RAID volumes that are spanned across Intel VMD controllers:

- It is known how to access the Configuration BIOS menus.
- All pertinent vendor documentation has been reviewed to ensure that the hardware is configured to allow for more than one VMD domain has been established.
- The Intel VMD has been properly enabled within the BIOS.
- The appropriate Intel VROC RAID Upgrade Standard/Premium/Intel-SSD-Only Key has been installed.
- The appropriate number of drives have been installed to create the desired RAID volume.

To create an Intel VROC (VMD NVMe RAID) Volume, do the following:

1. Enter into the BIOS configuration the setup menu to access the Intel VROC (VMD NVMe RAID) HII user interface.
2. Navigate to and select "Intel® Virtual RAID on CPU".



3. Navigate to and select "Create RAID Volume"
4. Type in a volume name and press the <Enter> key, or press the <Enter> Key to accept the default name.
5. Select the RAID level by pressing the <Enter> to access the Menu and scroll through the available values, then press <Enter> key to select the desired RAID type.
6. Ensure that the user's vendor supports spanned RAID volumes. To enable, toggle the status from the default status off, to on.

To enable the RAID to be spanned over multiple controllers, highlight the < >, and press <Enter>. This will open a small menu box with two values, blank and X. Blank indicates a setting that is disabled. To enable RAID spanned over Intel VMD controllers, highlight X and press the <Enter> key.

7. Using the arrow keys to highlight the drives one by one by selecting the < > bracket on the line next to that drive's port number. Press <Enter> to open the selection menu which will be set initially to blank, use the down arrow key to highlight the X and press <Enter> to enable as part of the RAID.
8. Repeat Step 7 for each drive required in this RAID.

Unless the user has selected RAID 1, select the strip size by selecting the current value and pressing the <Enter> key. This will allow the user to scroll through the available values, then press the <Enter> key to set the desired value. RAID 1 will by default place this setting at 128k and cannot be modified.

9. Select the volume capacity and press the <Enter> key.

The default value indicates the maximum volume capacity using the selected disks and calculating that at a total of 95% of the value due to disk coercion. To enter a value smaller than the maximum based on RAID type selected, press the <Enter> key. Type in the value of the desired RAID volume, and press <Enter> to save. This value is calculated in bytes. For example, 700GB is 716800. ($700 \times 1024 = 716800$). The user may also perform the math based on the base of 1000 instead of base 1024. This is left to the discretion of the administrator.

10. Navigate to and select Create Volume, then press <Enter>.
11. The following message will be displayed: "You have selected NVMe drives that are connected to multiple Intel VMD controllers. Please note that if the user continues and creates a RAID volume with drives from multiple Intel VMD controllers that RAID volume will not be bootable in a Windows* OS environment. Press 'y' to create, 'n' to discard".
12. Press <Enter> to create the RAID volume.

Note: The message is a warning regarding Windows. This is a BIOS configuration platform and can be compatible with either Windows* Server or Linux distributions. However, spanned boot volumes in Windows* are not supported.

Note: At this point in the process if there are any significant discrepancies between the drives selected to be RAID members, a warning message will be displayed if one of the following conditions is encountered:

- There is a combination of SSDs and HDDs used
- There are at least two drives that have a size difference of more than 10%.

13. The RAID volume configuration can be verified by looking at the RAID Volume Information screen. If the RAID volume created successfully, it will be listed under the heading of RAID Volumes.



Other disks or portions of volumes that were not included within the RAID will be listed as part of the selections under Non-RAID Physical Disks. These may be used to create additional RAID volumes. If there are insufficient disks or space on the other disks to create a RAID volume, the user may receive an error when creating secondary volumes.

14. To exit the user interface, press <Esc>. Press <Esc> again, and the user will be returned to the main menu. Navigate to and select Reset, and press <Enter> to reboot the system and begin the Operating System installation.

4.4.4 Intel VROC (VMD NVME RAID) Driver Upgrade or Downgrade

If the Intel VROC (VMD NVME RAID) driver is downgraded to a previous version, the Intel VROC (VMD NVMe RAID) volumes will be maintained (not broken). When the Intel VROC (VMD NVME RAID) driver is upgraded to the latest release, any of the Intel VROC (VMD NVMe RAID) volumes will remain fully operational.

Note: Any exceptions to this will be clearly outlined in the customer release notes included in the release package.

4.4.5 Intel VROC (VMD NVME RAID) Support Maximum number of Drives in a RAID Volume

The Intel VROC 6.1 will support the ability to create RAID volumes from 2 up to 48 NVMe drives connected to the Intel VMD controllers (either directly connected, connected via NVMe retimer cards, connected via NVMe switches or a combination of the three).

4.4.6 Intel VROC (VMD NVME RAID) Support Maximum number of Switch Levels

The Intel VROC 6.1 will support the ability to cascade up to two levels of NVMe switches on a single Intel VMD domain.

4.4.7 Intel VROC (VMD NVME RAID) Support for Maximum Number of Arrays

The Intel VROC 6.1 product will support the ability to create a maximum of 24 Arrays on NVMe drives connected to the Intel VMD controllers. An Array can be created on a minimum of two drives.

4.4.8 Intel VROC (VMD NVME RAID) Support for Maximum RAID Members

The Intel VROC 6.1 will support the ability to create 2 RAID volumes per Array on NVMe drives connected to the Intel VMD controllers.

4.4.9 NVMe Hot Insert

Intel VROC (VMD NVME RAID) has the following limitation when hot inserting an NVMe drive:

- When hot inserting multiple drives, allow enough wait time between each disk for the platform to process each event
- When inserting an entire RAID volume, it is recommended that the system be shut off, the disks inserted and then the system powered back on. If an attempt is made to hot insert all of the drives of the RAID volume, the RAID volume may enter into a Fail state because of the above mentioned delay. After all of the drives are inserted, if the RAID volume remains in a failed state, an attempt to recover the RAID volume by use the Intel VROC 6.1 GUI to reset the volume to normal. This may allow the volume to either continue to operate, or force a volume rebuild. There is a risk that the RAID volume could become inoperable.



4.4.10 Intel NVMe Wear Leveling Recommendations

NVMe Wear Leveling refers to techniques used to prolong the service life of NVMe drives. This section will outline the recommended configurations (number of drives vs strip size) to maximize Wear Leveling on Intel NVMe SSDs when configured as part of RAID 5 volume. When creating an Intel VROC (VMD NVME RAID) RAID 5 volume, several configuration parameters can be selected, and the number of drives used along with the strip size chosen can have an impact on the wear leveling. The following table outlines the different options for number of drives vs strip size to achieve the optimal wear leveling on Intel NVMe SSDs. For default settings, please see the section on [Various Strip Sizes](#) in this document.

Strip Size No of drives	4	8	16	32	64	128
3	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
4	Optimal	Optimal	Optimal	Optimal	Suboptimal	Suboptimal
5	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
6	Optimal	Optimal	Optimal	Optimal	Optimal	Suboptimal
7	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
8	Optimal	Optimal	Optimal	Suboptimal	Suboptimal	Suboptimal
9	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
10	Optimal	Optimal	Optimal	Optimal	Optimal	Suboptimal
11	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
12	Optimal	Optimal	Optimal	Optimal	Suboptimal	Suboptimal
13	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
14	Optimal	Optimal	Optimal	Optimal	Optimal	Suboptimal
15	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
16	Optimal	Optimal	Suboptimal	Suboptimal	Suboptimal	Suboptimal
17	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
18	Optimal	Optimal	Optimal	Optimal	Optimal	Suboptimal
19	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
20	Optimal	Optimal	Optimal	Optimal	Suboptimal	Suboptimal
21	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
22	Optimal	Optimal	Optimal	Optimal	Optimal	Suboptimal
23	Optimal	Optimal	Optimal	Optimal	Optimal	Optimal
24	Optimal	Optimal	Optimal	Suboptimal	Suboptimal	Suboptimal

Note: It is left to the customer to determine the most effective combination of parameters (number of drives vs. strip size) to achieve their desired performance goals, usage models and drive endurance.

If a RAID volume is being migrated to RAID 5 (or a new RAID 5 volume is being created), the strip size chosen should be based off the most optimal performance as defined in the above table.



4.4.11 Intel VROC (VMD NVME RAID) 90 Day Trial Period

The Intel VROC 6.1 package comes with a 90 day trial period that will enable Intel VROC Premium mode (in Windows* only) for data volumes without requiring an Intel VROC Upgrade Key. This allows the user to test and experience Intel Premium mode for 90 days. The trial period will begin at the time that Intel VROC (VMD NVME RAID) is installed onto the system and Intel VROC (VMD NVMe RAID) volume is created.

4.4.11.1 Intel VROC (VMD NVME RAID) Trial Period Recommendations and Limitations

The following are key limitations with the Intel VROC (VMD NVME RAID) Trial Period feature

- The Intel VROC (VMD NVME RAID) Trial Period does not extend to the Intel VROC (VMD NVME RAID) UEFI HII environment. As a result, RAID volumes generated in the Intel VROC 6.1 GUI during this trial period will not be seen in the Intel VROC (VMD NVME RAID) UEFI HII user interface.
- Intel recommends **not** attempting to migrate a system drive into a RAID volume during the trial period. This is to prevent the system from becoming unbootable because trial period does not extend to the Intel VROC (VMD NVME RAID) UEFI HII environment.
- Intel recommends not using Intel VROC (VMD NVME RAID) Trial Period RAID volume for any mission critical data. It is only intended for evaluation purposes and the data cannot be guaranteed (either in the Intel VROC (VMD NVME RAID) UEFI HII environment or after the period expires).
- At this point and time, there is no option to disable the Intel VROC (VMD NVME RAID) Trial Period.

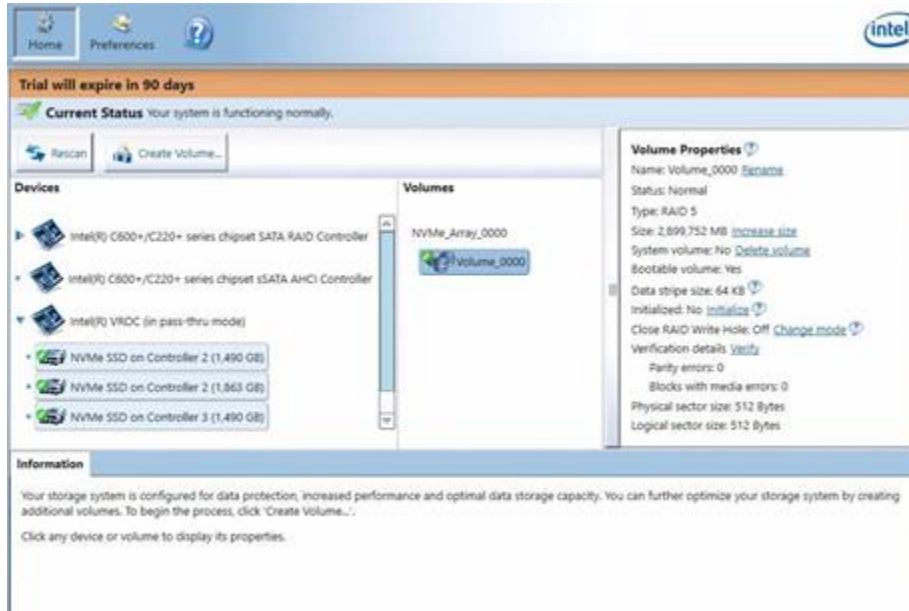
4.4.11.2 Intel VROC (VMD NVME RAID) Trial Period Usage

Since this Intel VROC (VMD NVME RAID) Trial Period mode only applies to platforms that do not have Intel VROC Upgrade Keys installed, the System Report will report that Intel VROC (VMD NVME RAID) is in Pass-thru mode.

Prior to initiating the trial period (by creating a RAID volume) the Intel VROC 6.1 GUI will show no indication of a trial period being activated. Once Intel VROC 6.1 GUI has been installed, the trial period will begin once the first trial RAID volume is created.

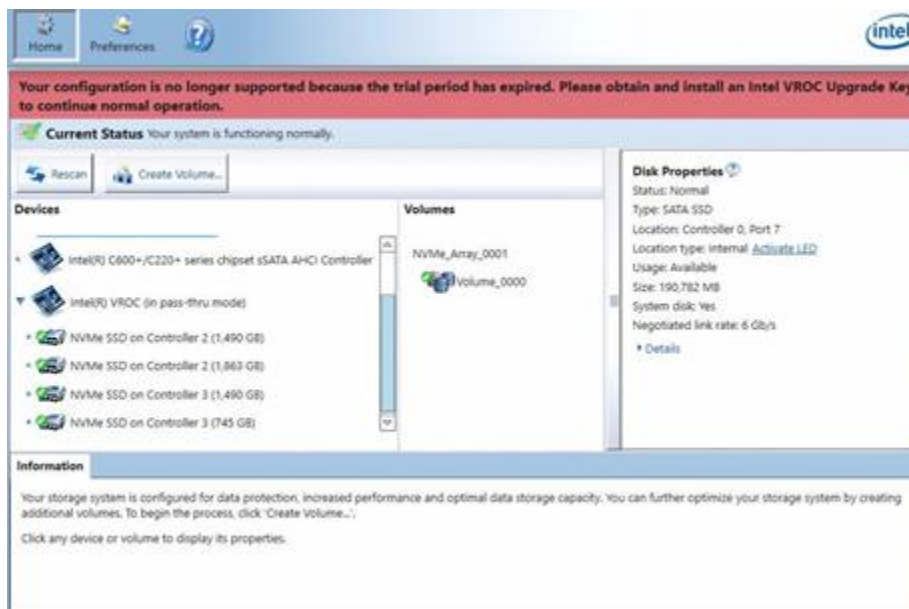


The status of the trial is displayed in the Intel VROC 6.1 GUI as:



During the trial period, the Intel VROC (VMD NVME RAID) UEFI HII will not display that RAID volume, and will show the attached drives as non-RAID disks.

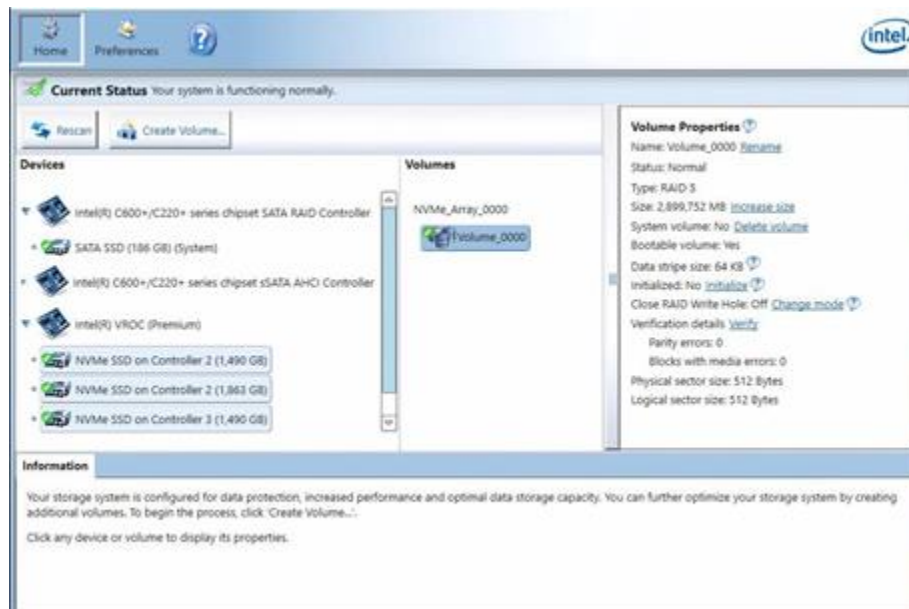
After 90 days, the trial period will expire and be displayed in the GUI as:



Note: At the end of the trial period, if a RAID volume is created, the expired license status will be displayed. The volume will be generated but may be unusable.



Once the Intel VROC Upgrade Key is installed in the platform (either during the trial period or after the trial period expires) the environment will become operational and all RAID volumes are accessible as normal.





5 Intel VROC 6.1 Common Features

This section outlines and describe those features that are common between the Intel VROC (VMD NVMe RAID), Intel VROC (NonVMD NVMe RAID) and Intel VROC (SATA RAID) packages.

The following is a summary of the key common features supported in both products.

Name	Key Features	
RAID	<ul style="list-style-type: none">• Support for Matrix RAID 0/1/5/10• Pass-through drives• Hot Plug with I/O• Hot Spare Disk• Auto Rebuild on Hot Insert• Check Pointing• UEFI using common metadata• SMART Support• RAID Volume Roaming• RAID Volume roaming between Linux* and Windows• On Line Capacity Expansion• RAID-Ready	<ul style="list-style-type: none">• Disk Coercion• Manual & Auto Rebuild• Instant Initialization• Volume creation/verify• Selectable Boot Volume• Email Alerting• RAID Level Migration (RAID 0, 1, or 10 to RAID 5)• Verify and Repair• RAID Write Hole closure• Native 4K Drives
Utilities	<ul style="list-style-type: none">• Install/Uninstall Utility• Configuration and Management Utilities	

5.1 Intel VROC 6.1 RAID Features and Functionality

This section outlines those features and functionality supported with the Intel 5.5 products.

5.1.1 RAID Levels

RAID 0 (striping)

RAID level 0 combines at least two (up to the max number) of drives supported by SATA(8)/sSATA(6) controller on the platform or up to 24 (NVMe) 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 drive simultaneously. However, data stored in a RAID 0 volume is not redundant. Therefore, if one drive fails, all data on the volume is lost.

The RAID 0 volume appears as a single physical drive with a capacity equal to (number of drives in the volume) X (the size of the smallest drive in the volume).

The maximum number of drives supported in a RAID 0 is up to the maximum number of drives supported by the platform

**RAID 1 (mirroring)**

RAID level 1 combines two 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 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 drive with a capacity equal to that of the smaller drive.

The maximum number of drives supported in a RAID 1 volume is 2 drives.

RAID 5 (striping with parity)

RAID 5 volume provides the capacity of $(N-1) \times$ smallest size of the drives, where $N \geq 3$ and:

- \leq max number of drives supported by SATA or sSATA controller
- ≤ 48 (NVMe)

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 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 combined with the built in 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 RAID 5 volume.

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.

The maximum number of drives supported in a RAID5 is the maximum number of drives supported by the platform.

RAID 10 (striping and mirroring)

RAID level 10 uses four 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 drive with a capacity equal to the two smallest drives of the four drive configuration (the only supported RAID 10 configuration). The space on the remaining two drives will be used for mirroring.

The maximum number of drives supported in a RAID 10 is 4.



5.1.2 Matrix RAID

The Intel VROC 6.1 family of products will support up to two logical RAID volumes on the same array. A RAID array simply refers to the set of disk drives that can be formed into a RAID volume. A RAID array can be created with a minimum of 2 drives.

When the second RAID volume is created, the creation process will automatically use all of the remaining available space.

Note: When utilizing a Matrix RAID configuration, avoid mixing a non-redundant RAID volumes (RAID0) with redundant RAID volumes (RAID 1/5/10). Any failed disk will fail the RAID0 volume and may result in the array not function properly.

5.1.3 Creating a RAID Volume

The Intel VROC 6.1 family of products will support the creation of RAID volumes in the following different ways:

- The Intel VROC PreOS environment
 - Intel VROC PreOS UEFI HII UI (*Only applies to Intel VROC (VMD NVMe RAID) and Intel VROC (SATA RAID) products*)
- The Intel VROC GUI in Windows*
- The Intel VROC Command Line Interface (CLI)

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

The configuration tools listed in the [Intel VROC 6.1 Environment](#) section can be used to create RAID volumes as follows.

Run **Rcfgxxxx.efi** from a USB key in an EFI Shell environment or **IntelVROCCLI.exe** in Windows* based 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 drives attached to the SATA Controller (Controller #0) on Port 0 and 1 with a strip size of 128 KB and a size of 120 GB:

```
C:\>rcfgsata.efi /C OEMRAID0 /DS 0.0 0.1 /SS 128 /L 0 /S 120
```

```
C:\> IntelVROCCLI.exe -C -l 0 -n OEMRAID 0-0-0-0 0-1-0-0 -s 128 -z 120
```

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

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 drives in the system connected to the SATA controller (configured to boot in legacy mode). The volume will be the maximum size allowable.

```
C:\>rcfgsata.exe /C OEMRAID0 (requires that only two disks can be attached to the system)
```

Note: At this point in the process if there are any significant discrepancies between the drives selected to be RAID members, a warning message will be displayed if one of the following conditions is encountered:

- There is a combination of SSDs and HDDs used
- There are at least two drives that have a size difference of more than 10%.



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

```
C:\>Rcfigxxx.efi /?
```

```
C:\> IntelVROCCLI.exe --help
```

5.1.3.2 Using the Intel® Rapid Storage Technology enterprise GUI

In this example, drives are connected to the CPU via the Intel VMD.

This example assumes that Intel VROC has been installed. The following example will step through the process of creating a 2-drive RAID 1 data volume.

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):

Start Menu ->All Programs -> Intel-> Intel Virtual RAID on CPU

2. From the Home screen, click the Create Volume button.
3. Select the controller to create a volume on.

In this case **Intel VMD controllers**. Select the RAID type **Real-time data protection (RAID 1)** and click **"Next"** to continue.

4. Specify the **Name** of the volume, select the two drives to be included in the volume, and click **"Next"**.

In this example, the name uses the default: **Volume_0000**.

Note: When configuring a volume, the application will only list the disks that meet the minimum 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 to the Creating a volume under help file in the GUI.

5. In the Confirm form, click **Create Volume**.

Note: At this point in the process if there are any significant discrepancies between the drives selected to be RAID members, a warning message will be displayed if one of the following conditions is encountered:

- a. There is a combination of SSDs and HDDs used
- b. There are at least two drives that have a size difference of more than 10%.

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

6. Click **OK** to close this dialog box.
7. Click **OK** to complete the process.

Under the Volumes section the new Array and RAID Volume are displayed. By selecting the RAID volume (**Volume_0000**), the "Properties" pane (right) will refresh to show the current status, properties, and available options of the newly created RAID volume.



5.1.4 Instant Initialization

The Intel VROC 6.1 family of products will support a newly created volume to be used immediately (no reboot required), protecting newly written data and creating parity data concurrently.

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

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

The VROC 6.1 product will initialize parity for the entire volume on all RAID 5 configurations after a dirty shutdown has occurred.

5.1.4.1 Intel® VROC 6.1 RAID Volume Initialization

The RAID volume initialization process ensures that the entire RAID volume has all of its RAID metadata writing. Under most conditions the initialization occurs in the background while data is being written to the volume. To ensure that all of the RAID volume is properly setup, it is recommended that the volume initialization be manually initiated. This is done through the Intel® VROC GUI.

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. From the Volumes pane, click the array or volume to which the user wants to modify.
3. The volume properties now display on the right.
4. In the pop up window Volume Initialization, select "Yes".

The initialization details are shown in the bottom **Information** pane as well as in the **Volume Properties** dialog.

5.1.5 Deleting a RAID Volume

RAID volumes can be deleted in three different ways. The method most widely used by end-users is the Intel VROC (SATA RAID) GUI. The second method is to use the Intel VROC Pre-OS environment (UEFI HII) user interface. The third way, used by OEMs only, uses the RAID command line configuration utility.

When a volume is deleted, the user creates available space that can be used to create new volumes. Note that the user cannot delete a system volume using this application because the operating system needs the system files to run correctly.

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

5.1.5.1 Using the Intel® VROC GUI

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. Under 'Home' click on the volume the user wants to delete.
The user will be presented with the volume properties on the right. Select the Delete volume link.
3. Review the warning message, and click "Yes" to delete the volume.

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



5.1.5.2 Using the Intel® VROC UEFI HII User Interface

Note: This section is OEM dependent. Where/how the OEM chooses to implement the UEFI UI is based on OEM preference. Use the following example for Intel CRB.

1. Enter the setup and configuration mode of the BIOS.
2. Navigate to EDKII Toolkit and press the <Enter> key.
3. Navigate to Intel® Virtual RAID on CPU and press the <Enter> key.

Alternate destinations may be Intel® VROC sSATA Controller or Intel® VROC SATA Controller.

4. Select the RAID volume the user wishes to delete.
Be careful not to delete a RAID volume the user does not wish to delete such as a boot RAID volume.
5. On the page that shows all of the RAID volume information, use the arrow keys to navigate to select the **Delete** option and press the <Enter> key.
6. Use the arrow keys to navigate to select **"Yes"** and press the <Enter> key.

5.1.5.3 Using the Intel® VROC CLI Tool

Using configuration tool, **IntelVROCCLI.exe**, in a Windows* based environment with the following command line flags to delete a RAID volume.

The following command line will instruct the utility to delete a RAID 0 volume named **OEMRAID0** on drives attached to the SATA:

```
C:\> IntelVROCCLI.exe -D -n OEMRAID
```

5.1.6 Degraded RAID Volumes

The Intel VROC 6.1 family of products will report RAID volumes as degraded when:

- A RAID 1 volume has one of its members is failed or is missing.
- A RAID 5 volume has one of its members is failed or is missing
- A RAID 10 volumes has at least one of its members is failed or is missing. It is possible for 2 drives to be either failed or messing and the RAID 10 will still be degraded.

5.1.7 Intel VROC GUI Support of Various Controllers

The Intel VROC 6.1 family of products will provide support for a single RAID management utility or GUI. This single Intel VROC GUI will display all supported controllers (SATA/sSATA, Intel VMD Controllers and/or (for Intel VROC (NonVMD NVMe RAID) supported Intel NVMe SSDs) and drives attached to them. The Intel VROC 6.1 GUI management tools will display each different controller and show all of the disks attached to that controller. The Intel VROC 6.1 GUI will allow creating/deleting/monitoring devices, RAID arrays and volumes on drives connected to the SATA, sSATA and Intel VMD controllers.



5.1.8 Various Strip Sizes

The Intel VROC 6.1 family of products will provide the ability to change strip size on existing volumes (migration required). Intel VROC 6.1 will support a strip size migration in conjunction with a RAID level migration.

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

Available Strip Size Configurations

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

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

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

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

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 (default value), 10 (default value)
128 KB	Best for audio and video editing.	RAID 0 (default value), 5

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



5.1.8.1 Setting the strip size when changing volume type

1. Under 'Home', click 'Create'.
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.

5.1.8.2 Setting the strip size when changing volume type

1. Under 'Home', in the storage system view, click the RAID volume that the user wants to modify.
The volume properties are now displayed on the right.
2. Click 'Change type'.
3. Make the necessary volume type and disk selections, and then select a new data strip size.
4. Click 'OK' to change the type of the existing volume.

The 'Manage' page refreshes and reports the new volume configuration.

5.1.9 Disk Coercion

Intel® VROC 6.1 family of products will provide support for Disk Coercion. When a RAID volume is created, this feature will analyze the physical disks and will automatically adjust (round down) the capacity of the disks to 95% of the smallest physical disk. This allows for the variances in the physical disk capacities from different vendors.

The VROC 6.1 UIs (Pre-OS, Windows* GUI and CLI tool) will provide an option to manually override this to be able to use all available disk space.

5.1.10 Flush Unit Access (FUA)

The Intel® VROC 6.1 family of products will honor all FUA commands received from the OS. This ensures that the cache on the attached drive is flushed when the command is received. This can have a negative impact on performance. For customers who are willing to accept the risk of potential data loss, they can improve the overall platform performance by having Intel® VROC ignore FUA commands received. This is done by creating the following Microsoft* Registry Keys.

- *Enable will trigger the Intel VROC (SATA RAID) driver to ignore FUA commands when received from the OS:*
`HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, IgnoreFUA, %REG_TYPE_DWORD%, 0x00000001`
- *Disable will configure the Intel VROC (SATA RAID) driver to continue to honor FUA commands:*
`HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, IgnoreFUA, %REG_TYPE_DWORD%, 0x00000000`

Note: It is left to the administrator (who is making the change) to know how to open up the Windows Registry editor and how to modify the entries and properly save the updates.



5.1.11 RAID Volume Roaming

The Intel VROC 6.1 family of products will support the ability to move all of the drives that make up a recognized RAID volumes between supported homogeneous controllers or platforms. This includes both Intel VROC (VMD NVMe RAID) and Intel VROC (SATA RAID).

Note: For Intel VROC, the new environment that the “old” Intel VROC RAID Volume is being moved to must match the configuration of the “old” platform/environment.

5.1.12 RAID Volume Roaming between Linux and Windows

The Intel VROC 6.1 family of products will support the ability to move RAID data volumes (configured appropriately) between Linux and Windows* environments and the RAID data volumes will be recognized and available for use. This also applies to a manufacturing environment where the RAID volume created for the OS installation needs to support either a Windows* environment or a Linux environment.

5.1.13 Support of Pass-Through Drive

The Intel VROC 6.1 family of products will support the ability to install to and to boot from a single pass-through drive (not part of a RAID array or volume).

Intel VROC 6.1 will also be able to support a platforms power management states (i.e. Power-on, Restart, Sleep and Hibernate) on single pass-through drives that contain the system OS.

5.1.14 Failed Drive Reinsertion

The Intel VROC 6.1 family of products will support the ability to recognize a failed drive re-inserted into the system. If array is properly configured, Intel VROC 6.1 will attempt to bring the drive into the array and rebuild the volume with that drive. If not, Intel VROC 6.1 will mark the drive as failed or off line in the GUI.

5.1.15 LED Management

The Intel VROC 6.1 family of products will support the ability to perform basic LED management of the status LEDs within SFF-8485 (or equivalent) compliant backplane and with the blink patterns as defined in SFF-8489.

For details on how to implement the Intel VMD Method, refer to the *LED Management for PCI Express SSDs with Intel RAID* document (IBL Doc ID: 334005-004US).

Intel VROC 6.1 LED Management provides the capability to configure different LED behaviors for each of the supported controller configurations:

1. NVMe LED management via Intel VROC,
2. SATA LED management via the PCH controller (SATA/sATA) in RAID mode
3. Custom NVMe LED Management via .DLL files generated by the platform OEMs.

For the platform the PCH supports LED management through SGPIO signaling.



The LEDs management capabilities are outlined as follows:

Event	Behavior	Configuration options	Default Setting
"Activate LED"	The ability to identify a specific device in an enclosure by blinking the status LED of that drive in the defined Locate pattern. The duration of this Locate pattern can be customized to meet the customers' needs.	Adjustable time of blinking. The range can be from 1 to 3600 seconds. Any value entered that falls outside of this range will automatically be interpreted as 12 seconds.	12 seconds
Drive fails	<p>The ability to identify a drive that is in a Failed state by illuminating the status LED of that device in a defined Fail pattern. This applies to RAID volumes in a degraded or failed state.</p> <p>A drive is generally marked as "failed" (causing the Fail LED light to illuminate) when communication to that drive has been interrupted and recovery attempts have failed to reestablish a connection. Some examples of (and by no means all) conditions that can cause a drive to be marked as failed:</p> <ul style="list-style-type: none"> • Signal errors on the line • Drive has stopped responding • The drive Bad Block management table is full 	<p>The Fail pattern will be illuminated:</p> <ol style="list-style-type: none"> 1. From the time a failed drive is identified until it is: <ul style="list-style-type: none"> -physically removed OR -the RAID volume, that contains the failed drive, is either deleted or physically removed 2. From the time a <ul style="list-style-type: none"> - non-Failed drive that is part of a RAID volume is removed OR - failed drive is identified and removed <p>UNTIL:</p> <ul style="list-style-type: none"> - a new drive is inserted into the same slot the failed drive resided in OR -the platform is rebooted 	Option 1
Managed unplug	When a managed hot unplug is being administered, the status LED of that drive will blink in the defined Locate pattern until drive is physically ejected.	None	Enabled
RAID volume Initialization or Verify and Repair Process	When a RAID volume is in either the Initialization or Verify and Repair processes, the status LEDs will blink in the defined Rebuild pattern on all drives in RAID volume until the process completes.	<ol style="list-style-type: none"> 1. Enabled – status LEDs will blink in the Rebuild Pattern 2. Disabled – status LEDs will not blink 	Enabled
RAID volume Rebuilding options	When a RAID volume is in a Rebuild state, the status LED(s) will blink in the defined Rebuild pattern on either the specific drive being rebuilt or on the entire RAID volume rebuilt	<ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives) 	Disabled (only on one drive)
RAID volume is migrating	When a RAID volume migration is in process, the status LEDs will blink in the defined Rebuild pattern on all drives until the process is complete	<ol style="list-style-type: none"> 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs) 	Enabled

Note: It is left to the administrator (who is making the changes) to know how to open up the Windows Registry editor, how to modify the entries and properly save the updates. These options are available and can be customized via the Windows* registry with the following keys:



- The Registry Key that is used to set the “Activate LED” Locate blink pattern duration is “LedLocateTimeout”

HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, LedLocateTimeout, %REG_TYPE_DWORD%, 0x0000000C

The value entered for this Registry Key is in Hexadecimal. In this example, 0x0000000C is translated to 12 seconds.

- The Registry Key that is used to customize the Drive fails behavior is “LedFaultOnEmptyPortBehavior”

HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, LedFaultOnEmptyPortBehavior, %REG_TYPE_DWORD%, 0x00000000

The Registry key set to a value of 0x00000000 will select configuration option 1. The value of 0x00000001 will select configuration option 2.

- The Registry Key that is used to customize the RAID volume Initialization or Verify and Repair Process settings is “LedBehaviorForInitilazing”

HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, LedBehaviorForInitilazing, %REG_TYPE_DWORD%, 0x00000001

This Registry Key set to a value of 0x00000000 will disable this configuration option. The value of 0x00000001 will enable this option

- The Register Key that is used to customize the RAID volume Rebuilding options settings is “LedBehaviorForRebuilding”

HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, LedBehaviorForRebuilding, %REG_TYPE_DWORD%, 0x00000000

This Registry key set to a value of 0x00000000 will disable this configuration option and set the LED behavior to blink the drive that is being rebuilt. This Registration key set to a value of 0x00000001 will enable this configuration option and will blink all of the status LEDs of the RAID volume being rebuilt.

- The Registry Key that is used to customize the RAID volume is migrating configuration option is “LedBehaviorForMigrating”

HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, LedBehaviorForMigrating, %REG_TYPE_DWORD%, 0x00000001

This Registry Key set to a value of 0x00000001 will enable this configuration option to blink the Status LEDs in a Rebuild pattern while the RAID volume is migrated. This Registry Key set to a value of 0x00000000 will disable this configuration options and the Status LEDs will not blink during a RAID volume migration.

Note: Intel VROC 6.1 LED Management only applies to drives that reside within a supported drive backplane (NVMe and/or SATA). Drives that are connected either by an I/O cable, PCIe add-in card or plugged directly into the motherboard (M.2) will not have LED Management support.

Note: Intel VROC 6.1 LED Management does not include drive activity LED management.



5.1.15.1 Failed Drive Indicator

When a member drive of a RAID volume is marked as failed, the Intel VROC 6.1 GUI will indicate that the disk is failed and if applicable set the status LED to a failed state.

5.1.15.2 Activate LED

The Intel VROC 6.1 family of products will support the ability to blink the status LED on specific drive when the Activate LED link is selected in the Intel VROC 6.1 GUI.

The assumption is made that the platform is properly configured with an enclosure that supports LED management.

1. Run Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. From the Home screen, in the “Devices” pane, select the device and navigate to the “Disk Properties” pane to run Verify and Repair on. In the Volume Properties pane, click on Verify by Verification details.
3. Locate the “Activate LED” link and click on the link. View the Status LED of the enclosure blink in the pattern specified by SFF-8489.

Note: For the SATA/sSATA controller, LED management is only supported when these controllers are in RAID mode.

5.1.15.3 Rebuild LED

The Intel VROC 6.1 family of products will support the ability to blink the status LEDs on all drives that are members of the degraded RAID volume as part of the RAID rebuild process. This occurs automatically on supported platforms and no user interaction is required.

5.1.15.4 UEFI LED Management

The Intel VROC 6.1 family of products will support LED management within the UEFI environment. This includes Locate, Rebuild and Fail. The LED management in the UEFI environment will only support the default values specified in the table above. Any changes that are made to the configuration in the OS environment will not be reflected in the UEFI environment. This means that the LED behavior seen in the UEFI environment may very well be different from what is seen in the OS environment (if the default values have been modified). The ability to initiate the “Locate” option is available in the Intel VROC 6.1 UEFI HII. Rebuild and Fail blink patterns will occur automatically based on the condition of the RAID volume state or member drives state.

5.1.16 Modify Volume Name

The Intel VROC 6.1 family of products will be able to update the name of a RAID volume in the normal state when requested by the Intel VROC 6.1 GUI or the Command Line Interface (CLI) tool.

5.1.17 Hot Plug (Surprise and Managed)

The Intel VROC 6.1 family of products will support the ability to Hot Plug (remove and replace) disk drives on properly configured, operational, platform whether or not I/O is being processed. Hot-Plug, also referred to as hot swap, is a feature that allows disk drives (SATA or NVMe) to be removed or inserted while the system is powered on



and running under a Windows* operating system. As an example, Hot-Plug may be used to replace a failed drive that is in an externally-accessible drive enclosure.

Managed Hot Plug support only applies to Intel VROC for both Pass-thru drives and drives in a RAID volume. Intel VROC (SATA RAID) currently does not contain support for managed Hot Plug for SATA drives attached to the SATA or sSATA controllers.

Managed Hot Plug refers to the process of informing the system that a drive will be removed. To accomplish this, Intel VROC 6.1 has an option in the Intel VROC GUI that is used to provide this notification.

1. Open the Intel VROC 6.1 GUI
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
In the "Devices" pane, select NVMe drive to be removed
2. In the "Disk Properties" pane, look for "Location type" and click on "Remove Disk"
3. A message window will pop up asking "Are you sure you want to remove this disk?" Click "Yes"

Intel® VROC 6.1 supports Surprise Hot Plug on both Pass-thru drives and drives in a RAID volume. Surprise Hot Plug refers to removing a drive at any point and time.

Note: To be able to take advantage of this feature, the system and platform BIOS must have this feature enabled. Please consult the platform design documentation for additional information.

Note: Surprise and Managed Hot Plug is only supported in platforms that have Hot Plug capable backplanes. Using the data cable and/or power cable plugging can result in unknown behavior.

5.1.17.1 How to Enable the BIOS for Hot Plug

This series of instructions will guide users through the Intel BIOS configuration for enabling a RAID with a spare drive addition. This allows for configurations, such as a data volume using a file storage service, using multiple drives, to have backup drives ready in case a failure occurs, enabling Auto Rebuild to assist in the recovery.

5.1.17.1.1 Enabling Surprise Hot Plug for Intel VROC

The following steps are an example of how to enable surprise hot plug in the Intel CRB platform BIOS:

1. From the boot options menu, select the option that will allow the user to enter the BIOS setup.
2. With EDKII Menu highlighted, press <Enter>.
3. Navigate to Socket Configuration, then press <Enter>.
4. Navigate to the selection that reads PCIe Hot Plug. Press <Enter> to open the options menu, the user may toggle this setting between disabled and enabled. Select Enable, and press <Enter> to set selection.

Note: This will enable the feature for all NVMe drives that are associated with the system. This is a distinction from Hot Plug as it works with PCH drives. Please consult the user's platform documentation because each BIOS is different and the steps taken to enable this feature may be different.



5.1.17.1.2 Enabling Surprise Hot Plug for PCH

1. From the boot options menu, select the option that will allow the user to enter the BIOS setup.
2. With the EDKII Menu highlighted, press <Enter>.
3. Navigate to Platform Configuration, and press <Enter>.
4. Navigate to PCH Configuration, and press <Enter>.
5. Navigate to PCH SATA or PCH sATA Configuration (as appropriate) and press <Enter>.

Note: Each port may be individually enabled for Hot Plug. The user may turn this feature on for all or none as appropriate as the administrator. Please consult the user's platform documentation because each BIOS is different and the steps taken to enable this feature may be different.

5.1.18 Hot Spare Disk

The Intel VROC 6.1 family of products will support the ability to set a drive as a hot spare that would automatically be used to rebuild a failed or degraded RAID volume without any user interaction. This support is provided in the Intel VROC GUI as well as in the Pre-OS images.

5.1.18.1 Global Hot Spare

The Intel VROC 6.1 family of products will support the Hot Spare disks definitions as Global hot spares. This means that the hot spare can be applied to any RAID array under the following conditions:

- A Hot Spare(s) defined on the SATA controller will be available for any RAID Array on the SATA controller
- A Hot Spare(s) defined on the sATA controller will be available for any RAID Array on the sATA controller
- A Hot Spare(s) defined on a VMD domain will be available for any RAID array on any of the VMD domains in the system.

5.1.18.2 Mark a Disk as a Spare

Marking a disk as a "spare" allows the user to designate an available disk as the default destination for automatic volume rebuilds in the event of a failed, missing or "at risk" (SMART event) array disk.

The Action is only available for non-system disks in a normal state. The "spare" disk must be connected to the same controller as the disk that it is supporting. The maximum number of "spare" disks is determined by the maximum number of disks supported by the controller.

Since all Hot Spare drives are considered Global (and not dedicated to a specific RAID Array), the size of the disk, being designated as a Hot Spare, is not checked.

Setting a disk as a "spare" disk can be accomplished in the following ways:

- The Intel® VROC GUI
- The Intel® VROC provides Pre-OS environment (UEFI HII) user interface support
- The Intel® VROC CLI tools



5.1.18.2.1 Mark a Disk as Spare in the GUI

The Intel VROC 6.1 family of products will provide an option in the GUI to mark a disk as a Hot Spare. The assumption is made that the appropriate additional drive(s) have been physically installed in the system. The installed drive(s) are also assumed to meet the size requirements to be identified as spare drives for the RAID array they will be associated to. To identify a drive as a Hot Spare in the GUI use the following steps:

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. In the Devices pane that displays the drives, select the drive to be marked as a spare
3. In the Disk Properties pane click on "Mark as Spare"

5.1.18.2.2 Mark a Disk as Spare in UEFI HII

The Intel VROC 6.1 family of products will provide an option to mark a disk as spare in the UEFI HII UIs.

The assumption has been made that the appropriate additional drive(s) have been physically installed on the system. The installed drives are assumed to meet the size requirements to be identified as spare drives for the RAID array they will be associated to.

1. Enter the setup and configuration mode of the BIOS.
2. Navigate to EDKII Toolkit and press the <Enter> key.
3. Navigate to Intel® Virtual RAID on CPU and press the <Enter> key.
Alternate destinations may be Intel® VROC sSATA Controller or Intel® VROC SATA Controller.
The non-RAID disk will be listed below the RAID volume in a category below.
4. Highlight the target drive desired and press the <Enter> key.
This screen presents the Physical Disk Information Page. Depending on the nature of the RAID array, there may be two options available. Mark as Spare and Mark as Journal.
5. For the intended purpose of marking as a Spare drive, highlight Mark as Spare, and press the <Enter> key.
A window will be displayed asking "Are you sure the user wants to mark the disk as Spare?"
"Marking disk as Spare will remove all data on the disk."
6. To proceed, highlight Yes, and press <Enter>.

5.1.18.2.3 Return Hot Spare Disk Back to Available

The Intel VROC 6.1 family of products will support the ability to return a drive marked as a Hot Spare back to available. When a drive is reset back to available, it can be used as a normal drive.

Return Hot Spare Disk Back to Normal in GUI

To return a drive marked as a Hot Spare back to available in the GUI, use the following steps:

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. In the Devices pane that displays the drives, select the drive that is marked as a spare
3. In the Disk Properties pane click on "Return to available"



Return Hot Spare Disk Back to Normal in UEFI HII

To return a drive marked as a Hot Spare back to available in the UEFI HII, use the following steps:

1. Enter the setup and configuration mode of the BIOS.
2. Navigate to EDKII Toolkit and press the <Enter> key.
3. Navigate to Intel® Virtual RAID on CPU and press the <Enter> key.
Alternate destinations may be Intel® VROC sSATA Controller or Intel® VROC SATA Controller.
The non-RAID disk will be listed below the RAID volume in a category below.
4. Highlight the target drive that is designated as a Hot Spare press the <Enter> key.
This screen presents the Physical Disk Information Page. Depending on the nature of the RAID array.
Select Reset to available and press the <Enter> key.
5. To proceed, highlight Yes, and press <Enter>.

5.1.19 Drive Requirements

This section outlines the requirements and expectations for drives used in an Intel VROC RAID Array/Volume.

5.1.19.1 Clean Drive Status

Intel VROC 6.1 RAID management expects that all drives used in or added to an Intel VROC RAID Array/Volume start as new or clean drives. This means that the drive does not contain RAID metadata from a previous RAID volume. Using drives that are not new or clean can result in various unexpected behaviors. This applies to drives that are used as Hot Spares and for drives used in a RAID volume rebuilds process.

5.1.19.2 Minimum Disk Size Requirements

Intel VROC 6.1 RAID management has minimum disk size requirements for certain RAID management activities.

5.1.19.3 Disk Addition: Minimum Required Disk Size

The Intel VROC 6.1 family of products will support the ability to add a disk to a RAID Array. The size of the disk added must be at least as large as the smallest disk in the existing RAID array. If the disk does not meet this minimum size requirements it will not be allowed to be added.

5.1.19.4 Targeted Disk Replacement: Minimum Required Disk Size

The Intel VROC 6.1 family of products will support the ability to replace a targeted RAID array member. The capacity of the replacement member must be greater than or equal to the sum of the greatest allocated LBA and the metadata size of the targeted member.

5.1.20 Auto Rebuild

The Intel VROC 6.1 family of products will provide support for the ability to automatically rebuild a failed or degraded RAID volume. This feature will begin when a member disk of the array has failed and a suitable replacement disk with sufficient capacity is available. As soon as the failure occurs the rebuild process will begin automatically, using the marked Hot Spare disk, without user intervention.



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

- Another free disk is plugged into the same directly attached physical location as the failed drive.
- The newly inserted disk size is at least as large as the amount of space used per disk in the current array.
- The newly inserted disk must be the same type (SATA or NVMe) as the disk being replaced or the rebuild will not start.
- If the newly inserted disk contains Intel® VROC metadata with current status of member being offline or contains no Intel® VROC metadata.
- The newly inserted disk has not reported a SMART event.
- Automatic rebuild support will default to OFF for Intel® VROC and can be enabled through the Intel VROC GUI.

The following table summarizes the functionality:

Controller	Auto Rebuild Support	Action
Intel VMD or Supported Intel NVMe SSDs /SATA/sSATA	No Hot Spare previously marked	No auto rebuild: Manual steps required to rebuild array using new disk
Intel VMD or Supported Intel NVMe SSDs /SATA/sSATA	Auto rebuild conditions described above are met.	Auto rebuild starts without any user intervention
Intel VMD or Supported Intel NVMe SSDs /SATA/sSATA	One or more of the above conditions was not met.	No auto rebuild: Manual steps required to rebuild array using new disk

5.1.20.1 Auto Rebuild to a Reinserted RAID Member

The Intel VROC 6.1 family of products will support the ability to automatically initiate RAID rebuild when a RAID volume member drive is removed and reinserted (into the same slot it was removed from) after the RAID volume becomes degraded. If the RAID member drive is reinserted after a few seconds and that drive is seen as still intact with the rest of the RAID volume, after the reinserted RAID member drive is processed (discovered, enumerated, updated, etc.) the volume may go directly back to Normal. This feature is independent of the Rebuild on Hot Insert feature. As a result this functionality is independent of how Rebuild on Hot Insert is configured.

5.1.20.2 Auto Rebuild on Disk Missing/Fail Event

When an Array has a Hot Spare configured and a RAID member disk is removed (or fails) from a redundant RAID volume, Intel VROC 6.1 will claim the spare drive for that RAID volume and automatically begin to rebuild the RAID volume.

5.1.20.3 Auto Rebuild on Hot Insert

The Intel VROC 6.1 family of products will support the ability to initiate an automatic RAID rebuild when a new physical disk (of the appropriate size) is hot inserted into the same slot the failed RAID member drive was removed from. A RAID volume member drive that has failed (or may soon fail) can be hot removed. When the system is configured appropriately, that failed member drive can be hot removed and if the new drive (of appropriate size) is hot inserted into the same slot, the rebuild process will begin automatically. The rebuild process will begin only when:

- The new disk drive inserted meets the minimum size requirements of the:
 - a. drive removed
 - b. array containing the degraded RAID volume
- The remaining healthy member disks are of the same type of drive
- The newly inserted disk is healthy (no recorded SMART event)
- The newly inserted disk is a pass through disk (not a member of a different RAID volume)



The Intel VROC 6.1 GUI shall give the user the ability to enable and disable this feature.

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. In the “Devices” pane select the controller that contains the RAID member disks (or the drives that will be used to create the RAID volume)
3. In the “Controller Properties” pane, look for “Rebuild on hot insert” and “Enable” or “Disable”. This feature is disabled by default.

Rebuild on hot insert is only applicable for a redundant RAID volume that is in a degraded state. (See RAID Volume States for specific definition of degraded state.)

This feature is only applicable to a single operational OS session. If the OS is rebooted, once the drive is inserted, the RAID Rebuild process will need to be restarted manually.

Note: Rebuild on Hot Insert is disabled by default.

5.1.20.4 Error Threshold Monitoring/Handling

The Intel VROC 6.1 family of products will support the ability to initiate an automatic RAID rebuild to a marked hot spare drive when a drive SMART event alert has occurred that indicates a failure (Windows* Only).

5.1.20.5 Manually Invoked Rebuild

The Intel VROC 6.1 family of products will support a manual method to initiate a RAID volume rebuild if a hot spare has not been configured or is not available or Auto Rebuild has been disabled. For the Rebuild process to begin, the new drive being added to the Array/Volume must:

- Meet the minimum size requirements of the:
 - a. drive being replaced
 - b. array containing the degraded RAID volume
- The remaining healthy member disks are of the same type of drive
- The new disk is healthy (no recorded SMART event)
- The new disk is a pass through disk (not a member of a different RAID volume)

5.1.20.6 Rebuild Resumption

The Intel VROC 6.1 family of products will support continuing a RAID Volume rebuild after a system shutdown, Sleep state or Hibernation state.

5.1.21 Verify and Repair

The Intel VROC 6.1 family of products will provide support for RAID volume verification and repair. This feature identifies any RAID metadata inconsistencies or bad data on a RAID 0, RAID 1, RAID 5, or RAID 10 and reports the number of inconsistencies or number of blocks with media errors found during RAID volume data verification. Those inconsistencies are automatically repaired. The Verify and Repair process operates on the entire RAID volume, and can only be run on one RAID volume at a time.



The following describes what occurs for each RAID level:

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

Intel VROC 6.1 will clear the LBA information from a Bad Block List (BBL) when a successful write command is issued to that LBA. This also applies when a RAID volume is being initialized.

5.1.22 Starting Verify and Repair Process

Following are the different methods in which the Verify and Repair process is either automatically started or manually initiated.

5.1.22.1 Dirty Shutdown

The Verify and Repair process will automatically begin after the platform encounters an unplanned/unexpected power outage (Dirty Shutdown). This will occur when data is being read from or written to a RAID volume. If the RAID volume is intact, the RAID volume may come up as Normal.

5.1.22.2 Manual Initiation

The Intel VROC 6.1 family of products require that a RAID volume be initialized prior to manually starting the "Verify and Repair" process. Follow the below steps to start RAID volume data verification and repair process.

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU
2. From the Home screen, in the "Volume" pane, select the Volume to run Verify and Repair on. In the Volume Properties pane, click on Verify by Verification details.
3. In the pop-up window select "Verify"

Note: For RAID 0 the verification process starts once the user clicks 'verify'. For RAID 1, 5 and 10, 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 a repair, the user can select this box and then click 'verify'.

4. The verification progress is shown the Information pane at the bottom.

When the verification process is complete and the volume status is set to normal, under the volume properties to the left, the user can view the number of verification errors, verification errors repaired and blocks with media errors that were found.



5.1.23 Read Patrol

The Intel VROC 6.1 family of products will provide support for Read Patrol, which checks the RAID volumes for errors that could result in a failure. The checks are done periodically in background and will verify all sectors of all RAID volumes on SATA and sSATA controllers as well as volumes behind Intel VMD. If an issue is discovered an attempt at corrective action is taken. Read Patrol can be enabled or disabled manually.

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

Note: Read Patrol does not extend to drives marked as Hot Spares.

The VROC 6.1 driver shall execute the Read patrol functionality on RAID volumes that are in a normal state (not degraded rebuilding or migrating). For a redundant volume if the Intel VROC Driver detects a bad block error it will attempt to recover from that error by reassigning the bad block and writing the recovered data to the reassigned location. Read patrol will be run on all volumes simultaneously.

For each power-cycle the VROC 6.1 driver will record the event but not store all errors across all RAID volumes being scrubbed. The VROC 6.1 driver will display the following read patrol error information: unrecovered LBAs and its corresponding RAID Volume and physical disk.

The Intel VROC 6.1 GUI will display an option to enable or disable the Read Patrol feature and the feature will be "disabled" by default.

5.1.24 Check Pointing

The Intel VROC 6.1 family of products will support the ability to perform Check Pointing to be able to track forward progress on read patrol, array rebuilds and volume migration if interrupts occur. Upon a system restart, the operation will restart from the last valid stage reached.

5.1.25 Bad Block Management

The Intel VROC 6.1 family of products will provide support for Bad Block Management.

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

This feature is supported in the Pre-OS as well as the Windows* driver.

5.1.26 Intel RAID Write Hole Closure

The Intel VROC 6.1 family of products will support the ability to close the RAID Write Hole scenario in RAID 5 configurations. This applies to Intel VROC Enabled platforms.

RAID Write Hole (RWH) is a fault scenario, related to parity based RAID. It occurs when a power-failure/crash and a drive-failure (e.g., strip write or complete drive crash) occur at the same time or very close to each other. When


these system crashes and disk failures are correlated events, this can lead to silent data corruption or irrecoverable data failure due to lack of atomicity of write operations across member disks of a parity based RAID volume.

Intel VROC RWH closure feature eliminates such vulnerability by using Partial Parity Logging (PPL) to non-volatile memory of the member disks.

5.1.26.1 RAID Write Hold Closure Example and Man Page

RAID 5 read modify write sequence failure scenario:


1. Receive host write request
2. Read old data and read old parity
3. Partial parity = old data XOR old parity
3. New parity = new data XOR partial parity
4. Write New Data and Write New Parity
5. Complete the write request to the OS



- PPL is an intermediate calculation, saved to NV media, to recover from RWH; fewer writes than prior art
- May be recorded on the parity drive (rotated on member drives)

RAID 5 read modify write sequence failure scenario:

1. Receive host write request
2. Read old data and read old parity
3. Partial parity = old data XOR old parity
3. New parity = new data XOR partial parity
4. Store PPL to one of the RAID member drives (parity drive for a given strip)
5. Write New Data and Write New Parity
6. Complete the write request to the OS



RWH is a potential condition when using RAID5 that during the time that a RAID strips is being written, a power loss is encountered. With HWRAID and battery backed DRAM, logging could be used to recover. With Intel VROC 6.1 and RWH, a journaling SSD can be added to preserve the Partial Parity and reduce the potential data loss issue

5.1.26.2 RAID Write Hole Closure (RWH) Details

The Intel VROC 6.1 family of products will support the option to configure this feature in the Intel VROC 6.1 GUI while the system is running, without requiring a platform reboot. This feature is set to disabled by default.

It is important to realize that when this feature is enabled

- The system may experience a performance impact due to the extra writes.
- The volatile cache of devices, which don't have power loss protection (Power Loss Imminent (PLI) feature), will be turned off to prevent possible data loss that may be caused by a surprise system power loss or a drive hot removal. If a device's volatile cache cannot be turned off, RWH Closure will not be enabled.
- RAID volume rebuild times may increase because of the device cache being disabled.

Note: RWH cannot be enabled if a RAID volume is in the process of migrating.

5.1.26.2.1 RAID Write Hole Closure (RWH) Windows

The Intel VROC 6.1 driver will refrain from rebuilding a RAID 5 volume before or during the RWH recovery process.

If the RWH condition occurs during a recovery process and the recovery has not completed, Intel VROC 6.1 will attempt to resume the recover process from where it was interrupted.

If the RWH condition occurs during the process of the OS going into hibernation mode (S4) the RWH recovery process will be able to fix the failure condition for all the data being written during hibernation.



5.1.26.2.2 RWH Pre-OS

The Intel VROC 6.1 UEFI drivers support the ability to recover from the RAID 5 volume invalid state caused by a RWH condition for all enumerated RAID 5 volumes during system boot.

Note: When creating a RAID 5 Volume with RWH enabled and the volume must be usable in either Windows* or Linux, the RWH closure feature must be enabled with Distributed PPL. The Linux environment does not support RWH closure using a Journaling Drive.

To enable RWH with “Journaling Drive” options, the user needs to complete the create process with RWH as “Disabled”, then go back one page. At this point the user selects an available disk and then set its properties to “Journaling Drive”.

5.1.26.2.3 RWH Backwards Compatibility

For RAID 5 volumes on drives created in previous versions of the product, Intel VROC 6.1 will be able to leverage the older RWH closure mechanisms.

5.1.26.3 RWH Recovery

The Intel VROC 6.1 drivers enter into a RWH recovery state when a RAID 5 volume encounters a RWH failure condition. There are two levels of this RWH recovery. The first level is the recovery after a dirty shutdown. In this case the recovery process goes through the intent log (PPL) to make the RAID volume consistent for in-flight stripes that might be inconsistent. The second level recovery is when dirty shutdown is followed by a drive failure. In this case, the recovery process references the PPL to first make things consistent despite a member drive not being present.

The RWH recovery process will resume if the system is shut down or restarted.

During a RWH recovery process:

- RAID volume rebuild is disabled. Once the recovery is complete and the RAID 5 data is consistent, then we can begin the rebuild process.
- Verify and Repair on a dirty shutdown is disable. The Verify and Repair process is covered by the RWH recovery process.

5.1.27 Controller Default Values

The Intel VROC 6.1 family of products support the ability to set or change the controller's default parameters (that are not directly associated with an existing RAID volume):

- Rebuild on Hot Insert
- Read Patrol

These default values can be set or changed for the platform's manufacturing environment. There are two ways to make these changes. The first method is to utilize the CLI utility, included in the Intel VROC 6.1 package, to define the default behavior. The second method is to use the Windows* registry key option.



5.1.27.1 Setting Default Rebuild on Hot Insert (ROHI)

The default behavior for the ROHI option is “disabled”. To change this behavior the CLI tool can be used with the following option:

- `--manage --change-rohi enable|disable --controller controllerName | --controllerMode SATA|sSATA|VMD`

Where the “controllerName” will be the name of the controller as shown when using the `-l` option of CLI command. Either “`--controller controllerName`” or “`--controllerMode SATA|sSATA|VMD`” can be use. For example:

- `IntelVROCCLI.exe --manage --change-rohi enable --controllerMode SATA`

The following Windows* registry setting can be used as well:

- *HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, RebuildOnHotInsert, %REG_TYPE_DWORD%, 0x0000000 (disabled) | 0x00000001(enabled)*

5.1.27.2 Setting Default Read Patrol:

The default behavior for the Read Patrol option is “disabled”. To change this behavior the CLI tool can be used with the following option:

- `--manage --read-patrol enable|disable --controller controllerName | --controllerMode SATA|sSATA`

Where the “controllerName” will be the name of the controller as shown when using the `-l` option of CLI command. Either “`--controller controllerName`” or “`--controllerMode SATA|sSATA|VMD`” can be use. For example:

- `IntelVROCCLI.exe --manage --read-patrol enable --controllerMode SATA`

The following Windows* registry setting can be used as well:

- *HKLM, System\CurrentControlSet\Services\iaStorE\Parameters\Device, ReadPatrol, %REG_TYPE_DWORD%, 0x0000000 (disabled) | 0x00000001(enabled)*

5.1.28 Driver Dirty Shutdown Recovery

The Intel VROC 6.1 drivers will permanently recover from the RAID 5 volume invalid state caused by Dirty Shutdown condition occurrence for all RAID 5 volumes in the system which were exposed to I/O interruption (e.g. dirty shutdown).

5.1.28.1 Driver DS Recovery Caused by Hot-plug

The Intel VROC 6.1 driver will permanently recover from the RAID 5 volume invalid state caused by Dirty Shutdown condition occurrence when the RAID 5 volume is discovered by the driver after hot-plug of all the member drives.

5.1.28.2 Driver DS Recovery Caused by Driver Start

The Intel VROC 6.1 drivers will permanently recover from the RAID 5 volume invalid state occurrence, caused by a Dirty Shutdown condition, when the RAID 5 volume is discovered by the driver during system boot.



5.1.28.3 Driver DS Recovery Caused by Enabling Drives

The Intel VROC 6.1 drivers will permanently recover from the RAID 5 volume invalid state caused by Dirty Shutdown condition occurrence when the RAID 5 volume is discovered by the driver after enabling all the member drives except the failed drive in the device manager.

5.1.28.4 DS Condition and Hibernation Process

If the Dirty Shutdown condition has occurred during the process of the OS going into hibernation mode (S4) the recovery process will be able to fix the Dirty Shutdown condition for all the data being written during hibernation.

5.1.28.5 Interrupted PPL Write - DS Recovery

If the PPL write request has been interrupted and PPL was not fully written the Dirty Shutdown recovery will not be performed for this particular RAID 5 I/O request.

5.1.28.6 Failed DS Recovery - Verify Volume

If Dirty Shutdown recovery fails and RAID volume was initialized, the Intel VROC drivers will perform a 'Verify and Fix' operation on the entire RAID volume.

5.1.28.7 Failed DS recovery - initialize volume

If Dirty Shutdown recovery fails and RAID volume was not initialized, the Intel VROC drivers will perform an 'Initialize' operation on the entire RAID volume.

5.1.28.8 Dirty Stripe Journaling

The Intel VROC 6.1 family of products support Dirty Stripe Journaling (DSJ). DSJ is used to help speed up RAID 5 write power loss recovery by storing the write stripes that were in progress at the time of the failure. The DSJ allows rapid recovery without having to rebuild the entire volume. The DSJ is only utilized when disk write cache is DISABLED.

5.1.28.9 Partial Parity Logging (PPL)

The Intel VROC 6.1 family of products will provide support for Partial Parity Logging (PPL). PPL is used to record the results of XOR'ing old data with old parity. PPL is currently saved as part of the RAID member information and is only utilized when writing RAID 5 parity. It helps protect against data loss when a power failure or a system crash occurs by allowing data to be rebuilt by utilizing the PPL information.

5.2 Intel VROC 6.1 Pre-OS Features and Functionality

This section addresses the usage of the different components and outline the system or platform requirements needed to properly support the usage of the Intel® VROC 6.1 family of products.

5.2.1 Intel VROC 6.1 Unified Extensible Firmware Interface (UEFI) Drivers

The Intel VROC 6.1 family of products provide UEFI drivers to support the UEFI BIOS environment. The UEFI drivers provide an HII interface to support a BIOS level menu-driven configuration tool that is accessed in the BIOS setup. The Intel VROC 6.1 UEFI images included support Intel VMD and SATA/sSATA controllers set to RAID mode.



5.2.1.1 Specification References

This document is not intended to be a go-to document for the UEFI specification. The specification is owned by the UEFI working group and detailed information regarding UEFI can be found in documents published by that organization. The Intel VROC 6.1 UEFI driver implementation conforms to the UEFI specification and is in compliance with version 2.3.1.

Note: At the time that this product was developed, the below specifications were the latest and the versions that were used in the design.

Table 3-1: UEFI Specifications and Location

Specification	Location
UEFI Specification version 2.3.1	http://www.uefi.org/specsandtesttools
UEFI Platform Initialization Specification version 1.2	http://www.uefi.org/specsandtesttools
UEFI Shell Specification version 2.0	http://www.uefi.org/specsandtesttools
UEFI EDK II Build Specification 1.28	https://edk2-docs.gitbooks.io/edk-ii-build-specification/content/#edk-ii-build-specification

5.2.1.2 Intel® VROC 6.1 UEFI User Interface

An **HII-compliant** user interface is provided for the pre-boot configuration of the RAID system. .

- The HII UI is integrated within the UEFI driver binary provided
- Per the UEFI specification, we publish the HII UI as string and forms packages
- The HII UI is accessible from within the UEFI BIOS (How the user accesses it from within the BIOS is OEM-dependent upon implementation)
- The text string 'Intel(R) Virtual RAID on CPU' or 'Intel(R) VROC SATA Controller' or 'Intel(R) VROC sSATA Controller' will be displayed as the selection to enter the HII UI
- Some OEMs may want to hard assign where the Intel VROC UEFI GUI will be located within their BIOS

The Intel VROC 6.1 UEFI Driver FORMSET_GUID is:

FORMSET_GUID { 0xd37bcd57, 0xaba1, 0x44e6, { 0xa9, 0x2c, 0x89, 0x8b, 0x15, 0x8f, 0x2f, 0x59 } }
{D37BCD57-ABA1-44e6-A92C-898B158F2F59}

5.2.2 UEFI System BIOS Requirements for Platform Compatibility with Intel VROC UEFI

This section covers what the OEM/BIOS Vendor is required to accomplish in order to ensure that the platform is compatible with the Intel VROC UEFI driver. See [Relevant Specifications](#) for the information on what the Intel VROC 6.1 UEFI driver was designed to support.



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

The Intel VROC 6.1 UEFI drivers require the following protocols/functions to be provided by the BIOS:

EFI_BOOT_SERVICES:

- LocateHandleBuffer
- OpenProtocol
- CloseProtocol
- WaitForEvent
- HandleProtocol
- FreePool
- AllocatePages
- AllocatePool
- InstallMultipleProtocolInterfaces
- UninstallMultipleProtocolInterfaces
- Stall
- CopyMem
- LocateProtocol

EFI_RUNTIME_SERVICES:

- SetVariable
- GetVariable
- GetTime

Other Protocols:

- EFI_ACPI_TABLE_PROTOCOL (or EFI_ACPI_SUPPORT_PROTOCOL (EDK117))

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

If the OEM plans to use the Intel VROC HII-compliant UI, then **the following protocols/functions are required to be provided by the BIOS:**

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

5.2.2.3 Protocols Provided by the Intel VROC 6.1 UEFI Drivers

The Intel VROC 6.1 UEFI drivers provide the following protocols:

- *Driver Binding Protocol*
- *Component Name Protocol (English only)*
- *Component Name 2 Protocol (English only)*
- *Driver Supported EFI Version Protocol*
- *Device Path Protocol*
- *Config Access Protocol*
- EFI_BLOCK_IO_PROTOCOL
 - For Logical Devices
- EFI_STORAGE_SECURITY_PROTOCOL
 - For Non-RAID disks that support TCG Feature Set
- EFI_EXT_SCSI_PASS_THRU_PROTOCOL:
 - All SCSI commands are supported (for ATAPI devices)



5.2.2.4 How-to-Enable the Platform with Intel VROC UEFI Driver/HII_GUI

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

5.2.2.4.1 Platform UEFI BIOS

1. Ensure that the UEFI System BIOS meets UEFI Specification 2.3.1 compliance
2. The BIOS must provide the following protocols:
 - EFI_Boot_Services Protocols (see section 4.4)
 - EFI_Runtime_Services Protocols (see section 4.5)
 - EFI_HII Protocols (see section 29) Required for the Intel VROC UEFI UI

5.2.2.4.2 Download and Integrate the Intel® VROC 6.1 UEFI Package

Download the latest kit from the Intel VIP (Validation Internet Portal) website.

From the kit select the Pre-OS zip file which will contain all of the Intel VROC and Intel VROC (SATA RAID) Pre-OS images and tools. Go to the `efi_standalone_vroc_rs` (Intel VROC) or `efi_sata` and `efi_ssata` (Intel VROC (SATA RAID)) directories to find the UEFI driver binary file.

1. Select and extract the binary file based on the planned integration method:
 - VMDVROC_1.efi and VMDVROC_2.efi – Include these binaries when building the BIOS image to support Intel VROC.
 - VMDVROC_1.depex and VMDVROC_2.depex – Include these files when building the BIOS image to support Intel VROC.
 - SataDriver.efi – Include this binary when building the BIOS image to support Intel VROC (SATA RAID) Controller
 - SataDriver.depex – Include this file when building the BIOS image to support Intel VROC (SATA RAID) Controller
 - sSATADriver.efi – Include this binary when building the BIOS image to support Intel VROC (SATA RAID) sSATA Controller
 - sSATADriver.depex – Include this file when building the BIOS image to support Intel VROC (SATA RAID) sSATA Controller
2. Use the proper integration tools based on the binary file selected above.

5.2.2.4.3 Verify Compliance

There are tools provided to help verify that this process has been completed successfully. Reference [Intel VROC Pre-OS Components](#) and [Intel VROC \(SATA RAID\) Pre-OS Components](#) for these files and instructions.

5.2.2.5 Known Compatibility Issues with the UEFI Self Certification Test (UEFI SCT) Tool

The following UEFI 2.3.1 SCT tests will appear as FAIL in reports generated using the “Report Generation” tool of the SCT framework. The “Report Generation” tool is the only method that should be used to determine if tests fail. Do not determine test failing test results by viewing the raw log files. The “Report Generation” tool will discard any test results that failed due to an invalid system configuration.



5.2.2.5.1 Bootable Image Support Test\Block IO Protocol Test

EFI_BLOCK_IO_PROTOCOL.Reset - Reset() returns EFI_SUCCESS with ExtendedVerification being TRUE

- *Test Index:* 5.7.5.1.1
- *Test GUID:* 61EE3A34-62A2-4214-B076-5073B177156C
- *Reason:* The Intel® VROC UEFI driver does not support Reset – EFI_UNSUPPORTED is returned.

EFI_BLOCK_IO_PROTOCOL.Reset - Reset() returns EFI_SUCCESS with ExtendedVerification being FALSE

- *Test Index:* 5.7.5.1.2
- *Test GUID:* 98530F3D-8BD8-44A1-9D06-08039FDFEC63
- *Reason:* The Intel® VROC UEFI driver does not support Reset – EFI_UNSUPPORTED is returned.

EFI_BLOCK_IO_PROTOCOL.ReadBlocks - ReadBlocks() returns EFI_SUCCESS with valid parameter

- *Test Index:* 5.7.5.3.1
- *Test GUID:* 9EFE26C2-C565-478A-A0B4-05A8FD2E7E3E
- *Reason:* Test called ReadBlocks() with a BufferSize of 0 so EFI_BAD_BUFFER_SIZE is returned. The UEFI 2.3.1 specification states for ReadBlocks, "The size of the Buffer in bytes. This must be a multiple of the intrinsic block size of the device."

5.2.2.5.2 EXT_SCSI_PASSTHRU_PROTOCOL Test

EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath - call BuildDevicePath with NULL DevicePath.

- *Test Index:* 5.7.8.2.1

5.2.2.5.3 HII Test\HII Config Access Protocol Test

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig - RouteConfig() returns EFI_NOT_FOUND if no target was

- *Test Index:* 5.18.6.2.3
- *Test GUID:* 1F99EBC8-0253-455F-88AC-9E2BA6DCD729 found with the routing data.
- *Reason:* Intel® VROC UEFI driver does not support RouteConfig – EFI_UNSUPPORTED is returned. RouteConfig is not supported so that Intel® VROC HII form values are only modified by the Intel® VROC drivers.

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig - RouteConfig() returns EFI_INVALID_PARAMETER with Configuration been NULL

- *Test Index:* 5.18.6.2.1
- *Test GUID:* 495C99F3-0231-45A5-AFFA-D25C6F9A191C
- *Reason:* is caused by not using EFI HII Configuration Access Protocol (see section 31.4 in UEFI 2.4 specification) in Intel VROC UEFI Drivers The Intel VROC software does not use it as it is not necessary.

HII_CONFIG_ACCESS_PROTOCOL.RouteConfig- RouteConfig() returns EFI_SUCCESS with valid parameters

- *Test Index:* 5.18.6.2.4
- *Test GUID:* 1A15DF85-6CC1-43F2-9B86-218BD5FDF4A0
- *Reason:* is caused by not using EFI HII Configuration Access Protocol (see section 31.4 in UEFI 2.4 specification) in Intel VROC UEFI Drivers. The Intel VROC software does not use it as it is not necessary.



5.2.3 UEFI Health Protocol

The Intel VROC 6.1 family of products supports the UEFI Health Protocol as defined in the UEFI specifications listed above. Intel VROC 6.1 provides support for `EFI_DRIVER_HEALTH_PROTOCOL.GetHealthStatus()`. The types of Health status messages that Intel VROC 6.1 will report, will match those messages that are displayed in the Intel VROC 6.1 UEFI HII. A health status of failed will be returned when:

- A RAID volume is in a failed state
- A RAID volume is in a degraded state
- A drive has encountered a SMART event
- A drive is marked as:
 - Failed
 - Unknown
 - Unsupported
 - Incompatible
 - Offline

Note: Support for `EFI_DRIVER_HEALTH_PROTOCOL.Repair ()` is currently not supported.

5.2.4 Determining the Version of the Intel® VROC UEFI Driver

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

5.2.4.1 Using the UEFI Shell

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

```
Shell:>drivers
```

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

```
"CD 0000000B B -- 1 2 Intel(R) VROC xx.x.x.xxxx SATA Driver"
```

5.2.4.2 Intel VROC (SATA RAID) Device Information Display in UEFI

The Intel VROC 6.1 product release package includes the support for the Intel VROC (SATA RAID) UEFI driver to report the physical port for devices connected to the SATA/sSATA controllers. In some of the releases of the Intel VROC UEFI driver, the device information provided (i.e. to the UEFI shell environment) was based off of the enumeration values created during the discovery of the devices attached. To support a manufacturing environment that relies on this information to identify the physical port the device is connected too, the UEFI driver now reports out the physical port value instead of the enumerated value. The data is displayed as a set of 3 values in the following order: X-Y-Z

X – 0: passthru disk, 1: volume

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

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



5.2.4.3 Intel VROC (SATA RAID) UEFI Boot Volume Configuration

This example shows how to configure a bootable RAID volume using PCH with SATA drives for UEFI. This procedure should only be used for a newly-built system or reinstall of the operating system.

To enter the setup and configuration mode of the BIOS, do the following:

1. Press the key indicated for entering setup and select boot options when the menu appears.
 2. Navigate to Intel VROC (SATA RAID) Controller, and press <Enter>.
 3. Highlight Create RAID Volume, and press <Enter>.
 4. Highlight the Volume Name field.
If the user wishes to change this value from the default, press <Enter>. Modify the name to the preferred name, and press <Enter> to save the new name.
 5. Highlight the RAID Level. Default setting is RAID0.
 6. To change this to another RAID configuration, press <Enter> and select desired RAID level, then press <Enter> to save the user's selection.
 7. To select the drives that will be part of the RAID, highlight the < > next to the desired drive and press <Enter>.
A small menu will pop up with the setting currently at a value of blank, indicating not enabled. To include this drive within the RAID, highlight the X and press <Enter> to enable this drive as part of the array.
 8. Repeat Step 7 with each drive needed until all desired drives have been included within the RAID.
 9. Unless the user is setting up a RAID 1, the user will be allowed to adjust the Strip Size for the configuration. RAID 1 will be set at 128KB, and cannot be adjusted. To modify this setting, highlight the present value, press <Enter>, and select the preferred value. Press <Enter> to save the new strip size.
 10. The capacity of a RAID Volume will initially display the maximum value allowed. This can be modified to a smaller value to create additional RAID volumes. If the user wishes to alter this value, select the Capacity and press <Enter>. Enter the desired value up to the maximum and press enter to save. The value will not change if it exceeds allowed capacity.
 11. To finalize creation of the RAID volume, highlight Create Volume, and press <Enter>.
- Note:** At this point in the process if there are any significant discrepancies between the drives selected to be RAID members, a warning message will be displayed if one of the following conditions is encountered:
- There is a combination of SSDs and HDDs used
 - There are at least two drives that have a size difference of more than 10%.
12. The user will see the created volume on the following screen.
 13. To save the configuration, press <Esc> and the user will be presented with the following message.
"Changes have not been saved. Save Changes and exit? Press 'Y' to save and exit, 'N' to discard and exit, and 'ESC' to cancel." Press <y> to save and exit.

WARNING: This step is extremely important, failure to save at this point will cause the user's RAID volume will not be saved. The user must save in order for the RAID process to properly complete.



After saving the user's RAID volume, the system will require a reboot before the RAID volume may be used for a boot or data volume. The user may finish other configuration items now, but if the user is ready to proceed with OS installation, this is the time to perform a reboot.

Once the system reboot has been completed, the user may format the RAID volume and begin system installation.

5.2.5 Pre-OS Installation of the Intel® VROC Driver

The Intel VROC (SATA RAID) and Intel VROC (VMD NVMe RAID) drivers can, and should, be used to install the Windows* OS when the platform is configured for RAID mode. For conditions where the Windows* OS is already installed, the Intel VROC drivers can be loaded/installed from the OS environment.

5.2.5.1 Pre-OS Driver Installation Using the “Load Driver” Method

1. During the Operating system installation, after selecting the location to install the Windows* OS click on 'Load Driver' button to install the desired Intel VROC drivers.
2. When prompted, insert the media with the Intel VROC driver files and press Enter.
3. The user can find the media and browse to the folder where the files are located.
4. Follow the steps to load the driver and return the installation.
5. Continue the installation.

5.2.5.2 Selectable Boot Volume

Intel VROC 6.1 family of products will support the ability to select any volume as the OS boot volume, provided the volume size meets the minimum size requirements for the OS being installed. The OS installer will be able to install the operating system onto the specified RAID volume.

5.3 Intel VROC 6.1 Monitoring and Management

This section outlines those components that are used for installing and managing the Intel VROC 6.1 products.

5.3.1 Intel VROC 6.1 GUI

The Intel VROC 6.1 GUI is a Windows* based application that provides users monitoring and management capabilities for the Intel VROC 6.1 storage subsystem. It offers a wide range of monitoring and management activities for the Intel VROC 6.1 based RAID subsystem.

Note: For the SATA Controller in AHCI mode, there is no management or monitoring capabilities offered in the GUI.

The RAID subsystem is defined as drives attached to the Intel VM controller or the PCH controller in RAID mode.

The Intel VROC 6.1 GUI requires (at a minimum) Microsoft* .NET 3.5.



5.3.1.1 Intel VROC 6.1 Installation

The Intel VROC 6.1 GUI, drivers and supporting infrastructure are distributed through a single installation executable.

When the Intel VROC 6.1 installer is executed it will perform the following tasks:

- If it detects that there is a newer version of the Intel VROC (VMD NVMe RAID)/Intel VROC (Legacy NVMe RAID) or Intel VROC (SATA RAID) solutions already installed it will display a message requesting permission from the user to overwrite the newer version of the software.
- It will overwrite older versions of the installed software components.
- It will install the driver in accordance with Microsoft's PnP device detection method.
- It will create the following shortcuts within the start menu path:
start->programs->Intel -> Intel(R) Virtual RAID on CPU
- It will provide the user a mechanism to list all available installation options.
- It will provide an option to specify a custom install path. If the specified install path is illegal the Intel VROC 6.1 installer will cancel installation.
- It will implement an option to store the steps of an installation to a log file. The default location is %temp%\Intel\Logs. The file name is IntelVROC.txt. If the installer sees that the file already exists, the log information is appended to the existing file.
- It will create a cached installer used for uninstalling. The location of this uninstaller is %programdata%\Package Cache\<some guid>. The exact path to the cached installer is stored in the product's Uninstall registry key. This registry key is HKLM\Software\Microsoft\Windows\CurrentVersion\Uninstall\<some guid> (in the 32-bit hive) where the Display Name value is the Intel VROC (SATA RAID)'s display name. Also the BundleCachePath value contains the full exact path to the cached installer. If the BundleCachePath value does not exist, the wrong registry key is being looked at, because there should be 2 keys, if on a 32-bit system.
- It will implement an option to install the GUI components only. When invoked with this option the Intel VROC 6.1 installer software will install the UI; Rapid Storage Technology enterprise Help UI System Service and Notification Area Icon.

The product installer shall install the following components:

- Intel VROC (VMD NVMe RAID) drivers for Intel VMD/Intel VROC (NonVMD NVMe RAID)
- Intel VROC (SATA RAID) driver for SATA
- Intel VROC (SATA RAID) driver for sSATA
- Intel VROC 6.1 System Service
- Intel VROC 6.1 Notification Area Icon

5.3.1.1.1 Default Installation Path

The Intel VROC 6.1 installer software will install 'Intel VROC Solution' to a default folder unless specified otherwise.

The default path shall be C:\Program Files (x86)\Intel\Intel(R) Virtual RAID on CPU

5.3.1.2 Intel VROC (SATA RAID) Driver

The Intel VROC 6.1 installer will install separate driver binaries for the SATA and sSATA controllers.



5.3.1.3 Intel VROC (VMD NVMe RAID) Driver

The Intel VROC 6.1 installer will install a single Intel VROC (VMD NVMe RAID) driver binary for all VMD devices enabled in the system.

5.3.1.4 Intel VROC (NonVMD NVMe RAID) Driver

The Intel VROC 6.1 installer will install a single Intel VROC (NonVMD NVMe RAID) driver binary for Intel NVMe SSD found in a supported platform (NonVMD capable).

5.3.1.5 Intel ASM Installation Checkbox

The Intel VROC 6.1 Product installer will provide an option (checkbox) to install Intel ASM (only supported on VMD capable platforms).

5.3.1.6 Intel ASM Installation Command Line Option

The Intel VROC 6.1 installer shall provide an option (command-line) to install the Intel ASM (only supported on VMD capable platforms).

5.3.2 Getting Started

The Intel VROC 6.1 software package provides high-performance Intel VROC package that provides support for:

- NVMe drives attached to the Intel VMD controller (on VMD capable platforms)
- Intel NVMe SSDs attached to the CPU (on NonVMD capable platforms)
- SATA drives attached to the platform PCH in RAID mode

5.3.2.1 RAID Enabled Systems

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

5.3.2.2 AHCI Enable 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 the user's computer. The Intel VROC 6.1 product will contains support for platforms that will have the PCH controller in AHCI mode. The following Intel VROC (SATA RAID) AHCI mode driver features are supported on AHCI-enabled systems:

- Native command queuing
- Hot plug
- Disks of more than two terabytes (if that size is supported by the Intel VROC UEFI pre-OS driver)

The Intel VROC (SATA RAID) AHCI mode product will also be able to support a platforms supported power management states (i.e. Power-on, Restart, Sleep, Hibernate, and OS installation) on drives attached to the controller in this configuration.

Note: Installing with this option will overwrite any driver previously installed. Intel recommends using the Microsoft* AHCI that is included with each OS release.



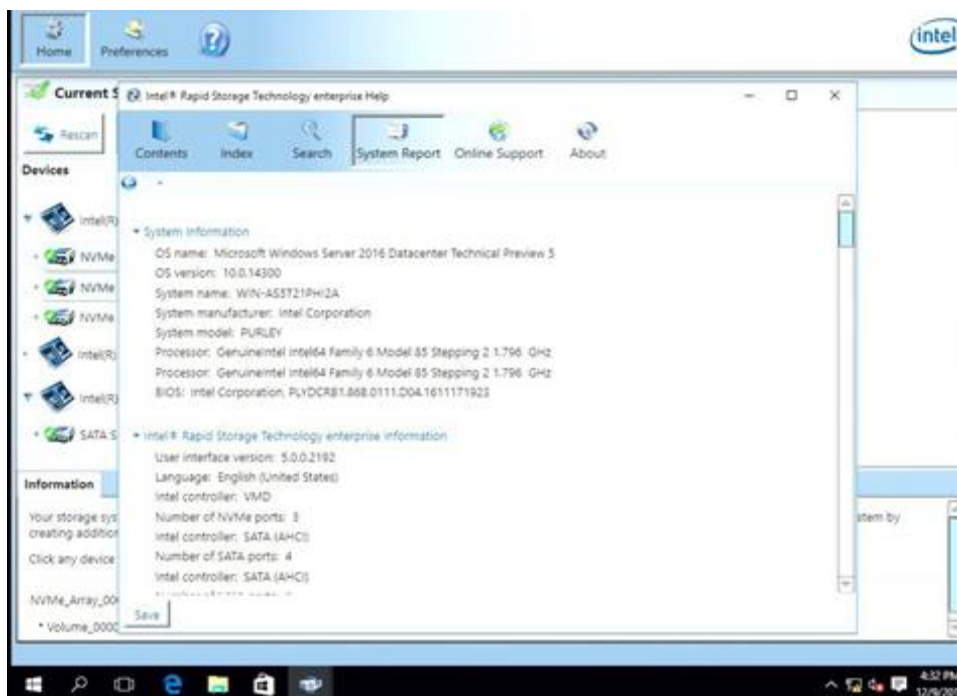
5.3.3 Component Version Detection

The Intel VROC 6.1 GUI at launch will determine the version of the Pre-OS environment and driver being used by the system. The version information of all components will be compared and the lowest common denominator used to determine the features visible in the GUI.

5.3.4 Using System Report

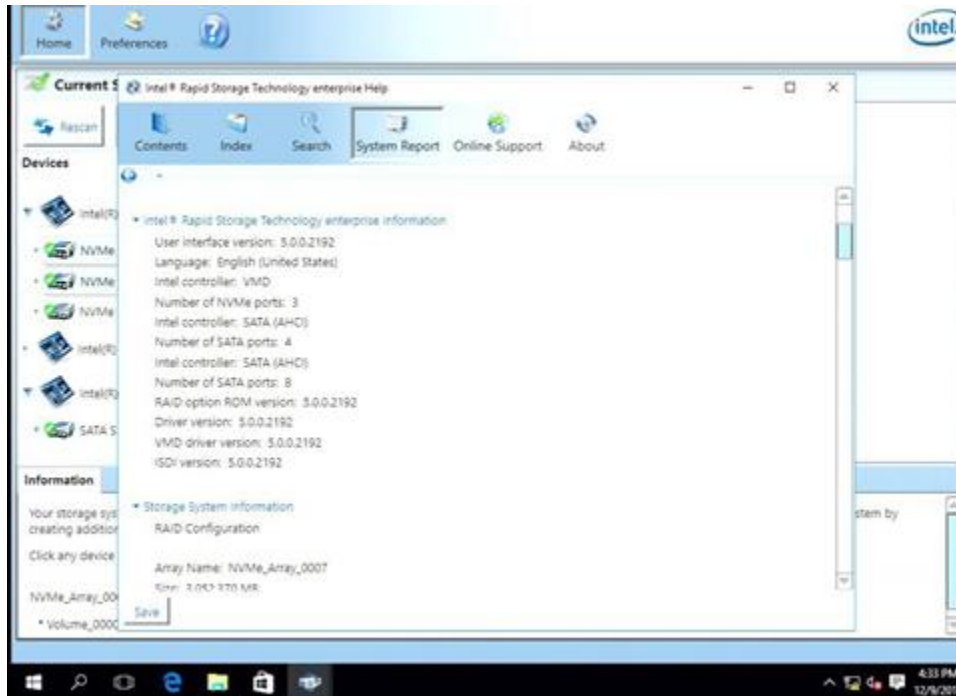
The System Report from the Intel VROC GUI will provide the detailed information about the system. It contains platform information, VROC component information and information on all the attached devices. To obtain this information open the Intel VROC GUI, and select the "?" button to open the Help window.

Select the **System Report** button.



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

To view the versioning information for the VROC components, scroll down to make all entries visible.



5.3.5 Notification Area





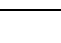
The notification area (also called the system tray) is located on the user's desktop. The taskbar contains the notification area icon for Intel VROC 6.1 product. The icon provides storage system status and notifications such as volume and disk events based on a change of state.

The notification area icon will automatically display in the notification area once the Intel VROC 6.1 package is installed. Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Settings changes are applied on a per user basis, and do not affect other users' settings.

5.3.5.1 Opening the Application from the Notification Area

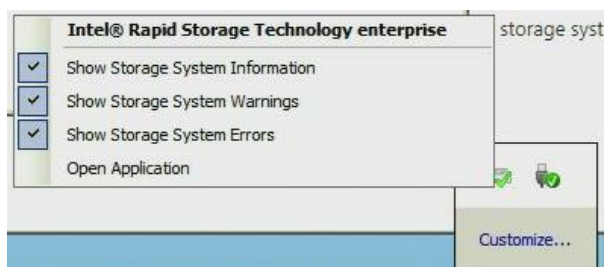
- Right-click the icon.
- Click 'Open Application'.

The notification area icon can be in the following states:

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

5.3.5.1.1 Selecting System Notifications

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



Note: To hide the notification area icon, deselect 'Show the notification area icon' under System Preferences.

5.3.5.1.2 Reviewing Notifications

Hover over the icon at any time to view the storage system status or the progression of an operation. A pop-up window will notify the user 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.



5.3.5.2 Understanding the Application

The Intel VROC 6.1 application allows the user to optimize and maintain a healthy storage system by creating volumes, customizing performance settings and managing storage system elements. This section provides the user with a general overview of the storage system configuration and an individual review of all the areas contained in this application.

5.3.5.2.1 Storage System Configuration

The Intel VROC 6.1 storage system combines hardware capabilities with Intel VROC 6.1 family of products to provide flexible data storage units on the user's computer. Each data storage unit, or RAID configuration, consists of three elements that include:

1. The physical disk:
 - a. NVMe SSDs
 - b. SATA disks
2. The RAID array
3. The RAID volume

When at least one RAID volume is present on the system, these elements are represented in the storage system. The drives on controller information is represented in the “Devices” pane to the left. The arrays and volumes are represented in the “Volumes Properties” pane in the center. The properties pane to the right show the properties information of what may be selected (i.e. Volume_0000).






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 NVMe or 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 the user used to create that volume form an array. Refer to the Creating Additional Volumes topic for details on how the user 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 the user created a volume for data protection, then the user's storage system may include a RAID 1 volume spanning two NVMe or SATA disks, which mirrors data on each disk.
- **Disks**
A disk (i.e., hard disk, solid state drive or Intel NVMe SSD) physically stores data and allows read/write data access. If a disk is used to create a volume, it becomes an array disk because it has been grouped with other disks to form an array.
The storage system can also include ATAPI devices, which cannot be used to create a volume. They are a mass storage device with a parallel interface, such as CD-ROM, DVD/Blu-ray disc, or tape drive.
When selecting drives to be RAID volume members, a warning message will be displayed if there are any significant discrepancies between two or more selected drives. The warning message will be displayed if one of the following conditions is encountered:
 - There is a combination of SSDs and HDDs used
 - There are at least two drives that have a size difference of more than 10%.

5.3.5.2.2 Navigation


The Intel VROC 6.1 application is organized into three main areas depicted by the top navigation buttons: Home, Preferences and Help.

Note: Depending on the user's computer's configuration and available hardware, Create and Accelerate may not be available.



Home

Home is the main area of operation, where the user can select to "Scan" for hardware changes or "Create Volume."



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

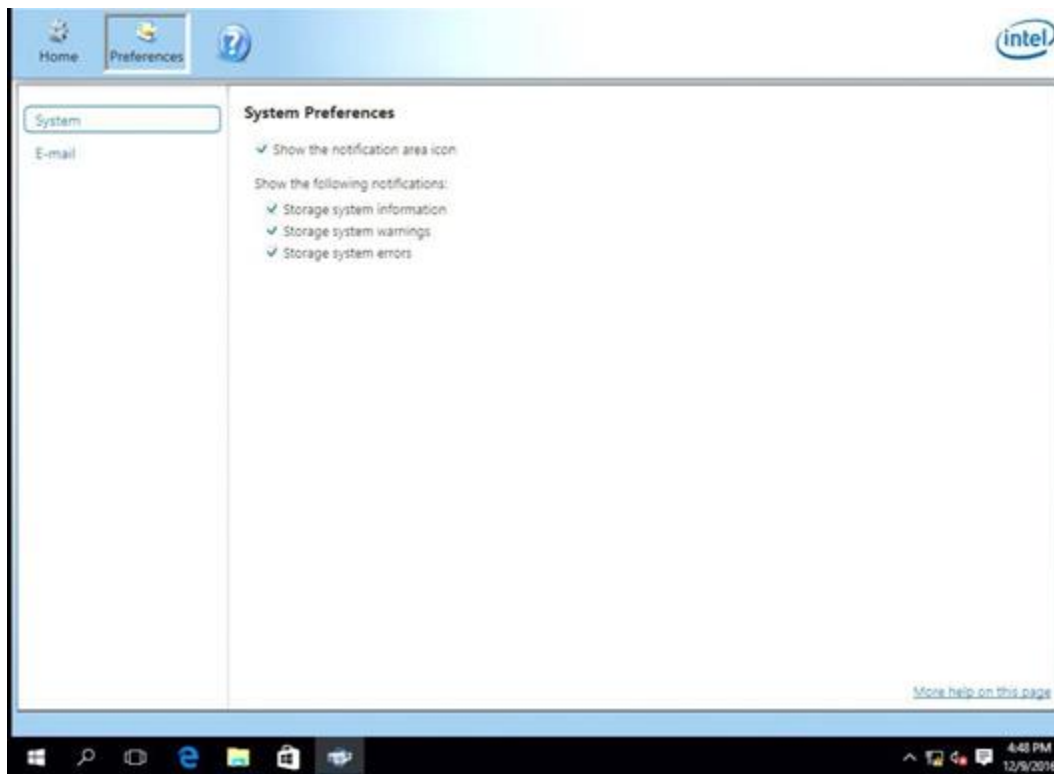


Information

Your storage system is configured for data protection, increased performance and optimal data storage capacity. You can further optimize your storage system by creating additional volumes. To begin the process, click 'Create Volume...'.
Click any device or volume to display its properties.

Preferences

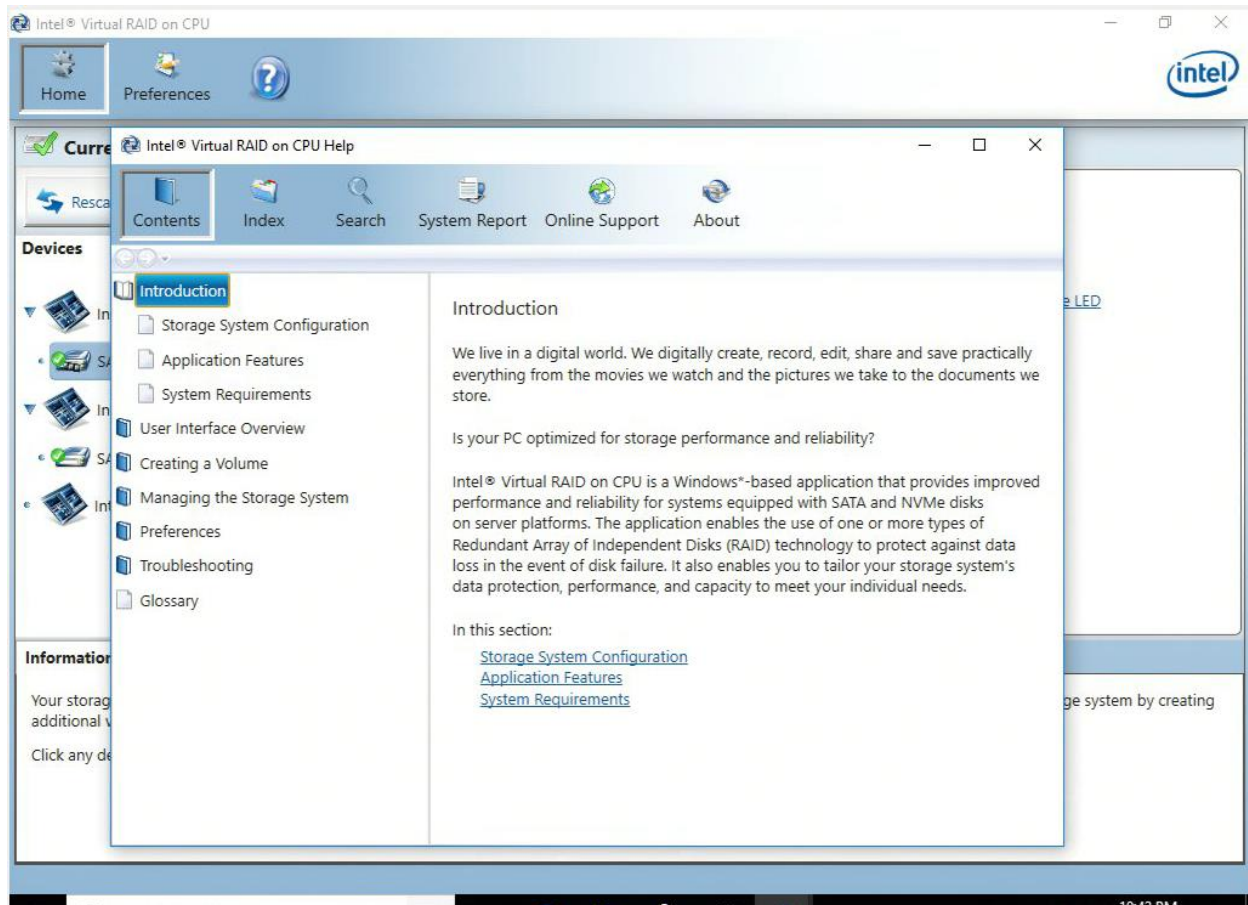
The 'Preferences' area allows the user to customize system settings by enabling the display of the notification area icon, and by selecting the type of notifications that the user wants the application to display.





Help

The "Help" area provides the detailed information on how to use the Intel VROC 6.1 GUI as well as key system information.



5.3.5.3 Storage System Status







Anytime the Intel VROC 6.1 GUI is launched, the application opens to the 'Status' area. This is where the general state and health of the user's 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 the user's storage system.



5.3.5.3.1 Understanding the Status

To get the full benefits of what Intel VROC 6.1 GUI has to offer, it is critical to maintain a healthy storage system. The application helps the user track and reports any disk or volume related problems that could put the safekeeping of the user's data at risk.

The storage system can be in the following states:

Icon	Description
	Reports that the system is functioning as expected, supported 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, the user 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 the user rectify any storage system conditions.
	Reports that storage system data may be at risk due to a problem detected on one or more supported disks.
	The Manage subsection displays any disk or volume states reported by the storage system that may require the user's attention in order to keep data fully protected and accessible. Details or a recommended action are provided to help the user fix any storage system problems. For example, if the master disk in a recovery volume is reported as failed, we would recommend that the user rebuild the volume to another disk. Note: In this state, we recommend that the user backup any accessible data before taking action.
	The Manage subsection displays any disk or volume states reported by the storage system that require the user's immediate attention in order to keep data fully protected and accessible. Details or a recommended action are provided to help the user 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 the user rebuild the volume to another disk. Note: In this state, we recommend that the user backup any accessible data before taking action.



5.3.5.4 Volume States

Volume Type	Normal	Degraded	Failed
		Refer to Troubleshooting Degraded Volumes and Caching Issues for more information.	Refer to Troubleshooting Failed Volumes and Caching Issues for more information.
RAID 0		Not applicable	
Single-disk (cache)			
Single-disk (data)		Not applicable	
RAID 1			
RAID 5			
RAID 10			

5.3.5.5 RAID-Ready Mode

A RAID-Ready system refers to a system that has been configured to support the Intel VROC 6.1 family of products. The system BIOS has the appropriate pre-boot drivers and has been configured for RAID mode. Intel VROC 6.1 will allow a pass-through drive that contains DATA or a System OS to become a RAID volume. This can be done while the system remains operational.

5.3.5.6 RAID Ready Migration

- Pass-through to 2-XX drive RAID0 (*Intel VROC (VMD NVMe RAID)/Intel VROC (NonVMD NVMe RAID)*)
- Pass-through to 2-8 drive RAID0 (*Intel VROC (SATA RAID)*)
- Pass-through to a 2 drive RAID1
- Pass-through to a 3 to XX drive RAID5 (*Intel VROC (VMD NVMe RAID)/Intel VROC (NonVMD NVMe RAID)*)
- Pass-through to a 3 to 8 drive RAID 5 (*Intel VROC (SATA RAID)*)
- Pass-through to a 4 drive RAID 10

5.3.5.7 Intel VROC Graphical User Interface (GUI)

The Intel VROC 6.1 will provide support for graphical user interfaces for management of RAID arrays and volumes on NVMe SSDs directly connected via PCI express slots controlled/managed by the VMD enabled ports. This includes a GUI in the Operating System as well as a web-based GUI for remote management, Intel ASM.

5.3.5.7.1 Intel VROC Pass-Thru Mode GUI Support

The Intel VROC GUI will provide support for LED Locate and Managed Hot Plug to be selected by the user when VMD is enabled but the Intel VROC RAID Upgrade Hardware Key is absent from the platform. RAID functionality is disabled in this configuration except for RAID 0 using Intel PXX08 NVMe SSDs.

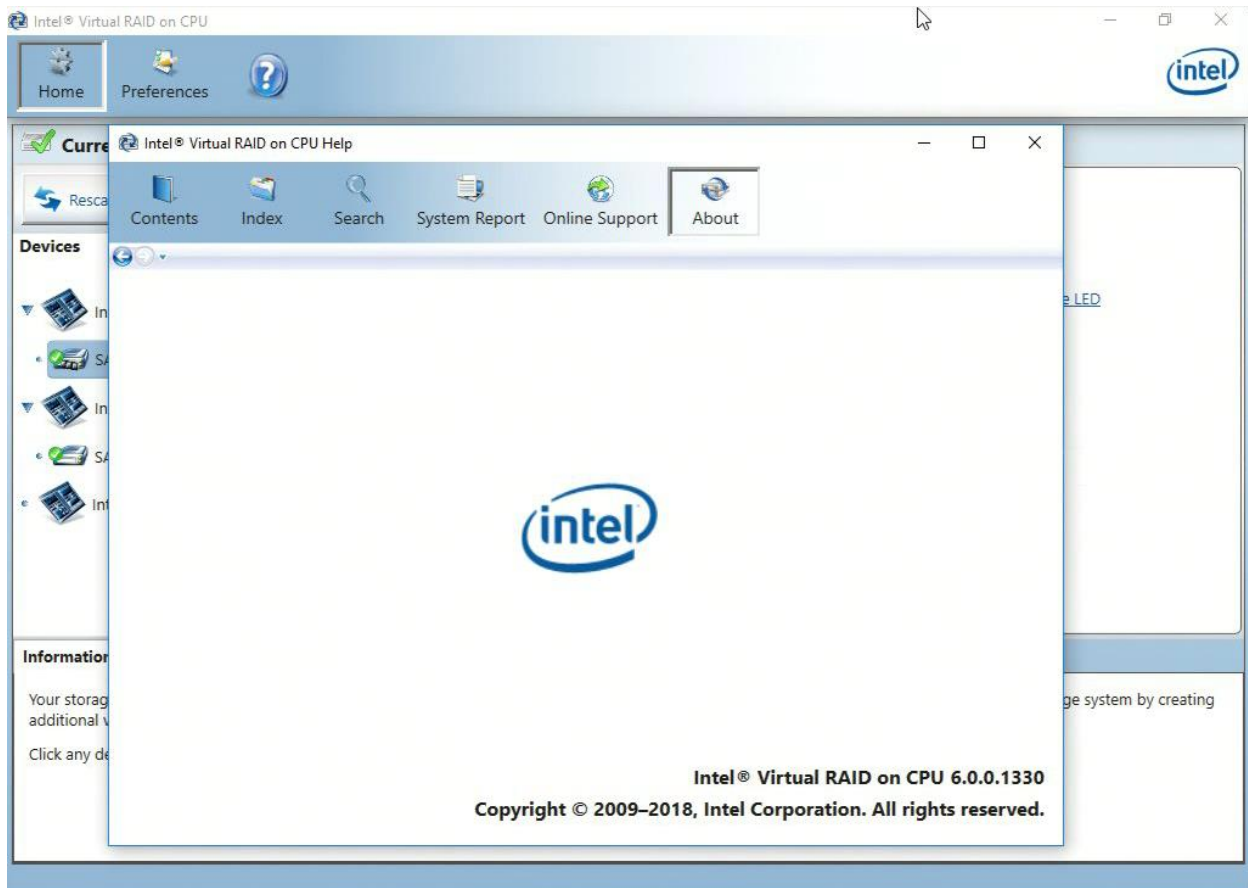


5.3.5.7.2 Intel VROC Enabled Mode GUI Support

The Intel VROC GUI will provide support for RAID managed, LED Locate and Managed Hot Plug to be selected by the user when Intel VMD is enabled and an Intel VROC RAID Upgrade Hardware Key is present in the platform.

5.3.5.8 Intel® VROC Graphical User Interface (GUI) Version

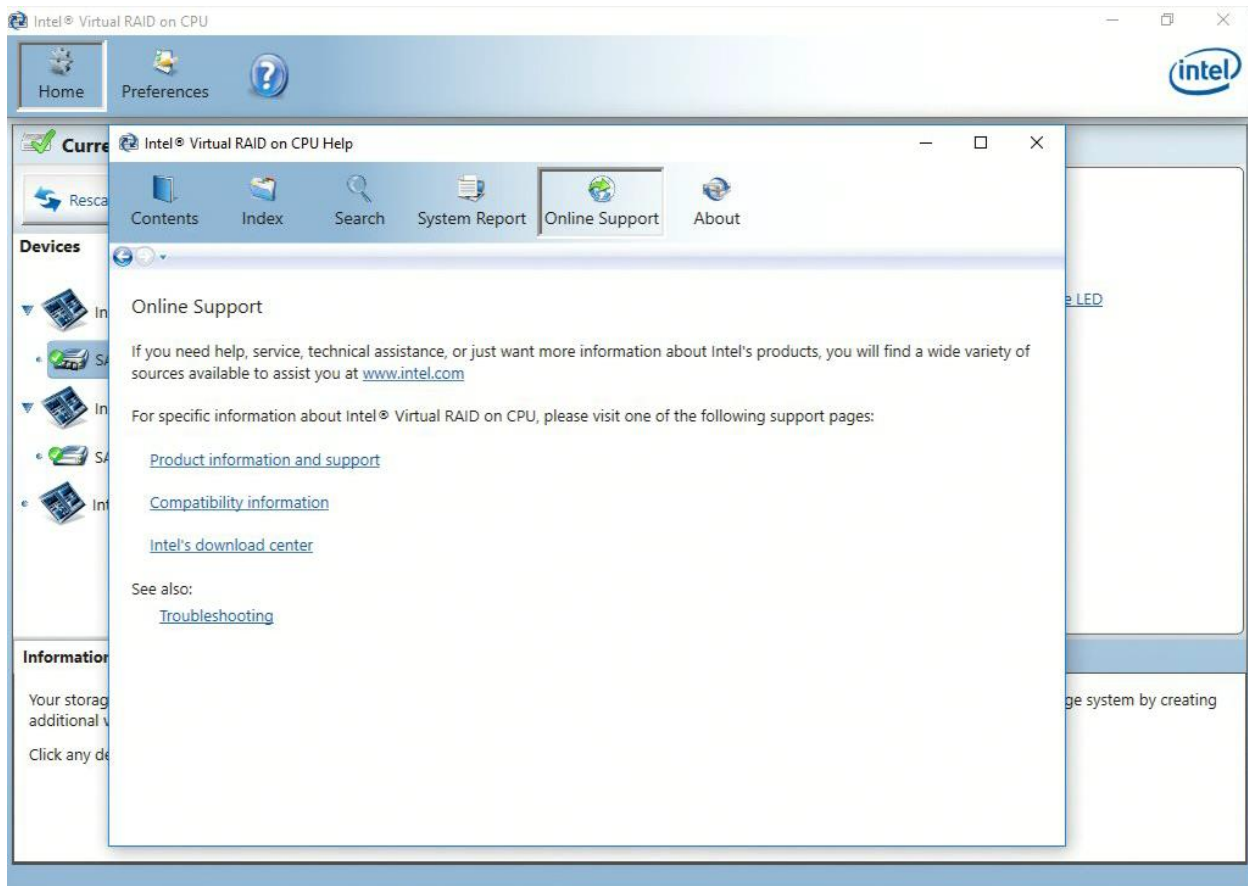
The specific version of the Intel VROC GUI is obtained from the Help window by clicking on “About”.





5.3.5.9 Intel VROC 6.1 GUI Online Support

The Intel VROC 6.1 GUI [Help] Menu, Submenu [Online Support] when selected will display a pop-up window with the support URLs as shown in the figure below:



5.3.6 Managing Arrays

The user must be logged in as administrator to perform the actions listed in this section.

The user can manage arrays by clicking a selected array in the storage system view under "Volume" name. This allows the user to review the properties and access all actions associated with that array, such as adding a disk or increasing a volume size.



5.3.6.1 Array Properties

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

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

5.3.6.2 Adding a Disk to an Array

The user can add one or more disks to an existing array to increase the system storage capacity. This feature can be useful if the user wants 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 disks are connected to the computer and available.
- The disk to be added to an existing array must meet the following criteria:
 - a. It must be on the same controller as the existing array (SATA, sSATA or Intel VMD). This cannot be used to cross controllers (SATA to sSATA or SATA to Intel VMD).
 - b. It must be at least the same size as the “smallest” drive in the existing array. A drive that is smaller (then the smallest drive in the array) cannot be added.

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

If the user performs a driver upgrade or downgrade while the data migration is in progress and then restarts the computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the user will not be able to restart the system until the migration is completed otherwise the operating system cannot load. If the user is migrating a data volume when a reset occurs, the user will need to reverse (roll back) to the previous driver used, and then restart the computer to return to a normal state.

This action can also be performed from Manage Volume.

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

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



5.3.6.3 Creating a Volume

Creating an Intel VROC RAID volume can be performed on supported platforms in the following configurations:

- NVMe SSDs that are managed by the Intel VMD or
- Intel NVMe SSDs attached to a CPU (that does not contain the Intel VMD)
- SATA drives attached to the SATA/sSATA controller (configured in RAID mode)

This can be accomplished by selecting a volume type under 'Create'. We recommend the user become familiar with the minimum requirements in this section before starting the volume creation process.

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

5.3.6.3.1 Volume Requirements

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

1. The user is logged on as an administrator.
2. 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).
3. Two or more supported disks, including the operating system disk are connected, and in a normal state, and unlocked (only applies to password-protected disks).
4. Each of the disks that are to be part of the RAID volume may not have any SMART events.

Enabling More Disks

When configuring a volume, the application will only list the disks that meet the requirements listed below. For example, a locked disk connected to the user's 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.

1. If the first selection is a system disk, any additional disks selected must be of equal or greater size to ensure that all the system files are migrated to the new volume.
2. If the first selection is a non-system disk, and a system disk is then selected, the former must be of equal or smaller size to ensure that all the system files are migrated to the new volume.
3. A system volume can be greater than 2TB. If the user's first selection is a system disk, the total size of the other disks shall not allow the volume size to exceed 2TB. Exception: If the user is creating a volume using disks that have no existing data, and the user's operating system is a 64-bit Edition, the application will allow a volume to be greater than 2TB.
4. The disks used to create a volume must have the same type of connection (NVMe/SATA).

5.3.6.3.2 Creation Process

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

Selecting a Volume Type

Before the user can create a volume, a decision on how to enhance the storage system based on the user's needs is required. Depending on the available hardware, the user may have the option to combine volume types by creating

more than one volume on a single array. Refer to the Creating Additional Volumes for more information on this type of configuration. Below is an overview of the five volume types that the user 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, the user can create a volume to protect data or optimize disk performance by clicking one of the two options listed in the Create subsection. When choosing this option, the application automatically configures the volume using the only two disks available and assigns default settings. Refer to the applicable procedure described in "Completing the Process" for details.

Creating a Custom Volume

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

Note: When selecting drives to be RAID volume members, a warning message will be displayed if there are any significant discrepancies between two or more selected drives. The warning message will be displayed if one of the following conditions is encountered:

- There is a combination of SSDs and HDDs used
- There are at least two drives that have a size difference of more than 10%.



RAID 1 : Real-time data protection

Combines two disks where each disk stores an exact copy of the data to appear as a mirror of each other.

Disks required	2
Advantage	Full data redundancy and excellent fault-tolerance; increased read transfer rate.
Disadvantage	Storage capacity is only as large as the smallest disk; slight decrease in write transfer rate.
Application	Typically used in workstations and servers to store critical data. Available in specific mobile configurations.

**RAID 0 : Optimized disk performance**

Combines two or more disks and breaks down data into units that are spread across the array disks.

Disks required

- SATA controllers: 2 to maximum number of ports in a single SATA Controller
- Intel VMD controllers: 2 to 24

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.

**RAID 5 : Efficient data hosting and protection**

Combines three or more disks where data and parity are striped across the array disks in a rotating sequence. Parity is a mathematical method for recreating lost data to a single disk.

Disks required

- SATA controllers: 3 to maximum number of ports in a single SATA Controller
- Intel VMD controllers: 3 to 24

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 and Internet and Intranet servers.

**RAID 10 Balanced performance and data protection**

Combines four disks to create a combination of RAID types 0 and 1. The data is striped across a two-disk array, forming a RAID 0 component. Each disk in the RAID 0 array is mirrored by a disk in the RAID 1 array, storing an exact copy of all the data.

Disks required

4

Advantage

Combines the read performance of RAID 0 with the fault-tolerance of RAID 1, resulting in increased data access, 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.



Configuring the Volume

Once the volume type is selected, the user is ready to configure the RAID volume.

RAID Volume

1. Type a new volume name if the default name change is desired.
2. Select the required number of disks.
3. Select the disk from which to keep the data, if desired.

The user can only keep data from one disk. If data is required to be kept from more than one disk, the user 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:

1. Select the array allocation by using the slider.
2. Select a data strip size.
3. Enable or disable the volume write-back cache.
4. Select the check box to initialize the volume. The user can choose to perform this action at a later time.

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

Completing the Process

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

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

Creating a Custom Volume

WARNING: The user can only keep existing data from one of the disks selected to create a volume. It is recommend that a backup of all valuable data be done before proceeding.

If the user performs a driver upgrade or downgrade while the data migration is in progress and then restarts the computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the system will not be able to be restarted until the migration is completed otherwise the operating system cannot load. If the user is migrating a data volume, the user will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

1. Under 'Confirm' review the selected configuration.
2. Click 'Create Volume' if the user wants to create the volume using the selected configuration. The process starts immediately.
3. Click 'OK' to confirm.
Once completed, a dialog displays to notify the user that the volume was successfully created.
4. Click 'OK' to close the dialog.



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

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

1. Once the migration status reports 100% complete, restart the user's computer for the operating system to recognize the new volume size.
2. Create a new partition or extend the existing partition to utilize the new volume space using Windows Disk Management*.

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

5.3.6.3.3 Creating Additional Volumes

Creating multiple volumes on a single array

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

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

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

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

**Supported RAID Volume Combinations on a Single Array**

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

Creating Additional Volumes on a New Array

The user 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 the user wants to keep data, if desired.
The user can only keep data from one disk. If the user wants to keep data from more than one disk, back up all valuable data prior to creating a volume is required.
7. Make any necessary changes in the "Advanced" section.
8. Review the selected configuration.
Click 'Back' or an option in the left pane if the user wants to make changes.
9. Click 'Next'.
10. Click 'Finish' to start the creation process.



RAID Volume Creation with Data Preservation

Intel VROC 6.1 will support the ability to preserve the data from one of the disks used for the volume creation. A non-RAID disk can be migrated to a RAID volume while retaining the existing data on that disk.

Note: When creating a system boot volume (the Intel VROC 6.1 family of products only), the maximum strip size supported is 128K. The system will also use the entire Array space for the volume, meaning that Disk Coercion is not used in this case.

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

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

- Individual pass-through to 2 - 8 drives for RAID 0 (*Intel VROC (SATA RAID)*)
- Individual pass-through to 2 - 6 drives for RAID 0 (*Intel VROC (SATA RAID)*)
- Individual pass-through to 2 - 48 drives for RAID 0 (*Intel VROC (VMD NVMe RAID)/Intel VROC (NonVMD NVMe RAID)*)
- Individual pass-through to 2 drive RAID 1
- Individual pass-through to 4 drive RAID 10
- Individual pass-through to 3 to 8 drive RAID 5 (*Intel VROC (SATA RAID)*)
- Individual pass-through to 3 to 6 drive RAID 5 (*Intel VROC (SATA RAID)*)
- Individual pass-through to 3 to 48 NVMe drive RAID 5 (*Intel VROC (VMD NVMe RAID)/Intel VROC (NonVMD NVMe RAID)*)

5.3.6.4 Increasing Volume Size

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

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

This option is only available if:

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

WARNING: If the user performs a driver upgrade or downgrade while the data migration is in progress and then restarts the user's computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the user will not be able to restart the system until the migration is completed. Otherwise the operating system cannot load. If the user is migrating a data volume when a system restart occurs, the user will have to reverse (roll back) that the last driver version, and then restart the computer to return to a normal state.



Increasing the volume size from Manage Array

- Under 'Home', in the storage system view, click the array the user wants to manage. The array properties are now displayed on the right.
- Click 'Increase size' next to the volume name.

If more than one volume is present on a single array, the user will need to increase the size of each volume, one at a time.

- Click 'Yes' to confirm.

CAUTION: Once the data migration starts, the operation cannot be canceled.

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

Increasing the volume size from Manage Volume

- Under 'Home', in the storage system view, click the volume whose size the user wants to increase. The volume properties are now displayed on the right.
- Click 'Increase size'.
- Click 'Yes' to confirm.

CAUTION: Once the data migration starts, the operation cannot be canceled.

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

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

5.3.6.5 Enabling Disk Data Cache

Enabling the disk data cache for all disks on the array allows the user to enable cache memory physically present on the disks and use it to speed up data access. This action is only available by selecting the Array in the 'Volumes' pane. The array must be selected because the data cache must be in the same state across all disks that are part of a single array.

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

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

- Under 'Home', in the storage system view, click the array the user wants to manage. The element properties are now displayed on the right.
- In the Advanced section, click 'Enable' or 'Disable' depending on the option available.
- Click 'Yes' to confirm.
- The page refreshes and now displays the new setting.



5.3.6.6 Managing Volumes

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

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

5.3.6.7 Volume Properties

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

RAID Volume Status Table

Status	Description
Normal	Indicates that volume data is fully accessible.
Locked	Indicates that at least one array disk is locked with a password. The volume is visible because at least one other array disk is unlocked. Refer to Unlocking Password-Protected Disks for instructions on unlocking disks.
Degraded	Indicates that one array disk is missing or has failed. A RAID 0 volume cannot be in this state because of the striping configuration.
Failed	1. RAID 0 volume: indicates that one or more array disks are missing or have failed. 2. RAID 1 volume: indicates that both array disks are missing or have failed. 3. RAID 5 or 10 volume: indicates that two or more array disks are missing or have failed.
Incompatible	Indicates that the volume was moved to another system that does not support the volume type and configuration.
Inaccessible	Indicates that data on the accelerated volume cannot be accessed because it is missing, or that the accelerated volume data is not synchronized with the data on the cache volume.
Unknown	Indicates that an unknown error was detected.

Busy Volume States Table

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



General Parameters Table

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

5.3.6.8 Renaming a Volume

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

- Under 'Home', in the storage system view, click the volume that the user wants to rename. The volume properties are now displayed on the right.
- Click 'Rename'.
- Type a new volume name, and then click 'OK'.

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

5.3.6.9 Rebuilding a Volume

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

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



5.3.6.9.1 Rebuilding from 'Status' (Manually)

Verify that the volume is reported as degraded in the Manage subsection. If the user has more than one volume listed in this section, the user will need to fix the issues reported one at a time.

1. Click 'Rebuild to another disk' next to the volume the user wants to rebuild.
2. In the Rebuild Volume dialog, select the disk that will replace the failed disk. Only compatible disks in a normal state will be displayed. Refer to Volume Requirements for more information.
3. Click 'OK' to confirm.

The volume starts rebuilding and the page refreshes displaying the progress of the operation. The user can use other applications during this time and the user will be notified when the process has successfully completed.

5.3.6.9.2 Rebuilding from 'Manage' (Manually)

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

5.3.6.9.3 Auto Rebuild to Missing Disk

The Intel VROC 6.1 driver will automatically start a rebuild to a rediscovered member drive of a redundant RAID volume unless a rebuild to a spare has already begun.

5.3.6.9.4 Rebuild Resumption

If a rebuild operation is interrupted the Intel VROC drivers will resume the rebuild at the last good known location.

5.3.6.10 Resetting Volume to Normal

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

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

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

1. Under 'Home', in the disk "Properties" pane, click 'Reset volume to normal'.
2. Click 'Yes' to confirm.
3. The page refreshes and the volume displays as normal. If the operation failed to return the volume to a healthy state, visit Intel's online support web site for more options.



5.3.6.11 Changing Volume Type

The user can choose to change the type of an existing volume based on the user's storage system needs. Refer to section 9 for more details.

Note: Before starting, refer to the system and volume requirements to determine which RAID types are supported by the user's computer and make sure the required number of disks are connected. The Intel® chipset provides support for the creation of all RAID volume types and for up to 8 SATA ports and 48 Intel NVMe SSD connected to the CPU. Changing volume type does not require re-installation of the operating system

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

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

If the user performs a driver upgrade or downgrade while the data migration is in progress and then restart the user's computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the user will the system not be able to be restart until the migration is completed otherwise the operating system cannot load. If the user is migrating a data volume, the user will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.

5.3.6.12 Increasing Volume Size

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

After creating a volume with an array allocation set to less than 100% in the Configure Volume step, the user will be able to increase the volume size by the amount of available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size. The new drive(s) added must be equal to or larger than the smallest drive in the Array

This option is only available if:

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

WARNING: If the user performs a driver upgrade or downgrade while the data migration is in progress and then restart the user's computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the user will the system not be able to be restart until the migration is completed otherwise the operating system cannot load. If the user is migrating a data volume, the user will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.



5.3.6.12.1 Increasing the Volume Size from Manage Array

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

5.3.6.12.2 Increasing the Volume Size from Manage Volume

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

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

5.3.6.13 Adding a Disk to a Volume

The user can add one or more disks to an existing array to increase the system storage capacity. This feature can be useful if the user wants 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. The user cannot add an external disk to an array that includes internal disks, and vice versa. In specific advanced system configurations, this condition may not apply.

Refer to Connecting a Disk under Managing Disks for more information on installing SATA disks on the user's computer.

WARNING: Any existing data on the available disk used to increase the array size will be permanently deleted. Backup all the data the user wants to preserve before completing this action.

If the user performs a driver upgrade or downgrade while the data migration is in progress and then restart the user's computer, the driver will not be able to recognize the volume or the data on it. If the user is migrating a system volume, the user will the system not be able to be restart until the migration is completed otherwise the operating system cannot load. If the user is migrating a data volume, the user will have to reverse (roll back) that last performed driver update, and then restart the computer to return to a normal state.



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

1. Under 'Home', in the storage system view, click the volume to which the user wants to add a disk. The element properties are now displayed on the right.
2. Click 'Add disk'.
3. Select the disk the user wants to use to increase the array capacity.
4. Click 'Add Disk'.

CAUTION: Once the data migration starts, the operation cannot be canceled.

Once the migration has completed, restart the user's computer for changes to take effect. Then use Windows Disk Management* to increase the partition size on the volumes for which a disk was added, or add another partition.

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

5.3.6.14 Enabling Volume Write-back Cache

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

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

Enabling the volume write-back cache

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

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

Disabling the volume write-back cache

1. Under 'Home', in the storage system view, click the volume for which the user wants to disable the write-back cache. The volume properties are now displayed on the right.
2. In the Advanced section, click 'Disable', and then click 'Yes' to confirm.

The page refreshes and the write-back cache is now disabled.



5.3.6.15 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 the user attempts to start a verification process for a volume that has not been initialized, the user will be prompted to do so.

Initializing a Volume

1. Under 'Home', in the storage system view, click the volume that the user wants to initialize. The volume properties are now displayed on the right.
2. Click 'Initialize'.
3. Click 'OK' to start the initialization process. Caution: Once the data migration starts, the operation cannot be canceled.

Initializing a Volume when Verifying Data

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

Note: While initialization is in progress, the user can view the status in the notifications area by hovering over the Intel Rapid Storage Technology enterprise icon, or in the application under the Home pane.

WARNING: The initializing process could take a while depending on the number and size of the disks. The user can continue using array disks and other applications during this time. Closing the application, or powering off and restarting the user's computer will not disrupt the progress of this operation.

5.3.6.16 Verifying and Repairing Data

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

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

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



5.3.6.17 RAID Level Migrations

The RAID level migration feature in the Intel VROC 6.1 products will enable the ability to convert the contents of a drive (behind Intel VMD or attached to the SATA and/or sATA controller) into a RAID volume (RAID 0, RAID 1, RAID 5, or RAID 10). The RAID level migration feature also supports the ability to migrate from a one RAID volume to another.

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

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

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

Note: When migrating to a RAID volume that is greater than 2 TB, the partition will be migrated from an MBR partition to GPT.

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

Change Type from	To (SATA/sATA)	To (NVMe (VMD/Legacy))
2-disk RAID 1	2-disk RAID 0 3 to max ports on controller RAID 0 3 to max ports on controller RAID 5	2-disk RAID 0 3 to 48-disk RAID 0 3 to 48-disk RAID 5
2-disk RAID 0	3 to max ports on controller RAID 5	3 to 48-disk RAID 5
3-disk RAID 0	4 to max ports on controller RAID 5	4 to 48-disk RAID 5
4-disk RAID 0	5 to max ports on controller RAID 5	5 to 48-disk RAID 5
4-disk RAID 10	4 to max ports on controller RAID 5	4 to 48-disk RAID 5

Note: Additional information on RAID level migration can be found in the Intel VROC GUI help

The Intel VROC 6.1 GUI provides the ability to show the migration process.

5.3.6.18 Migrating From one RAID Level to Another

RAID level migration allows an existing RAID configuration to be migrated to another RAID configuration. The following migrations are possible.

NOTES:

- 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® VROC product requirements document);
- For the migration options to be accessible, the minimum required drives for the RAID level have to be met.



Perform the following procedure:

1. Start the Intel Virtual RAID on CPU GUI application (run as administrator).
Start Menu -> All Programs -> Intel -> Intel (R) Virtual RAID on CPU
2. From the Volumes pane, click the array or volume to which the user wants to modify.
The volume properties now displays on the right.
3. Click 'Change type'.
4. In the 'Change Volume Type' dialog, type a new name if the user wants 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, the user can view the migration progress under status.

When the Status field indicates volume as 'Normal', the migration is complete

5.3.7 Managing Disks

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

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

5.3.7.1 Disk Properties

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

Parameter	Value
Port	Reports the port number to which the disk or device is attached.
Port location	Reports whether the port is internal or external.
Usage	Array disk: a disk that has been grouped with other disks to form an array containing RAID volumes.
	Spare: the disk has been designated as the destination drive for automatic volume rebuilds in the event of a failed, missing or at risk array disk. For RAID 0 volumes, automatic rebuilds will only occur when one of its array disks is reported as at risk.
	Available: the disk is physically connected to the computer, healthy, and available to be used in an array or as a spare disk. WARNING: Assigning an available disk to an array or marking it as a spare will permanently delete any existing data on that disk.



Parameter	Value
	Unknown: the disk is available but contains metadata that cannot be displayed in the operating system. Even though the disk is reported as normal, the user will need to clear and reset the disk to make the disk available.
Status	Normal: the disk is present, functioning as expected.
	At risk: an impending error condition was detected on the disk and it is now at risk of failure.
	Missing: the disk is not present or physically connected to the computer.
	Failed: the disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.
	Offline: indicates that an array disk contains metadata but no longer fits the Array or RAID Volume. e.g. Drive-1 and Drive-2 are in a RAID 1 configuration. Drive-1 is removed and replaced with Drive-3. The RAID 1 is rebuilt. Drive-1 is reinserted into the system. Drive-1 will be marked as "offline" because it is no longer an active member of the RAID volume. The VROC GUI (that contains all three of these drives) will not do anything with this driver until RAID metadata is cleared. If Drive-1 is removed and taken to another system, That VROC GUI will show the drive as "online" as part of a degraded RAID1 volume.
Status	Reports the total capacity of the disk in megabytes (MB) in the disk properties
Size	and in gigabytes (GB) in the storage system view.
Serial number	Reports the manufacturer's serial number for the disk.
Model	Reports the model number of the disk.
Firmware	Reports the version of the firmware found in the disk.
System disk	Reports whether the disk contains system files that are required to start and run the operating system.
	Reports whether the disk is protected with a password.
Disk data cache	Reports whether the disk supports this feature.
Native command queuing	Reports the data transfer rate between the SATA controller and the SATA disk. The supported rates are: SATA 1.5 Gb/s (generation 1) SATA 3 Gb/s (generation 2) SATA 6 Gb/s (generation 3) The data transfer rate reported is based on the Intel® Chipset and SATA disks present in the user's system.
Physical sector size	Reports the size of physical sectors on the disk (bytes).
Logical sector size	Reports the size of logical sectors on the disk (bytes).



5.3.7.2 Marking a Disk as Spare

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

Marking a disk as a spare allows the user to designate an available disk as the default destination for automatic volume rebuilds in the event of a failed, missing or at risk array disk. However, for RAID 0 volumes, automatic rebuilds will only occur if one of its members is reported at risk

1. Under 'Home', in the storage system view, click the disk that the user wants to mark as a spare.
The volume properties are now displayed on the right.
2. Click 'Mark as spare'.
3. Click 'OK'.

Note: RAID 1, 5 and 10 volumes can use one or more spares.

WARNING: When marking a disk as a spare, any existing data on that disk is permanently deleted. Back up all data the user wants to preserve before starting this action.

If the user's system is running a version of the Intel® VROC (SATA RAID) Legacy OROM that does not support disks that are 2TB or larger, the user can reset such a disk to available, but disallow the marking of it as a spare.

5.3.7.3 Resetting a Disk to Available

After a disk was marked as spare, the user can choose to make that spare disk available again and use it differently. Once available, the disk can be used to create a volume or be added to an existing volume if all other requirements are met.

1. Under 'Home', in the storage system view, click the disk that the user wants to reset to available.
The volume properties are now displayed on the right.
2. Click 'Reset to available'.
The page refreshes and the disk usage is now reported as available.

5.3.7.4 Resetting a Disk to Normal

The user can reset a disk to normal when the storage system reports one of the following disk statuses:

At Risk

A disk is reported at increased risk of failing in the near future that could be due to a slow degradation over time. The user can choose to ignore this alert at this time by resetting the disk to normal, but it may re-appear if the disk continues to assert this condition. We recommend that the user contact the manufacturer for more information to prevent potential data loss.

Failed

A disk has failed to properly complete read and write operations in a timely manner, and data may be lost. We recommend that the user replace the failed disk as soon as possible to return the overall storage system to normal. In this state, data may be lost, but the user can try resetting the disk to normal and attempt a data recovery. If the disk operations continue to fail, the disk will return to a failed state immediately.



If the failed disk is an array disk, refer to Appendix F: Troubleshooting for guidelines on rebuilding a failed or degraded volume.

1. Under 'Home', in the "Device" pane, locate the disk reported as at risk or failed.

The user can also perform this action from Disk Properties to the right...

2. Click 'Reset disk to normal'.

The page refreshes instantly, returning to a normal state.

Note: Completing this action clears the event on the disk and does not delete existing data. However, ignoring early warning signs of disk failure may result in data loss.

5.3.7.5 Resetting a Disk to Online/Normal

After a disk was marked "offline", the user can bring that drive back "online" by removing the Metadata on the drive under drive properties. Once back "online", the disk can be used to create a volume or be added to an existing volume if all other requirements are met.

1. Under 'Home', in the storage system view, click the disk that the user wants to bring back "online".

The Disk properties are now displayed on the right.

2. Click "Clear metadata" on the Status line, next to "Offline"

3. A window will pop up explaining that the data on the disk will be deleted. Click "Yes" to complete the process.

Note: If there is critical data on the drive, take the drive to another system to recover the data before clearing the metadata.

The page refreshes and the disk usage is now reported as "Normal"

This process can also be used via the Intel VROC CLI and also in the Pre-OS environment.

5.3.7.6 Connecting a Disk

Installing new hardware is one of the steps the user may have to take to keep the storage system healthy or to extend the life of a computer that is running out of storage space.

The Intel VROC 6.1 family of products will provide NVMe and SATA drive hot plug support, which is a feature that allows disks to be removed or inserted while the computer is turned on and the operating system is running. As an example, hot plugging may be used to replace a failed external disk.

Our application provides support for SATA 1.5 Gb/s (generation 1), SATA 3 Gb/s (generation 2), and 6 Gb/s (generation 3) data transfer rates. The rate support depends on the Intel® chipset and SATA disks present in the user's system.

Follow these procedures to replace or connect a disk in case the user needs to power off the user's computer:

Replacing a Disk

1. Power off the user's computer.
2. Replace the disk that reports a problem.
3. Turn the user's computer back on.



If the replaced disk was part of an array, the user will need to follow the procedure provided in Appendix F: Troubleshooting based on the volume state and type.

Note: To install an external disk, plug it into the user's computer and connect the power cord.

To remove and install an internal disk, the user should be comfortable opening the user's computer case and connecting cables. Follow the manufacturer's installation guide to complete this procedure. If the user is replacing the system disk, the user will have to re-install the operating system after the user connects the disk because the system disk contains the files required to start and run the user's computer.

Installing a New Disk (to increase storage space)

1. Power off the user's computer.
2. Connect the new disk.
3. Turn the user's computer back on.

During the system startup, the application's PreOS should automatically detect the new disk if it is installed correctly. Once the user opens the application, verify under 'Home', in the storage system view that the new disk displays. The user can then access management options by clicking that disk.

5.3.7.7 Managing Ports

A port is a connection point on the user's computer where the user can physically connect a device, such as a SATA disk, an ATAPI device or an Intel NVMe SSD. A port transfers I/O data between the device and the computer.

If a port is reported as empty in the storage system view, the user can use that port to connect a new device in order to increase the storage system capacity. Currently, the maximum number of internal ports that can be used to connect SATA devices is eight (per controller).

The port properties listed below display to the left of the storage system view under 'Manage' and report values specific to the element selected in the view.

Parameter	Value
Port	Reports the port number to which the disk or device is attached.

5.3.7.8 Detecting a Spare Drive

The Intel VROC 6.1 family of products will provide support to detect spare drives.

5.3.8 Disk Monitor Service

Intel VROC 6.1 family of products will support the ability to provide a disk monitoring service. The service will be active by default and executed as a system service. The service will monitor the system for SMART and RAID volume state changes events. The changes will be logged in the system log.



5.3.9 NVMe Deallocate/SATA TRIM Command (RAID 0, 1 and 10)

Note: This feature is not an end-user visible feature. There is no Intel® VROC 6.1 application or user interface control to configure the feature. Registry settings are provided for OEM use.

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

5.3.10 Un-installation

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

The uninstall application that is included with the Intel® VROC 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 VROC software or to disable the SATA RAID controller:

5.3.10.1 Uninstalling the Intel VROC Software (except the RAID Driver)

1. Run the Uninstall program from the following start menu link:

Start→All Programs→Intel -> Intel(R) Virtual RAID on CPU

The first dialog box that appears gives the user the option of un-installing all components of the Intel® Rapid Storage Technology enterprise software except the RAID driver.

2. Click 'OK'.

The next dialog box is a confirmation that the user would like to un-install all components of the software except the RAID driver.

3. Click 'Yes' to confirm.

All components of the software will be un-installed except the RAID driver. The user should no longer see any Start menu links to the UI application or a control panel applet for Intel® VROC. However, the RAID configuration should still function normally.



5.3.10.2 Disabling the RAID Driver by Disabling the RAID Controller

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

1. Enter System BIOS Setup and disable the Intel VMD or SATA RAID Controller. This setting may be different for each motherboard manufacturer. Consult the user's manufacturer's user manual if necessary. When done, exit Setup.
2. Reboot the system (The OS must have been installed on a disk not attached to the Intel VMD or SATA RAID controller). The user should no longer see the Intel VROC PreOS status screen during boot, and the user should no longer see the Intel VROC Controller(s) in Device Manager.
3. At this point, Windows* will no longer be using the RAID driver and the user will not have Intel RAID functionality. **All data contained in existing RAID volumes will no longer be accessible.** To re-enable Intel RAID functionality, re-enter System BIOS Setup and re-enable RAID mode.

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

5.4 Intel VROC 6.1 Power Management

Intel VROC 6.1 family of products will support all the following power management functions required by the OSs:

- Working state (S0)
- Sleep state (S3) including Hybrid sleep (S3 + hibernation file)
- Hibernation state (S4) and S4/S5 Fast startup (reduced hibernation file)

Supporting these power states applies to scenarios without hot-swapping the Intel VROC RAID Upgrade Hardware Key.

5.5 Intel VROC 6.1 Miscellaneous Features and Functionality

5.5.1 OS Installation

The Intel VROC 6.1 family of products will provide the OS appropriate driver files required for installation during the OS setup onto a drive or RAID volume attached to either the Intel VMD, SATA or sSATA controllers.

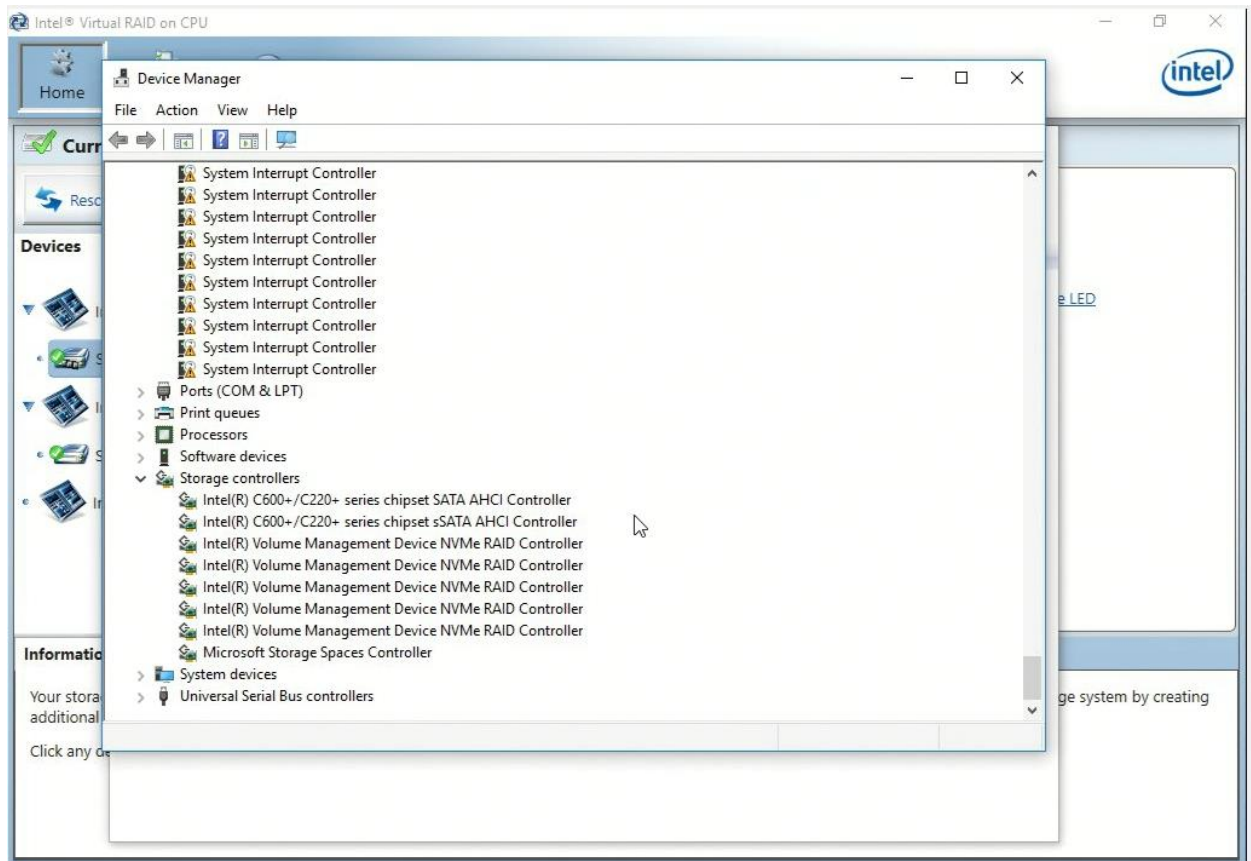
Note: Installing an OS onto an Intel VROC (NonVMD NVMe RAID) Volume is not supported.

5.5.2 Using the Device Manager

1. Open the Start Menu of the user's system and type the words 'Device Manager.'
2. Scroll down to the section that reads Storage Controllers.
3. Left Click on Storage Controllers to expand the menu.

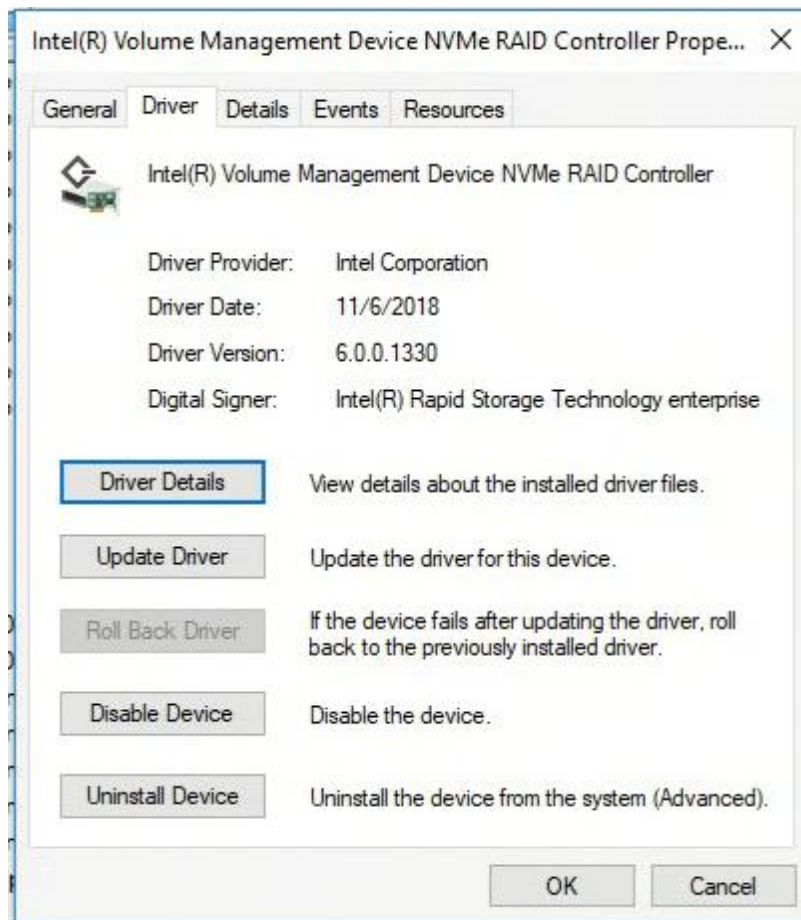


The exact wording will vary some depending on operating system involved, Windows* 10 show this as Microsoft* Storage Spaces Controllers.



4. Right Click on Microsoft* Storage Spaces Controllers and select Properties.
5. Select the Driver tab.

Driver Version denotes the version of the drivers that are installed on the system controller for this system.



5.5.3 Disk Write Cache

Intel VROC 6.1 family of products will support the ability to enable/disable Disk Write Cache through the Intel VROC GUI. Disk Cache will be disabled by default.

5.5.4 Write Back Cache

Intel VROC 6.1 family of products will support the ability to enable/disable Write Back Cache through the Intel VROC GUI. Write Back Cache will be disabled by default.

5.5.5 4K Native Sector Sizes Drives

Intel VROC 6.1 family of products will supports the following drives:

- 512 and 512b sector size drives
- Native 4K sector size drives for SATA/sSATA controllers and NVMe SSDs for Intel VROC.

Note: The only limitation is the inability to use the mixed 4K and 512(b) in one RAID volume.

Note: This feature is only supported in UEFI mode.



5.5.6 Intel VROC 6.1 Microsoft* WHQL Drivers

The production version of the VROC 6.1 drivers shall be released after passing Microsoft* WHQL logo requirements for storage devices under all supported Operating Systems.

5.5.7 Intel VROC 6.1 Microsoft* UWD Support

The Intel VROC 6.1 family of products meets Microsoft*'s UWD "D", "C" and "H" compliance criteria. Intel VROC 6.1 supports this feature only new (supported) platforms that have launched after Windows* 10 RS5's official release.

Note: UWD (specific) related issues will only be addressed if they occur (or can be reproduced) on one of these new (supported) platforms.

5.5.7.1 INF Installation Support

The Intel VROC 6.1 family of products supports the ability to install each of the drivers provided by using the INF file. All that needs to be done is to highlight the .inf file, right mouse click and select install. Once prompted to, please reboot the system to ensure the installation process takes full affect.

5.5.7.2 Windows* Application Store Support

The Intel VROC 6.1 family of products has provided/uploaded the Intel VROC 6.1 GUI application to the Microsoft* Application Store for download. All that must be done is to go to the Microsoft* Application Store, search for "Intel VROC", download and install. This process corresponds with the driver .inf installation process previously mentioned.

Note: The Intel VROC drivers and GUI can also be installed the standard way by running the installation executable application provided.

5.5.7.3 Stand Alone Intel VROC 6.1 GUI

The Intel VROC 6.1 family of products supports the ability to load the Intel VROC 6.1 GUI without running the installer or going to the Microsoft* Application Store. Within package there is a "Standalone" compressed file that can be copied to the target system, unzipped and executed from that location. As with the application uploaded to the Windows* Application Store, this process requires that all of the Intel VROC drivers have already been installed via the .inf installation process.

5.5.8 Intel SMART Event Support

The Intel VROC 6.1 family of products will provide support for SMART Alerts for NVMe and SATA disks. A SMART drive event response alert, indicating that a drive is in danger of failing, will initiate a rebuild to predefined hot spare disk. If a hot spare has not been defined, the RAID volume will become degraded.

5.5.9 Event Log Support

The VROC 6.1 family of products will log member disk error status to an event log stored in RAID metadata. Please see [Appendix B: Storage System Events](#) for more details.



5.5.10 2 Terabyte Disk Size

Intel VROC 6.1 family of products will be able to support drives that are larger than 2 Terabytes.

5.5.11 2 Terabyte Volume Size

The Intel VROC 6.1 family of products will be able to create and operate with RAID volumes that are greater than 2 Terabytes in size.

5.5.12 Email Alerting and Notification

The Intel VROC 6.1 family of products will support email notification of certain storage events.

5.5.12.1 E-mail Notification UI Visible Enable/Disable

Intel VROC 6.1 will support email notification of certain storage events. Please see [Appendix B: Storage System Events](#) for a full list of supported events. The Intel® VROC GUI will provide the interface for enabling/disabling and configuring the email notification feature. **The default setting in the GUI is 'disabled'.**

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

5.5.12.1.1 Configuration

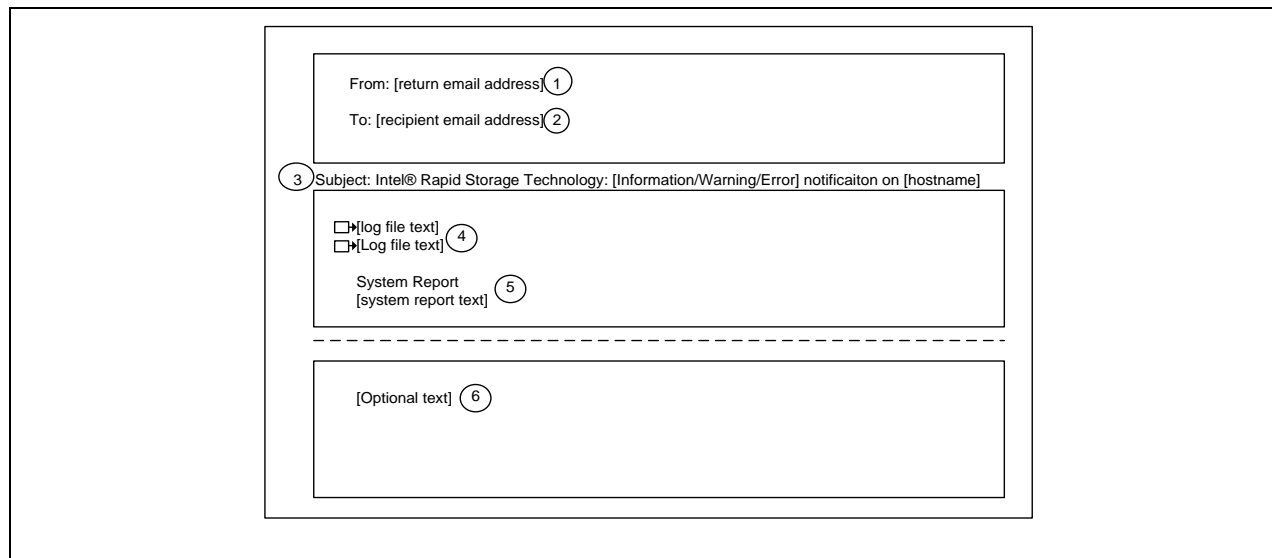
The Intel VROC 6.1 GUI will provide the interface to allow the user to configure the email alert notification feature via the 'Preferences' tab of the UI (**user must be logged on with administrative privileges**).

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

5.5.12.1.2 Email Message Format

- Message header:
 - Item1. Return email address: email address of the originating computer
 - Item2. Recipient email address: email address of computer receiving the email notification
 - Item3. Subject: system formatted subject content with product name, the storage system event level and the hostname of the originating computer

- Message body:
 - Item4. Log file text: contains the text of the event as it is displayed in the event log
 - Item5. System report: contains the system configuration information of the originating computer as seen in the Intel® VROC GUI Preferences page.
- Optional text:
 - Item6. This section is blank unless the originating computer's OS is in a language other than English. If the originating computer sends items 4 and 5 in non-English, the English translation of those two items will appear in this section (for test emails, only item 4 will be translated here)



5.5.12.1.3 Protocol Support

Email alert shall support SMTP host & SMTP port.

Note: Intel VROC 6.1 will check the SMTP Host and Port settings but does not check credentials for clients.

5.5.12.1.4 Error Conditions

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



5.5.13 Microsoft .NET Framework

The Intel VROC product installation application has removed the Microsoft .NET Framework as well as the Intel ASM component.

The following table shows how the removal of the Microsoft .NET Framework may impact the launching of the Intel VROC GUI, based off the Windows operating system installed:

Intel® VROC 6.1	
Server 2k12 R2	Install Latest .NET Framework
Server 2k16	Install Latest .NET Framework
Windows 2k19	No Impact
Windows 10 RS3	Install Latest .NET Framework
Windows 10 RS4	No Impact
Windows 10 RS5	No Impact
Windows 10 19H1	No Impact

If the system configuration requires the .NET Framework version to be updated and the system has internet access, a web installer can be used, which should go out and install the latest version. For example: (<https://support.microsoft.com/en-us/help/4054531/microsoft-net-framework-4-7-2-web-installer-for-windows>).

If the system is not connected to the Internet, then an offline version must be downloaded, moved to and installed on the system. The following are some additional instruction to help in this process:

1. Download the latest version of .NET Framework from Microsoft
2. Compress the downloaded image (to avoid potential undesirable side effect as outlined in <https://docs.microsoft.com/en-us/dotnet/framework/install/troubleshoot-blocked-installations-and-uninstallations#compat>)
3. Copy the compressed file to a USB drive
4. Copy the compressed file from the USB drive to the Download directory of the platform being configured
5. Uncompressed the file
6. Run the executable file as administrator

For more information please refer to <https://dotnet.microsoft.com/>.

Once the latest version of the .NET Framework is installed, rerun the Intel VROC product installation application. This helps ensure that all components will start properly.

5.6 Intel VROC Package Updates

Intel VROC 6.1 family of products supports upgrading to new release packages. Each Intel VROC release package contains the PreOS components, the OS drivers, the product installation application and tools. Intel recommends that all of the Intel VROC components, installed on the supported platform, be upgraded to the latest version of the product. When performing a package update, it is important to follow the recommended upgrade order.

1. Upgrade the Intel VROC OS driver package first
2. Upgrade the BIOS with the corresponding Intel VROC PreOS image second



It is more challenging to update the platform BIOS than it is to update the OS drivers and tools. The Intel VROC OS drives and tools are designed to be backwards compatible with older Intel VROC PreOS images (unless otherwise outlined in the products release notes). The Intel VROC PreOS is not backwards compatible with the older versions of the Intel VROC OS drivers. The reason for this is because there is information passed from the Intel VROC PreOS environment to the Intel VROC OS driver. The Intel VROC OS driver relies on this information to function properly. If this information happens to change (as part of a bug fix or feature enhancement), an older Intel VROC OS driver may not initialize properly, not function and could lead to a system failure.

5.6.1 Intel VROC Package Upgrade

The Intel VROC package upgrade is the process of transitioning to the latest/newest version of the Intel VROC PreOS and/or Intel VROC OS components. Table 1 show the Intel VROC PreOS versions to Intel VROC OS version support matrix to help with upgrading the platform.

Table 2: Intel VROC Compatibility Matrix

	Intel VROC PreOS 5.5	Intel VROC PreOS 6.0	Intel VROC PreOS 6.1
Intel RSTe 5.5	Supported	Not supported	Not supported
Intel VROC 6.0	Supported	Supported	Not supported
Intel VROC 6.1	Supported	Supported	Supported

5.6.2 Intel VROC Package Downgrade

The Intel VROC package downgrade is the process of transitioning to an older version of the Intel VROC PreOS and/or Intel VROC OS components. Intel does not recommend downgrading. If it is deemed necessary to downgrade, do so using the follow order of steps:

1. Upgrade the BIOS with the corresponding Intel VROC PreOS image first
2. Upgrade the Intel VROC OS driver package second



6 Intel® VROC (SATA RAID) Key Features

The following is a summary of the key features of this product that are supported on Intel VROC (SATA RAID) PCH/SATA devices. Not all legacy Intel VROC (SATA RAID) features below will be supported on Intel VROC RAID for NVMe devices.

Name	Key Features	
RAID	<ul style="list-style-type: none"> • NCQ (SATA) • Bad Block Management¹ 	<ul style="list-style-type: none"> • Read Patrol • SGPIO • Dirty Stripe Journaling¹ • Partial Parity Logging (PPL)¹

Note: 1. Items that are supported on both Intel VROC (SATA RAID) PCH/SATA and Intel VROC

Intel VROC (SATA RAID) is designed to support both M.2, SATADOM, and standard form factor SATA drives.

6.1 Intel VROC (SATA RAID) Pre-OS Features and Functionality

6.1.1 Intel VROC (SATA RAID) RAID Legacy Option ROMs

Intel® VROC (SATA RAID) will provide a Legacy RAID OROM images that support standard Int13 BIOS environments. These images provide a BIOS level menu-driven configuration tool for managing Intel VROC (SATA RAID) volumes outside of an OS. There are individual images for each of the PCH controllers on the Intel® C620 and C230 series chipset platforms:

- There is a single Intel VROC (SATA RAID) Legacy Option ROM image that will support the AHCI/SATA controller for each of the mentioned chipsets.
- For the Intel® C610/C620 series chipsets there is also a RAID Legacy Option ROM for the sSATA controller.
- The Intel VROC (SATA RAID) Legacy OROM will support 48bit LBAs for pre-boot platform IO. This ensures that the Intel VROC (SATA RAID) OROM is capable of providing native OROM boot support for 2TB disks and RAID volumes. It is up to the platform BIOS to provide a means of sending IO that utilizes the full LBA address space.

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

Upon re-boot, the user will see the option ROM status message on the screen.

1. Press CTRL-I to enter the Intel Rapid Storage Technology enterprise option ROM user interface.
2. In the Main Menu, select option #1 'Create RAID Volume'.
3. Enter the name the user wants to use for the RAID volume, then press Enter.
4. Select the RAID Volume Name and the RAID Level.
5. Press Enter to select the disks to be used by the array that the volume will be created on.
The key to select the disks will be shown at the bottom of the screen, usually this is the <Space> bar.
6. Press Enter when done.
7. Select the strip size (64 KB is the default for RAID 5) by using the arrow keys, then press Enter.
8. Enter the size for the RAID volume in gigabytes.
The default value will be the maximum size. If the user specifies a smaller size, the user will be able to create a second volume in the remaining space using the same procedure.
9. When done select arrow down to "Create Volume" and press enter to complete.



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

6.1.1.2 Legacy OROM Reset Disks to Non-RAID

When the 'Reset Disks to Non-RAID' option is selected in the Intel VROC (SATA RAID) Legacy OROM, the RAID metadata will be removed from the disk, returning it back to a pass-through drive.

6.1.1.3 Legacy OROM Rebuild Messaging

If the Intel VROC (SATA RAID) Legacy OROM detects that one or more RAID volumes are in the 'Rebuilding' state then a message will be displayed stating that the RAID Volume is in a rebuild state.

6.1.1.4 Legacy OROM Banner when locked HDDs are detected

The Intel VROC (SATA RAID) Legacy OROM will present a banner to the user if a locked HDD is detected in the system when it is powering on or if the system is resuming from a hibernate (S4).

6.1.1.5 Legacy OROM boot from password protected disk

When a system disk contains password protection the Intel VROC (SATA RAID) Legacy OROM will only boot from the disk if unlocked.

6.1.1.6 Legacy OROM boot from password protected volume

When a system volume contains member disks with password protection the Intel VROC (SATA RAID) Legacy OROM will only boot from the RAID Volume after all of the member disks are unlocked. If any member disk is password protected, the system may not boot properly.

6.1.1.7 Legacy OROM RAID Volume limit

The Intel VROC (SATA RAID) Legacy OROM environment supports a maximum of 4 RAID volumes on the SATA/sSATA controller

6.1.1.8 Legacy OROM Dirty Shutdown Recovery

The Intel VROC (SATA RAID) Legacy OROM for SATA and sSATA will be able to recover from the RAID 5 volume invalid state caused by Dirty Shutdown condition if all RAID5 volume disk members are enumerated during system boot.

6.1.1.9 Boot Device Name

The Intel VROC (SATA RAID) Legacy Option ROM will provide the name of any bootable device recognized by the Option ROM to the system BIOS

6.1.1.10 Migration Support

The Intel VROC (SATA RAID) Legacy Option ROM will support I/O access to a RAID volume that has a migration in progress.



6.1.1.11 Using the BCFS to Differentiate Platform SKUs

Intel VROC (SATA RAID) Pre-OS has support for BIOS Control Feature Set (BCFS) to enable OEMs the opportunity to customize the Intel® VROC (SATA RAID) feature offerings. OEMs can enable/disable the desired features per platform SKU directly in their BIOS code. By clearing or setting the corresponding bits of the '**Intel VROC (SATA RAID) Feature Capabilities**' register in the Intel chipset's SATA controller MMIO space, OEMs now have greater flexibility in determining what Intel® VROC (SATA RAID) features will be supported per platform model/SKU.

The following sections explain the use of each of the bits in the BCFS, also known as the Software Feature Mask bits.

BCFS Bit Number	Bit Meaning	Values		Additional Info
0	Enable/disable Raid0	Raid type disabled	0x0	If you disable all raid levels – all BCFS settings will be set back to default and OROM UI delay will be set to 2 seconds
		Raid type enabled	0x1	
1	Enable/disable Raid1	Raid type disabled	0x0	
		Raid type enabled	0x1	
2	Enable/disable Raid10	Raid type disabled	0x0	
		Raid type enabled	0x1	
3	Enable/disable Raid5	Raid type disabled	0x0	
		Raid type enabled	0x1	
4	RESERVED			
5	Enable/disable UI	Feature disabled	0x0	If you set it to 0 (disable) bits 10-14 are ignored
		Feature enabled	0x1	
6	Enable/disable unlock of HDD in OS	OS can unlock HDDs OS cannot unlock HDDs	0x1 0x0	
7	Enable/disable LED SGPIO	Feature disabled Feature enabled	0x0 0x1	
8	eSATA RAID usage	Allow all volume types to span internal and external ports.	0x0	
		Allow only R1 volume type to span internal and external ports. Other volume types – allow all eSATA-internal disks only if there are enough of them to create the volume	0x1	
9	RESERVED			



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

Note: This document does not cover details on how to setup a system BIOS. For that level of information please contact the user platform's BIOS vendor or the user's Intel field representative to put the user in contact with the appropriate Intel BIOS support personnel.

6.1.1.11.1 Configuring the Platform's RAID Related Features

When the BIOS has set the SATA Controller's mode to RAID, the following bits of the 'Intel VROC (SATA RAID) Feature Capabilities' register in the Intel chipset's SATA controller MMIO space will determine what RAID levels will be supported on the platform SKU:

Note: Clearing all RAID level related bits to '0' is an unsupported configuration. The Intel VROC (SATA RAID) Legacy OROM will ignore the BIOS settings and enable all RAID levels.

6.1.1.12 BCFS Bit Setting

6.1.1.12.1 Example Configuration

To configure a platform SKU that offers **only** RAID levels 0 and 10, the bits must be configured as follows:

Bit 0 == 1 (default)
Bit 1 == 0
Bit 2 == 1 (default)
Bit 3 == 0



6.1.1.12.2 Configuring the Behavior of the OROM UI and Banner

There are three (3) bits that control the behavior of the Intel VROC (SATA RAID) Legacy OROM UI and the Banner Splash Screen that are displayed during POST at system boot-up. Use the following bit configurations to determine whether or not the splash screen will be displayed during post and if so, how long the delay will be before the system continues the boot process:

Bits	Type	Reset/Default	Description
11:10	RWO	0h	OROM UI Normal Delay (OUD): Values of these two bits specify the delay of the OROM UI Splash Screen in a normal status. 00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs Note: If bit 5 == 0, then these values are disregarded Comment: Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.
5	RWO	1h	Intel VROC (SATA RAID) OROM UI (RSTOROMUI): If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.

6.1.1.12.3 Example Configuration

To configure a platform SKU that enables the OROM Banner Splash Screen to be displayed for 6 seconds, the bits must be configured as follows:

Bit 5 == 1 (default)
 Bit '10' == 0 (default)
 Bit '11' = 1

6.1.1.12.4 Configuring Intel® VROC UI Capabilities

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

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

Bits	Type	Reset/Default	Description
6	RWO	0h	HDD Unlock (HDDLK): If set to '1', then HDD password unlock is enabled in the OS.
10:12	RW	0x000	Delay on UI splash screen: 2 seconds 0x000 4 seconds 0x001



			6 seconds 0x010
			8 seconds 0x011
			10 seconds 0x100
			15 seconds 0x101
			30 seconds 0x110
			60 seconds 0x111

6.1.1.12.5 Example Configuration

Bits	Type	Reset/ Default	Description
15:14	RO	0h	Reserved.
13:12	RWO	0h	Reserved
11:10	RWO	0h	<p>OROM UI Normal Delay (OUD): Values of these bits specify the delay of the OROM UI Splash Screen in a normal status.</p> <p>00 – 2 secs (default and previous value) 01 – 4 secs 10 – 6 secs 11 – 8 secs If bit 5 == 0, then these values are disregarded</p> <p>Comment: Allow OEM to lengthen normal timeout of OROM splash screen so user has more time to hit CTRL+I on keyboard.</p>
9	RWO	0h	Reserved
8	RWO	0h	eSATA RAID usage: If set to '1', allow only R1 volume type to span internal and external ports. Other volume types – allow all esata-internal disks only if there are enough of them to create the volume If cleared to '0', then allow all volume types to span internal and external ports.
7	RWO	0h	Enable/disable LED SGPIO: If set to "1", then LED SGPIO support is enabled.
6	RWO	0h	Enable/disable unlock of HDD in OS: If set to "1", then unlock of HDD in OS is enabled.
5	RWO	1h	Intel® VROC (SATA RAID) OROM UI: If set to '1' then the OROM UI is shown. Otherwise, no OROM banner or information will be displayed if all disks and RAID volumes are Normal.
4	RWO	0h	Reserved
3	RWO	1h	RAID 5 Enable (R5E): If set to '1', then RAID5 is enabled



Bits	Type	Reset/Default	Description
2	RWO	1h	RAID 10 Enable (R10E): If set to '1', then RAID10 is enabled
1	RWO	1h	RAID 1 Enable (R1E): If set to '1', then RAID1 is enabled
0	RWO	1h	RAID 0 Enable (R0E): If set to '1', then RAID0 is enabled

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

Bit 6 == 0 (default)

6.1.1.13 Intel VROC (SATA RAID) UEFI Dirty Shutdown Recovery

The Intel VROC (SATA RAID) UEFI drive will be able to recover from the RAID 5 volume invalid state caused by Dirty Shutdown condition occurrence if all RAID 5 disk members are enumerated during system boot.

6.1.2 ATAPI on SATA/sSATA Controller

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

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

6.2 Intel VROC (SATA RAID) Miscellaneous Features and Functionality

6.2.1 SGPIO on SATA Controller (RAID Mode Only)

Intel VROC (SATA RAID) product will support enclosure management, compliant to **SFF-8485** standard (Specification for Serial GPIO Bus) as well as **SFF-8489** (Specification for Serial GPIO IBPI (International Blinking Pattern Interpretation)), to identify drive location or unit failures on the SATA or sSATA controllers.

6.2.2 Common Storage Management Interface (CSMI)

Intel VROC (SATA RAID) product will support the Common Storage Management Interface (CSMI) for reporting RAID configurations and SMP, SSP, STP pass-through. This interface can also be used to update drive Firmware on pass-through drives attached to the SATA/sSATA controllers in RAID mode. The following commands are supported by Intel VROC 6.1.

```
CSMI_GET_CNTRL_CFG
CSMI_GET_CNTRL_STATUS
CSMI_GET_DRIVER_INFO
CSMI_GET_PHY_INFO
CSMI_GET_RAID_CONFIG
```



CSMI_GET_RAID_INFO
CSMI_SATA_SIGNATURE
CSMI_SSP_PASSTHRU
CSMI_STP_PASSTHRU
SSP_PASSTHRU
ATA_DOWNLOAD_MICROCODE

6.2.3 NCQ (SATA/sSATA)

The Intel VROC (SATA RAID) product will support Native Command Queuing (NCQ - The ability of the SATA drive to re-order commands in order to maximize the efficiency of gathering data from the platters).

6.2.4 Intel VROC (SATA RAID) Maximum Number of SATA Drives in a RAID Volume

The Intel VROC (SATA RAID) product will support to create SATA RAID volumes from maximum number of drives equal to the number of ports on SATA controller on the platform.

6.2.5 Intel VROC (SATA RAID) Maximum Number of sSATA drives in a RAID Volume

The Intel VROC (SATA RAID) product will support to create sSATA RAID volumes from maximum number of drives equal to the number of ports on sSATA controller on the platform.

6.2.6 Intel VROC (SATA RAID) Maximum Number of RAID Volumes

The Intel VROC (SATA RAID) product will support to create maximum number of RAID volumes equal to double the number of maximum arrays supported by the platform (corresponding SATA/sSATA controller).

6.2.7 SCSI Pass-through

The Intel VROC (SATA RAID) product will support the following Extended SCSI Pass-through commands as per Unified Extensible Firmware Interface Specification v2.3.1, Errata D:

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

6.2.8 RAID Volume Read Cache

Intel VROC (SATA RAID) will support the ability to enable/disable RAID Volume Read Cache through the Intel VROC 6.1 GUI. RAID Volume Read Cache will be enabled by default.

6.2.9 Intel VROC (SATA RAID) Support for Staggered Spin-up

The Intel VROC (SATA RAID) product will provide support for staggered spin-up on the SATA and sSATA controller for those drives that support this feature.



6.2.10 Intel VROC (SATA RAID) Driver Downgrade

If the driver is downgraded to a previous version RAID volumes on SATA will not be permanently broken. After upgrading back to the current version RAID volumes shall be fully operational.

§



7 Intel VROC (NonVMD NVMe RAID) Key Features

The Intel VROC 6.1 family of products includes RAID support for Intel NVMe drives on platforms outlined in [Supported Chipset SKU for Intel VROC \(SATA RAID\)](#). This feature is Intel VROC (NonVMD NVMe RAID). This provides Data RAID solution on Intel NVMe SSDs plugged into a PCIe slot managed by the Intel CPU (that does not contain Intel VMD).

7.1 PCIe Intel NVMe Device Usage Model

The Intel VROC (NonVMD NVMe RAID) release package provides support for Intel NVMe SSD in the following usages:

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

7.2 Intel VROC (NonVMD NVMe RAID) Boot Support

The Intel VROC (NonVMD NVMe RAID) release package does not support booting from a RAID volume, so **Intel VROC (NonVMD NVMe RAID)** does not contain a PreOS UEFI driver. It is not required nor is it needed.

7.3 Intel VROC (NonVMD NVMe RAID) Installation

The Intel VROC 6.1 installation utility will only install Intel VROC (NonVMD NVMe RAID) on those platforms outlined in [Supported Chipset SKU for Intel VROC \(SATA RAID\)](#).

7.4 Intel VROC (NonVMD NVMe RAID) Pass-through IOCTL Support

The Intel VROC (NonVMD NVMe RAID) product provides a private API that allows user-space applications to send and execute NVMe commands to remapped PCIe NVMe devices. This new API is based on the new IOCTL definition to implement NVMe pass-through channel.

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



7.5 Intel NVMe SSDs Support List

This is the list of supported Intel NVMe SSD included with the Intel VROC (NonVMD NVMe RAID) release.

Intel NVMe SSD	Device Identification
Intel® SSD 760p, E 6100p and Pro 7600p Series	F1A6
Intel® SSD D7-P5601/P5501 Series	0A60
Intel® Solid State Drive DC P3700/P3600/P3500/P3520 and 750 Series	0953
Intel® Solid State Drive DC P3520 Series	0A53
Intel® SSD DC P4500/4600/4501/4601/4608 Series	0A54
Intel® SSD DC P4610 Series	0A55
Intel® Optane™ SSD 900P Series	2700
Intel® Optane™ SSD DC P4800X Series	2701

7.6 Intel VROC (NonVMD NVMe RAID) RAID Volume Support with Multiple Controllers

The Intel VROC (NonVMD NVMe RAID) release package provides support for the creation of RAID volumes that across multiple Intel® NVMe SSD controllers when attached to PCIe express slots managed by the CPU. When installing the Intel® VROC (NonVMD NVMe RAID) package, please take care to ensure that there are no Intel® NVMe SSDs attached to PCI express slots managed by the PCH (please refer to platform documentation). When there are Intel® NVMe SSD attached to the PCH, Intel VROC (NonVMD NVMe RAID) installation process will install the Intel VROC (NonVMD NVMe RAID) driver against those devices. Those devices will show up in the Intel VROC GUI. Intel VROC (NonVMD NVMe RAID) package should not impact the operation these devices in this configuration.

NOTE: Intel® VROC (NonVMD NVMe RAID) does not support spanning NVMe RAID volumes across CPU and PCH PCI express slots.

7.7 Intel VROC (NonVMD NVMe RAID) Feature Limitations

Intel VROC (NonVMD NVMe RAID) release package has the following feature limitations:

- No support for:
 - Hot Plug
 - RAID Volume Boot Support
 - Spanning across PCI Express slots managed by the CPU and those managed by the PCH
 - Third party NVMe SSDs
 - Intel NVMe SSDs plugged into PCIe slots managed by the PCH controller.
- Validated with a maximum of 8 Intel NVMe SSDs
- If used in a RAID volume, all member device must be on a PCI Express slot managed by the CPU
- To fully install Intel VROC (NonVMD NVMe RAID), at least one Intel NVMe SSD must be present



Appendix A: Related Documentation

Relevant Specifications

UEFI Specifications 2.4 (http://uefi.org)
UEFI Platform Initialization Specification version 1.2 (http://www.uefi.org/specsandtesttools)
UEFI Shell Specification version 2.0 (http://www.uefi.org/specsandtesttools)
ATA Command Set 2 (http://www.t13.org/Documents/UploadedDocuments/docs2009/d2015r2-ATAATAPI_Command_Set_-_2_ACS-2.pdf)
ATA8-ACS-8 (http://www.t13.org/Documents/UploadedDocuments/docs2007/D1699r4c-ATA8-ACS.pdf)
SATA 1.0 Specification (http://www.serialata.org)
SATA II Specification (http://www.serialata.org)
SATA 3 (http://www.sata-io.org/documents/SATA-Revision-3.0-Press-Release-FINAL-052609.pdf)
Serial Attached SCSI - 2 (SAS-2) (http://www.t10.org)

Relevant Documentation

CCL / IBL#	Title/Location
558771	Skylake Server Processor External Design Specification (EDS), Volume One
546955	Intel 100 Series Chipset Family Platform Controller Hub (SKL PCH-H) External Design Specification (EDS)
547817	Intel C620 Series Chipset Platform Controller Hub External Design Specification
566449	KabyLake Platform Controller Hub H External Design Specification V2_Rev2_0



Appendix B: Storage System Events

The following table lists storage system events detected by monitor service (IAStorDataMgrSvc):

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



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

NOTES:

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

The following table outlines the events that will be logged and their corresponding Event.

Event ID	Define of message in source code	Severity/Level	Message
4097	START_READ_PATROL_LOG	Informational	Started Patrol on RAID Volume %2.
4098	STOP_READ_PATROL_LOG	Informational	Stopped Patrol on RAID Volume %2.
4099	ERROR_READ_PATROL_LOG	Warning	Read Patrol detected error on RAID Volume %2 at LBA %3 on member disk %4 at PBA %5.
4100	RECOVERY_READ_PATROL_LOG	Informational	Read Patrol detected and recovered bad blocks on the RAID volume %2.
4101	NO_RECOVERY_READ_PATROL_LOG	Error	Read Patrol detected and failed to recover bad blocks on the RAID volume %2.
4102	BBM_INFO_LOG	Informational	Bad Block Management currently found %2 bad blocks and a max capacity for %3 bad blocks on array %4.
4103	DISK_ERROR_SMART_LOG	Error	Error log: Smart event occurred on disk %2.
4104	DISK_ADDED_ON_RUNTIME_LOG	Informational	Disk %2 has been hot plugged.
4105	DISK_REMOVED_ON_RUNTIME_LOG	Informational	Disk %2 has been hot unplugged.
4106	DISK_MARK_AS_SPARE_LOG	Informational	Disk %2 has been marked as spare.
4107	DISK_UNMARK_AS_SPARE_LOG	Informational	Disk %2 is no longer a spare.
4108	DISK_MARK_AS_OFFLINE_LOG	Informational	Disk %2 is offline array member.
4110	DISK_IN_FAILED_STATE_LOG	Error	Disk %2 is in failed state.
4111	DISK_IN_NORMAL_STATE_LOG	Informational	Disk %2 is in normal state.



Event ID	Define of message in source code	Severity/Level	Message
4112	DISK_MANAGE_UNPLUG_LOG	Informational	Disk %2 has been manage unplugged.
4113	CONTROLLER_READ_PATROL_ENABLE_LOG	Informational	Read Patrol has been enabled.
4114	CONTROLLER_READ_PATROL_DISABLE_LOG	Informational	Read Patrol has been disabled.
4115	CONTROLLER_ROHI_ENABLE_LOG	Informational	Rebuild on Hot Insert has been enabled.
4116	CONTROLLER_ROHI_DISABLE_LOG	Informational	Rebuild on Hot Insert has been disabled.
4117	VOLUME_CREATE_LOG	Informational	RAID volume %2 has been created.
4118	VOLUME_DELETE_LOG	Informational	RAID volume %2 has been deleted.
4119	VOLUME_RENAME_LOG	Informational	RAID volume %2 has changed name to %3.
4120	VOLUME_MIGRATION_START_LOG	Informational	Migration of RAID volume %2 has started.
4121	VOLUME_MIGRATION_FINISH_LOG	Informational	Migration of RAID volume %2 has successfully finished.
4122	VOLUME_MIGRATION_FAILED_LOG	Error	Migration of RAID volume %2 has failed.
4123	VOLUME_MIGRATION_SUSPEND_LOG	Informational	Migration of RAID volume %2 is suspended.
4124	VOLUME_MIGRATION_RESUME_LOG	Informational	Migration of RAID volume %2 is resumed.
4125	VOLUME_INITIALIZE_START_LOG	Informational	Initialization of RAID volume %2 has started.
4126	VOLUME_INITIALIZE_FINISH_LOG	Informational	Initialization of RAID volume %2 has successfully finished.
4127	VOLUME_INITIALIZE_FAILED_LOG	Error	Initialization of RAID volume %2 has failed.
4128	VOLUME_INITIALIZE_SUSPEND_LOG	Informational	Initialization of RAID volume %2 is suspended.
4129	VOLUME_INITIALIZE_RESUME_LOG	Informational	Initialization of RAID volume %2 is resumed.
4130	VOLUME_VERIFY_START_LOG	Informational	Verify process on RAID volume %2 has started.
4131	VOLUME_VERIFY_FINISH_LOG	Informational	Verify process on RAID volume %2 has successfully finished.
4132	VOLUME_VERIFY_FAILED_LOG	Error	Verify process on RAID volume %2 has failed.
4133	VOLUME_VERIFY_SUSPEND_LOG	Informational	Verify process on RAID volume %2 is suspended.



Event ID	Define of message in source code	Severity/Level	Message
4134	VOLUME_VERIFY_RESUME_LOG	Informational	Verify process on RAID volume %2 is resumed.
4135	VOLUME_VERIFY_AND_FIX_START_LOG	Informational	Verify and fix process on RAID volume %2 has started.
4136	VOLUME_VERIFY_AND_FIX_FINISH_LOG	Informational	Verify and fix process on RAID volume %2 has successfully finished.
4137	VOLUME_VERIFY_AND_FIX_FAILED_LOG	Error	Verify and fix process on RAID volume %2 has failed.
4138	VOLUME_VERIFY_AND_FIX_SUSPEND_LOG	Informational	Verify and fix process on RAID volume %2 is suspended.
4139	VOLUME_VERIFY_AND_FIX_RESUME_LOG	Informational	Verify and fix process on RAID volume %2 is resumed.
4140	VOLUME_REBUILD_START_LOG	Informational	Rebuild on RAID volume %2 has started.
4141	VOLUME_REBUILD_FINISH_LOG	Informational	Rebuild on RAID volume %2 has successfully finished.
4142	VOLUME_REBUILD_FAILED_LOG	Error	Rebuild failed on RAID volume %2.
4143	VOLUME_REBUILD_SUSPEND_LOG	Informational	Rebuild on RAID volume %2 is suspended.
4144	VOLUME_REBUILD_RESUME_LOG	Informational	Rebuild on RAID volume %2 is resumed.
4145	VOLUME_WBC_ENABLE_LOG	Informational	Write-back cache is enabled on RAID volume %2.
4146	VOLUME_WBC_DISABLE_LOG	Informational	Write-back cache is disabled on RAID volume %2.
4147	VOLUME_STATE_DEGRADED_LOG	Warning	RAID volume %2 is degraded.
4148	VOLUME_STATE_FAILED_LOG	Error	RAID volume %2 is failed.
4149	VOLUME_STATE_NORMAL_LOG	Informational	RAID volume %2 is normal.
4154	VOLUME_REBUILD_FAILED_BBL_FULL_LOG	Error	Rebuild of RAID volume %2 failed, because bad block table is full.
4150	HIBERNATE_SUSPEND_LOG	Informational	Platform is entering hibernation.
4151	SLEEP_SUSPEND_LOG	Informational	Platform is entering standby.
4152	RESUMING_LOG	Informational	Platform is resuming.
4153	POWER_UP_LOG	Informational	Platform is powering up.
4155	IO_FAILED_LOG	Error	IO on %2 has failed.
4156	IO_TIMEOUT_LOG	Error	IO on %2 has reached timeout.



Appendix C: External Hardware Capability

Enterprise SATA Drives

The following table identifies the current list of SATA drives used in validation and is subject to change without notice. Contact the user's factory representative for questions on any specific hardware item.

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

Non-Intel NVMe Drives

The following table identifies the non-Intel NVMe SSD used in validation and is subject to change without notice. Contact the user's factory representative for questions on any specific hardware item.



Vendor	Model	Added Date	Validation Completed	Comment (first release)
Samsung*	SM951	Q4'16	Y	5.1
Samsung*	SM961	Q4'16	Y	5.1
Samsung*	PM961	Q4'16	Y	5.1
Samsung*	PM953	Q1'17	Y	5.1
Toshiba*	XG3	Q4'16	Y	5.1
Micron*	9100 Series	Q4'16	Y	5.1
Lenovo*	Atsani	Q3'17	Y	5.2
Huawei*	ES3600P	Q3'17	Y	5.2 HF
Samsung*	PM963	Q3'17	Y	5.4
Toshiba*	PX04PMB	Q3'17	Y	5.4
Toshiba*	XG5	Q4'17	Y	5.3
Micron*	9200 Series	Q4'18	Y	6.0
Samsung*	PM983	Q4'18	Y	6.0
Western Digital*	SN720	Q4'18	Y	6.0

Note: All shipping Intel Data Center and Professional NVMe* SSDs are supported by Intel® VROC 6.2, except dual port NVMe* SSDs.



Expanders and Enclosures

The following is a list of Switch Vendors Intel has worked with to support Intel VMD LED Management[†].

- Broadcom*
- Microsemi*
- Pericom*
- Semtech*

[†]Contact your specific vendor to confirm the make/model that supports Intel VMD.

§



Appendix D: Intel C620 Series Chipset Legacy OROM Boot Option

INT 15 / AX=F300h / BX=0002h (Get Intel VROC (SATA RAID) OROM Boot Info)

Description

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

Inputs

- AX = F300h (Function)
- BX = 0002h (Sub Function)
- EAX = 0000F300h
- EBX = 4F450002h ('OE') + Sub Function
- EDX = 424F4F54h ('BOOT')

Normal Outputs

- CF = clear if successful
- EAX = 424F4F54h ('BOOT')
- BL = legacy_orm_boot_controller_selection:

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

- 0000h = No runtime space restrictions. BIOS indicates that both VROC (SATA RAID) and sSATA runtime code should provide full int13h support for Intel VROC (SATA RAID) devices.

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

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



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

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

Error Outputs

CF = set on error

AH = error code

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

§



Appendix E: VROC (SATA RAID) Port Bitmap Implementation

Legacy OROM

Intel VROC (SATA RAID) overloaded offset 0x08 of the PNP header, with a bitmap of which ports make up the PNP header. For example, a RAID5 volume whose member disks located on SATA ports 0,2,3 would have a value of 0x0D (0000_1101b) and a pass-thru disk on port 4 would have a value of 0x10 (0001_0000b).

Offset	Size	Value	Description
00h	1	'\$'	Signature byte 1
01h	1	'P'	Signature byte 2
02h	1	'N'	Signature byte 3
03h	1	'P'	Signature byte 4
04h	1	01h	Structure revision
05h	1	Varies	Length (in 16 byte blocks)
06h	2	Varies	Offset of next header (0000h if none)
08h	1	00h	Reserved
09h	1	Varies	Checksum
0Ah	4	Varies	Device identifier
0Eh	2	Varies	Offset to manufacturer string (optional)
10h	2	Varies	Offset to product name string (optional)
12h	3	Varies	Device type code
15h	1	Varies	Device indicators
16h	2	Varies	Boot Connection Vector (BCV)
18h	2	Varies	Disconnect Vector (DV)
1Ah	2	Varies	Bootstrap Entry Vector (BEV)
1Ch	2	0000h	Reserved
1Eh	2	Varies	Static resource information vector



UEFI Driver

The Intel VROC 6.1 UEFI driver, in an effort to provide similar functionality as in the legacy OROM, has implemented the Port Number value in the Device Path as a bitmap representing the physical disk connections that the Logical Disk represents. The LSB (least significant bit) represents port 0 and increases linearly. E.g. a single pass through disk on SATA port 3 (assuming the SATA ports are enumerated 0 – X) is represented by 0000_1000b (or 0x08).

EFI_DEVICE_PATH_PROTOCOL

For each logical disk that is exposed by the SATA RAID UEFI Driver, an EFI_DEVICE_PATH_PROTOCOL shall be created.

The Device Path Protocol for each logical disk shall be appended to the PCI SATA Controller Device Path.

The fields of the EFI_DEVICE_PATH_PROTOCOL shall be filled out differently depending on whether the device is an ODD or an HDD.

EFI_DEVICE_PATH_PROTOCOL Field	ATAPI (ODD)	HDD/Volume – Logical Device
Type	3 (Messaging Device Path)	3 (Messaging Device Path)
Sub-Type	18 (SATA)	18 (SATA)
Length	10	10
HBA Port Number	Port ID bitmap (bit #n set if device is on port #n)	Port ID bitmap (bit #n set if logical device contains device ID #n) Lowest Significant Bit (LSB) represents port 0.
Port Multiplier Port Number	0x8000 (directly connected)	0x8000 (directly connected)
Logical Unit Number	0	0 for pass-through devices, myVolRaidDevNum for RAID volumes, which is the n th volume created on the array.

§



Appendix F: Troubleshooting

This section explains how to resolve the most common problems that may occur while using the application. If the user has any questions regarding installing, using or maintaining this product, the user can also visit Intel's online support site which provides the user with self-help resources and electronic problem submission.

Failed Volumes

RAID 0 A RAID 0 volume is reported as failed when one of its members is disconnected or has failed. In both cases, the volume and its data are no longer accessible.	
Cause	Solution
Missing array disk	<p>Follow this procedure to recover data:</p> <ol style="list-style-type: none">1. Power off the user's computer and reconnect the missing disk.2. Turn on the user's computer. During the system startup, the volume status will display as 'Normal' in the Intel VROC PreOS (UEFI or Legacy Option ROM) user interface.3. Once the operating system is running, open the Intel VROC GUI from the Start menu or click the Intel VROC notification area icon.4. Under 'Status', verify that the volume and disks status display as 'Normal'. The user can also review this information under 'Manage'.
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, the user 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">1. Power off the user's computer and replace the failed NVMe or SATA disk with a new one that is of equal or greater capacity.2. Turn on the user's computer. During the system startup, the volume status will display as 'Failed' in the Intel VROC PreOS (UEFI or Legacy Option ROM) user interface.3. To access the Intel VROC UEFI HII user interface, enter into the BIOS Setup environment and navigate to the Intel VROC HII user interface menu option4. To access the Intel VROC (SATA RAID) Legacy Option ROM, press Ctrl-I to access the main menu of the option ROM user interface.5. Select Delete RAID Volume from the main menu.6. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.7. Press the 'Delete' key to delete the volume, then 'Y' to confirm.8. Create a new RAID 0 volume using the new disk. If the failed disk was part of the system volume, the user will also need to reinstall the operating system.



RAID 5 A RAID 5 volume is reported as failed when two or more of its members have failed.	
Cause	Solution
Two or more array disks failed	<p>In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, the user 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"> 1. Power off the user's computer and replace the failed NVMe or SATA disks with new ones that are of equal or greater capacity. 2. Turn on the user's computer. During the system startup, the volume status will display as 'Failed' in the Intel VROC PreOS (UEFI or Legacy Option ROM) user interface. 3. To access the Intel VROC UEFI HII user interface, enter into the BIOS Setup environment and navigate to the Intel VROC HII user interface menu option 4. To access the Intel VROC (SATA RAID) Legacy Option ROM, press Ctrl-I to access the main menu of the option ROM user interface. 5. Select Delete RAID Volume from the main menu. 6. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 7. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 8. Create a new RAID 5 volume using the new disks. If the failed disk was part of the system volume, the user will also need to reinstall the operating system.
RAID 10 A RAID 10 volume is reported as failed when two adjacent members are disconnected or have failed, or when three or four of its members are disconnected or have failed.	
Cause	Solution
Two adjacent array disks missing (visual inspection)	<ol style="list-style-type: none"> 1. Power off the user's computer and reconnect the missing disks. 2. The rebuild operation will start automatically. The user can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'.
Three or four array disks missing	<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"> 1. Power off the user's computer and reconnect the missing disks. 2. Turn on the user's computer. During the system startup, the volume status will display as 'Failed' in the Intel VROC PreOS (UEFI or Legacy Option ROM) user interface. 3. To access the Intel VROC UEFI HII user interface, enter into the BIOS Setup environment and navigate to the Intel VROC HII user interface menu option 4. To access the Intel VROC (SATA RAID) Legacy Option ROM, press Ctrl-I to access the main menu of the option ROM user interface. 5. Select Delete RAID Volume from the main menu. 6. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 7. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 8. Create a new RAID 10 volume using the new disks. 9. The user will then need to reinstall the operating system on the new volume.



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, the user 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">1. Power off the user's computer and replace the failed NVMe or SATA disks with new ones that are of equal or greater capacity.2. Turn on the user's computer. During the system startup, the volume status will display as 'Failed' in the Intel VROC PreOS (UEFI or Legacy Option ROM) user interface.3. To access the Intel VROC UEFI HII user interface, enter into the BIOS Setup environment and navigate to the Intel VROC HII user interface menu option4. To access the Intel VROC (SATA RAID) Legacy Option ROM, press Ctrl-I to access the main menu of the option ROM user interface.5. Select Delete RAID Volume from the main menu.6. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys.7. Press the 'Delete' key to delete the volume, then 'Y' to confirm.8. Create a new RAID 10 volume using the new disks.9. The user will then need to reinstall the operating system on the new volume.
--------------------------------	--

Degraded Volumes

RAID 1

A RAID 1 volume is reported as degraded when one of its members is disconnected or has failed. Data mirroring and redundancy are lost because the system can only use the functional member.

RAID 5

A RAID 5 volume is reported as degraded when one of its members is disconnected or has failed. When two or more array disks are disconnected or have failed, the volume is reported as failed.

RAID 10

A RAID 10 volume is reported as degraded when one of its members is disconnected or has failed, or when two non-adjacent members are disconnected or have failed. When two or more array disks are disconnected or have failed, the volume is reported as failed.



Cause	Solution
Missing array disk	<p>If the user can reconnect the missing disk, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> 1. Power off the user's computer and reconnect the missing disk. 2. Turn on the user's computer and the rebuild operation will start automatically. <p>If the user cannot reconnect the missing disk and a NVMe or SATA disk is available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> 1. If a NVMe or SATA disk is compatible, available and normal, follow this procedure to rebuild the volume: 2. Select the disk the user wants to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. The user can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. <p>Note: If there is no available disk present, the user will need to power off the user's computer and connect a new NVMe or SATA disk that is equal or greater capacity than the failed disk. Once the user's computer is back up and running the user can follow the rebuild procedure described above.</p>
Failed array disk	<p>We recommend that the user rebuild the degraded volume to a new disk to return the volume and overall storage system status to normal. However, the user can try resetting the disk to normal, which will prompt the volume to start rebuilding automatically. But if the read/write data access consistently fails, the disk will likely return to a failed state immediately and the user will need to rebuild the volume to another disk.</p> <p>If a NVMe or SATA disk is compatible, available and normal, follow this procedure to rebuild the volume:</p> <ol style="list-style-type: none"> 1. Under 'Status', click 'Rebuild to another disk'. 2. Select the disk the user wants to use to rebuild the volume, and then click 'Rebuild'. 3. The rebuild operation starts immediately. The user can follow the progress by hovering over the notification area icon or by reviewing the volume status under 'Status' or 'Manage'. 4. Once the operation successfully completed, the array disk and volume status will display as 'Normal'. <p>Note: If there is no available disk present, the user will need to power off the user's computer and connect a new NVMe or SATA disk that is equal or greater capacity than the failed disk. Once the user's computer is back up and running the user can follow the rebuild procedure described above.</p>



Other Volume States


Incompatible	
Cause	Solution
1) Indicates that the volume was moved to another system that does not support the volume type and configuration.	<p>In this situation, volume data is accessible to the operating system and can be backed up, but the volume cannot operate because the user's system does not support its RAID configuration. Here are the user's options:</p> <ul style="list-style-type: none"> • Reconnect the volume to the computer where the volume was originally created, and continue using it. • Delete the volume, and then create a new volume with a RAID configuration that is supported by the current system. Follow the procedure described above to delete the volume.



	WARNING: When a volume is deleted, all existing data on the member disks of the selected volume is permanently erased. It's recommended that the user backup all valuable data prior to beginning this action.
2) Indicates that the Intel VROC Upgrade Key is incorrect or missing.	In this situation, volume data may not be accessible to the operating system. Here are the user's options: <ul style="list-style-type: none">• Install the proper Intel VROC Upgrade Key
Unknown	
Cause	Solution
The volume is in an unexpected state due to a configuration error.	The application is unable to detect the exact nature of the problem. Try restarting the user's computer. If the error persists, back up all valuable data and delete the volume using the option ROM user interface. Refer to the user's manual accessible from the Online Support area for details on using the option ROM.
Missing volume	
Cause	Solution
A driver upgrade or downgrade was performed while a data migration was in progress.	<p>The driver cannot recognize the volume or read its data if a driver upgrade or downgrade was performed during a volume migration. Volume migrations occur after one of the following operations was initiated:</p> <ol style="list-style-type: none">1. Creation of a system volume or data volume while preserving data.2. Volume type change combined with disk addition to the new RAID configuration.3. Volume size increase.4. Disk addition to an existing array. <p>Troubleshooting a data volume</p> <ol style="list-style-type: none">1. If the data migration involved a data volume, the user will need to reverse the driver upgrade or downgrade operation and return to the original driver version. This will restore driver and volume compatibility.2. Once the operation has completed, restart the user's computer.3. Open the application and make sure that the volume displays again in the storage system view. Data migration operation should resume immediately. <p>Troubleshooting a system disk</p> <p>If the data migration involved a system disk or volume, it is highly likely that the user will not be able to start the user's system because the driver cannot read the system files. The following options may allow the user to load the operating system again:</p> <ol style="list-style-type: none">1. Restore to the last known good configuration.2. Boot from a flash drive that supports NTFS partitioning and includes the storage driver files.3. Bring the corrupt disk to another system, and then replace the storage driver files from a compatible driver version. Return the disk to the original system and try booting. <p>Troubleshooting a system volume</p> <p>If the data migration involved a system disk or volume, it is highly likely that the user will not be able to start the user's system because the driver cannot read the system files. The following options may allow the user to load the operating system again:</p> <ol style="list-style-type: none">1. Restore the last known good configuration.2. Bring all corrupted volume disks to another system, and then replace the storage driver files from a compatible driver version. Return the corrupted volume disks to the original system and try booting.

Disk Events

State	Cause	Solution
<p>At risk</p> 	<p>An impending error condition was detected on an internal or external disk and is now at risk of failure.</p>	<p>The application is detecting early warning signs of failure with a NVMe or SATA disk that result from a slow degradation over time. When a disk is reported at risk, the user can reset that disk to normal, but we recommend that the user contacts 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"> 1. Under 'Status', in the Manage subsection, locate the disk reported as at risk. The user 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. <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> <p>If the disk reported at risk is included in a RAID volume and a compatible spare disk is available, the rebuild process will start automatically. Once complete, the disk reported at risk becomes available and the user can reset it to normal to return to a healthy state.</p>
	<p>An unexpected error was detected on a disk that has RAID configuration data (metadata) on it.</p>	<p>In this state, it is likely that some or all of the disk data is inaccessible. After backing up any accessible data, the user 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"> 1. Under 'Status', in the Manage subsection, locate the disk reported as at risk. The user can also perform this action from Manage Disk, which is accessible by clicking the disk in the storage system view. 2. Click 'Clear and reset disk', and then click 'Yes' to confirm. 3. Once complete, the page refreshes with the disk returning to a normal state.
<p>Missing</p> 	<p>An array disk is not present or physically connected to the computer.⁹</p>	<p>Ensure that the disk is securely connected to the NVMe or SATA port and that the data cable is functioning properly. If the disk is lost or cannot be reconnected, the user will need to connect a new NVMe or 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.</p>

<p>Failed</p> 	<p>An internal or external disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.</p>	<p>Back up the user's data and we recommend that the user 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.</p> <p>In a failed state, disk data may be lost, but the user 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"> 1. Under 'Status', in the Manage subsection, locate the disk reported as failed. The user 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. <p>Note: If the failed array disk is part of a redundant volume, the volume will start rebuilding automatically as soon as the disk is reset to normal.</p>
---	--	---



Software Errors

Message	Cause	Solution
An unknown error occurred while running this application. If the problem persists, please restart the user's computer or try reinstalling the application.	This error may be related to: <ol style="list-style-type: none"> 1. Missing components 2. Corrupted application 3. Application unable to connect to the service 4. Application fails to start. 	Restart the user's computer or try reinstalling the application.
Intel® Rapid Storage Technology enterprise is trying to connect to the service.	The application is launched and is attempting to connect to the service in order to run.	If the connection succeeds, the application opens and is fully functional; if the connection fails, the error message described above is displayed. Try starting the service manually using Microsoft® Windows® Services, or follow the recommended solution listed above to resolve the problem.
The Intel® Rapid Storage Technology enterprise service cannot be started in safe mode.	The user's computer was started in safe mode and the operating system is running with a limited set of files and drivers. Intel Rapid Storage Technology enterprise cannot start or run in safe mode.	Once the user is done troubleshooting application or driver problems in safe mode, the user will need to exit safe mode, restart the user's computer, and then let the operating system start normally. The Intel Rapid Storage Technology enterprise service can now be started and open the application.
Multiple users cannot run the application at the same time.	One or more users are attempting to open the application while an instance of the application is already running.	Make sure only one instance of the application is running at a time.
An error occurred due to insufficient resources, and the operation could not be completed. Please try again later.	The Intel® Rapid Storage Technology enterprise driver does not have sufficient resources to execute the request. Another operation may be in progress and needs to complete before being able to handle a new request.	Wait a few moments, then try performing the action again.
An unknown error occurred during the volume creation process. Please try recreating the volume.	An unexpected error occurred during the operation, and the application cannot identify its origin. The volume could not be created.	Verify that the user's 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 the user's computer and then the operation.