



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

Technical Product Specification for Intel® VROC 9.3PV

Revision 027

October 2025



You may not use or facilitate the use of this document in connection with any infringement or other legal analysis. You may not use or facilitate the use of this document in connection with any infringement or other legal analysis concerning Intel products described herein. You agree to grant Intel a non-exclusive, royalty-free license to any patent claim thereafter drafted which includes subject matter disclosed herein.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

All product plans and roadmaps are subject to change without notice.

The products described 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.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer or learn more at intel.com.

Intel disclaims all express and implied warranties, including without limitation, the implied warranties of merchantability, fitness for a particular purpose, and non-infringement, as well as any warranty arising from course of performance, course of dealing, or usage in trade.

Altering clock frequency, voltage, or memory interface speeds may void any product warranties and reduce stability, security, performance, and life of the processor and other components. Intel has not validated processor running memory above Plan-Of-Record (POR) speed. DRAM/DIMM devices should support desired speed, check with DRAM/DIMM vendors for details. System manufacturers are responsible for all validation and assume the risk of any stability, security, performance, or other functional issues resulting from such alterations.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

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

Copyright© 2025 Intel Corporation. All rights reserved.

Contents

1	Introduction	13
	1.1 Intended Audience	13
	1.2 Reference OEM Platform Documentation	13
	1.3 Terminology	13
2	Product Overview	17
	2.1 Product Release Numbering Scheme	17
	2.2 Supported Operating Systems per Platform	18
	2.3 Supported Platforms, Chipsets and SKUs	18
3	New Features	20
	3.1 New Features Introduced with Intel® VROC 9.3 Release	20
	3.2 New Features Introduced with Intel® VROC 9.2 Release	20
	3.3 New Features Introduced with Intel® VROC 9.1 Release	20
	3.4 New Features Introduced with Intel® VROC 9.0 Release	20
	3.4.1 Added SATA RAID Support on Montage IOH	20
	3.4.2 OOB for NVMe SSDs behind VMD via MCTP	20
	3.4.3 NPEM Support for CPU-Attached SSD	20
	3.5 New Features Introduced with Intel® VROC 8.6 Release	21
	3.5.1 Windows Server 2025 Enterprise	21
	3.5.2 Windows 11 24H2	21
	3.5.3 Sanitize Command	21
	3.5.4 Storage System Event Changes	21
	3.6 New Features Introduced with Intel® VROC 8.5 Release	21
	3.6.1 Intel® On Demand Licensing Activation	21
	3.6.2 VROC RAID 1 Only SKU	21
	3.7 New Features Introduced with Intel® VROC 8.2 Release	22
	3.7.1 Support for Pre-Boot DMA	22
	3.7.2 UEFI PEI Module driver	22
	3.8 New Features Introduced with Intel® VROC 8.0 Release	23
	3.8.1 Intel® VROC Out-Of-Band Management (Server Platforms Only)	23
	3.8.2 LED Management on PCH-VMD	24
	3.8.3 LED Patterns Conflict Precedence	24
	3.8.4 Intel® VROC 8.0 New GUI Design	24
	3.8.4.1 Warning Message Displayed when GUI and Driver Version is Mismatching	24
	3.8.5 Physical Drive Location Information in UEFI	25
	3.8.6 Third SATA Controller (Server Platform Only)	25
	3.9 New Features Introduced with Intel® VROC 7.8 Release	25
	3.10 New Features Introduced with Intel® VROC 7.7 Release	25
	3.10.1 UEFI Firmware Management Protocol Support	25
	3.10.2 Intel® VROC (SATA RAID) Support of EFI_ATA_PASS_THRU_Protocol	26
	3.10.3 Disable Locate LED within Intel® VROC LED Management	26
	3.11 Features Introduced with Intel® VROC 7.6 Release	26
	3.11.1 Intel® VROC Pre-OS Environment RAID Volume Failure Recovery ...	27
	3.11.2 Intel® VROC (SATA RAID) UEFI Support for EFI_ATA_PASSTHRU ..	27
	3.12 Features Introduced with Intel® VROC 7.5 Release	27
	3.12.1 Intel® VMD 2.0 Features	28



- 3.12.1.1 Increasing MSIX Vectors to 64.....28
 - 3.12.1.2 Intel® VMD (PCH) Support28
 - 3.12.2 Number of Intel® VMD Increased30
 - 3.12.3 Native PCIe Enclosure Management (NPEM)30
 - 3.12.4 Limited Self Encrypted Drives.....30
 - 3.12.5 Intel® VROC 7.5 GUI UWD Upgrade Limitations30
 - 3.12.6 Intel® VROC 7.5 CLI31
 - 3.12.7 Unsupported Drives31
 - 3.13 Features Introduced with Intel® VROC 7.0 Release.....31
 - 3.13.1 Removal of AHCI Mode Support.....31
 - 3.13.2 Removal of Legacy Option ROM Support.....31
 - 3.13.3 Removal of Windows* 7 Support31
- 4 Intel® VROC Package Components.....32
 - 4.1 Intel® VROC CLI32
 - 4.2 Intel® VROC OS Installer32
 - 4.3 Intel® VROC (VMD NVMe RAID) Environment33
 - 4.3.1 Intel® VROC (VMD NVMe RAID) Pre-OS Components.....33
 - 4.3.2 Intel® VROC (VMD NVMe RAID) F6 (or “Load Driver”) Drivers.....35
 - 4.4 Intel® VROC (SATA RAID) Environment35
 - 4.4.1 Intel® VROC (SATA RAID) Pre-OS Components.....35
 - 4.4.1.1 EFI Support for the SATA/sSATA Controller35
 - 4.4.2 Intel® VROC (SATA RAID) F6 (or “Load Driver”) Drivers.....37
 - 4.4.2.1 Intel® VROC (SATA RAID) F6 (or “Load Driver”) Limitation37
- 5 Intel® VROC Key Features40
 - 5.1 Intel® Volume Management Device (Intel® VMD)40
 - 5.1.1 Multiple Intel® VMD Controllers41
 - 5.1.2 Intel® VMD Method of LED Management41
 - 5.1.3 Intel® VMD Error Management42
 - 5.1.4 Intel® VMD Enabling43
 - 5.1.5 Intel® VMD Surprise Hot-Plug of NVMe Devices.....44
 - 5.2 Intel® VROC (VMD NVMe RAID) Pre-OS Features and Functionality.....44
 - 5.2.1 Intel® VROC (VMD NVMe RAID) for UEFI.....45
 - 5.2.2 Intel® VROC (VMD NVMe RAID) UEFI Human Interface Infrastructure (HII).....45
 - 5.2.3 Intel® VROC (VMD NVMe RAID) UEFI Maximum Number of Drives..45
 - 5.2.4 Using the Intel® VROC (VMD NVMe RAID) HII to Create a RAID Volume46
 - 5.2.5 NVM Express Pass-Thru Protocol Support.....47
 - 5.2.6 Intel® VROC (VMD NVMe RAID) IOCTL Support48
 - 5.2.7 Intel® VROC (VMD NVMe RAID) Private Device/Member Info Protocol49
 - 5.2.7.1 Reference Documents.....49
 - 5.3 Intel® VROC Monitoring and Management50
 - 5.3.1 Intel® VROC Graphical User Interface (GUI)50
 - 5.4 Intel® VROC Features and Functionality50
 - 5.4.1 Intel® VROC RAID Modes of Operation for Upgrade Hardware Keys for RAID Support.....50
 - 5.4.2 Intel® VROC (VMD NVMe RAID) Volumes on a Single Intel® VMD Domain53
 - 5.4.3 Spanning Intel® VROC (VMD NVMe RAID) RAID Data Volumes across Intel® VMD Domains53

5.4.3.1	Intel® VROC (VMD NVMe RAID) on Pass-Thru Boot Drives	54
5.4.3.2	Intel® VROC (VMD NVMe RAID) Spanned RAID Volumes across Intel® VMD Controllers.....	54
5.4.4	Intel® VROC (VMD NVMe RAID) Driver Upgrade or Downgrade	56
5.4.5	Intel® VROC (VMD NVMe RAID) Maximum Number of Drives Supported in a RAID Volume	57
5.4.6	Intel® VROC (VMD NVMe RAID) Maximum Number of Switch Levels Supported	57
5.4.7	Intel® VROC (VMD NVMe RAID) Maximum Number of Arrays Supported	57
5.4.8	Intel® VROC (VMD NVMe RAID) Maximum RAID Members Supported	57
5.4.9	NVMe Hot Insert.....	57
5.4.10	Hot Plug Limitation in Windows*	57
5.4.11	Intel® VROC Out-Of-Band Management (Server Platforms Only)	58
5.4.12	Intel® NVMe Wear Leveling Recommendations.....	59
5.4.12.1	Intel® VROC Optimized Wear Leveling	60
5.4.12.2	Intel® VROC RAID 5 Default Settings	60
5.4.13	Intel® VROC (VMD NVMe RAID) 90-Day Trial Period	61
5.4.13.1	Intel® VROC (VMD NVMe RAID) Trial Period Recommendations and Limitations	62
5.4.13.2	Intel® VROC (VMD NVMe RAID) Trial Period Usage	62
5.4.14	Intel® VROC (VMD NVMe RAID) Pre-Load Installation	65
5.4.15	Intel® VROC Manual Tuning of CPU Utilization.....	65
6	Intel® VROC Common Features	67
6.1	Intel® VROC RAID Features and Functionality	67
6.1.1	RAID Levels	67
6.1.2	Matrix RAID	69
6.1.3	Creating a RAID Volume	69
6.1.3.1	Using RAID Configuration Utilities (DOS, UEFI Shell, and Windows*).....	70
6.1.3.2	Using the Intel® Virtual RAID on CPU GUI	70
6.1.4	Instant Initialization	71
6.1.4.1	Intel® VROC Automatic Initialization	72
6.1.4.2	Intel® VROC RAID Volume Initialization.....	73
6.1.5	Deleting a RAID Volume	73
6.1.5.1	Using the Intel® VROC GUI.....	73
6.1.5.2	Using the Intel® VROC UEFI HII User Interface	74
6.1.5.3	Using the Intel® VROC CLI Tool	74
6.1.6	Degraded RAID Volumes	74
6.1.7	Intel® VROC GUI Support of Various Controllers	74
6.1.8	Various Strip Sizes.....	75
6.1.8.1	Setting the Strip Size When Creating a Volume	76
6.1.8.2	Setting the Strip Size When Changing Volume Type.....	76
6.1.9	Disk Coercion.....	76
6.1.10	Minimum RAID Volume Size Recommendation	76
6.1.11	Flush Unit Access (FUA).....	77
6.1.12	RAID Volume Roaming	77
6.1.13	RAID Volume Roaming between Linux* and Windows*	77
6.1.14	Support of Pass-Through Drive	77
6.1.15	Failed Drive Reinsertion.....	78
6.1.16	LED Management	78
6.1.16.1	Failed Drive Indicator	81

6.1.16.2	Activate LED	81
6.1.16.3	Rebuild LED	82
6.1.16.4	LED Patterns Conflict Precedence	82
6.1.16.5	LED Management on PCH-VMD	82
6.1.16.6	UEFI LED Management	82
6.1.17	Modify Volume Name	82
6.1.18	Hot Plug (Surprise and Managed).....	82
6.1.18.1	How to Enable the BIOS for Hot Plug	83
6.1.18.2	Intel® VROC Pre-OS Hot Plug Support.....	84
6.1.19	Hot Spare Disk.....	84
6.1.19.1	Global Hot Spare	85
6.1.19.2	Mark a Disk as Hot Spare.....	85
6.1.20	Drive Requirements	87
6.1.20.1	Clean Drive Status	87
6.1.20.2	Minimum Disk Size Requirements.....	87
6.1.20.3	Disk Addition: Minimum Required Disk Size	87
6.1.20.4	Targeted Disk Replacement: Minimum Required Disk Size	87
6.1.21	Intel® VROC Rebuild	88
6.1.21.1	Intel® VROC Rebuild Components	88
6.1.21.2	Auto Rebuild	88
6.1.21.3	Manually Invoked Rebuild	90
6.1.22	Verify and Repair.....	91
6.1.23	Starting Verify and Repair Process	92
6.1.23.1	Dirty Shutdown	92
6.1.23.2	Manual Initiation.....	92
6.1.24	Read Patrol.....	93
6.1.25	Check Pointing	93
6.1.26	Bad Block Management	93
6.1.26.1	Marking a Bad Block.....	94
6.1.26.2	Full Bad Block Table Rebuild.....	94
6.1.26.3	Degraded Read.....	94
6.1.26.4	Read Patrol Bad Block Management.....	94
6.1.26.5	Multiple Bad Blocks Intel® VROC (SATA RAID) Read Patrol	94
6.1.26.6	Bad Block Management Logging.....	94
6.1.26.7	Bad Block Management Recovery	95
6.1.26.8	Intel® VROC GUI – Bad Block Reporting	95
6.1.27	Intel® VROC RAID Write Hole Closure.....	95
6.1.27.1	RAID Write Hole Closure (RWH) Details.....	96
6.1.27.2	RWH Recovery.....	97
6.1.28	Controller Default Values	97
6.1.28.1	Setting Default Rebuild on Hot Insert (ROHI).....	97
6.1.28.2	Setting Default Read Patrol	98
6.1.29	Driver Dirty Shutdown (DS) Recovery	98
6.1.29.1	Driver DS Recovery Caused by Hot-Plug.....	98
6.1.29.2	Driver DS Recovery Caused by Driver Start	98
6.1.29.3	Driver DS Recovery Caused by Enabling Drives	98
6.1.29.4	DS Condition and Hibernation Process	98
6.1.29.5	Interrupted PPL Write - DS Recovery	99
6.1.29.6	Failed DS Recovery - Verify Volume.....	99
6.1.29.7	Failed DS Recovery - Initialize Volume	99
6.1.29.8	Dirty Stripe Journaling.....	99
6.1.29.9	Partial Parity Logging (PPL)	99
6.1.29.10	Reading/Changing Controller Parameters	99

6.2	Intel® VROC Pre-OS Features and Functionality	100
6.2.1	Intel® VROC Unified Extensible Firmware Interface (UEFI) Drivers	100
6.2.1.1	Specification References	100
6.2.1.2	Intel® VROC UEFI User Interface.....	100
6.2.2	UEFI System BIOS Requirements for Platform Compatibility with Intel® VROC UEFI	101
6.2.2.1	Required Protocols/Functions to be Provided by the UEFI System BIOS.....	101
6.2.2.2	Optional Protocols/Functions to be Provided by the UEFI System BIOS.....	102
6.2.2.3	Protocols Provided by the Intel® VROC UEFI Drivers	102
6.2.2.4	How-to-Enable the Platform with Intel® VROC UEFI Driver/HII_GUI.....	103
6.2.2.5	Known Compatibility Issues with the UEFI Self Certification Test (UEFI SCT) Tool	104
6.2.3	UEFI Health Protocol	105
6.2.4	Determining the Version of the Intel® VROC UEFI Driver	106
6.2.4.1	Using the UEFI Shell.....	106
6.2.4.2	Intel® VROC (SATA RAID) Device Information Display in UEFI	106
6.2.4.3	Intel® VROC (SATA RAID) UEFI Boot Volume Configuration	106
6.2.5	Pre-OS Installation of the Intel® VROC Driver	108
6.2.5.1	Pre-OS Driver Installation Using the “Load Driver” Method	108
6.2.5.2	Selectable Boot Volume	108
6.2.6	Pre-OS RAID Failure Recovery Process	108
6.2.7	UEFI Firmware Management Protocol Support	114
6.2.7.1	GetImageInfo for NVMe SSDs.....	114
6.2.7.2	Set Image for NVMe SSDs.....	116
6.2.7.3	GetImageInfo for SATA SSDs	116
6.2.7.4	SetImage for SATA SSDs	117
6.2.7.5	Firmware Update Progress	117
6.2.8	Physical Drive Location Information in UEFI.....	118
6.2.9	Multiple Namespaces in NVMe SSDs	118
6.2.10	Disable Intel® VROC LED Management	118
6.2.10.1	VrocUefiConfig Spec to Disable Intel® VROC LED Management	118
6.3	Intel® VROC Monitoring and Management	119
6.3.1	Intel® VROC GUI	119
6.3.1.1	Intel® VROC Installation	120
6.3.1.2	Intel® VROC (SATA RAID) Driver.....	121
6.3.1.3	Intel® VROC (VMD NVMe RAID) Driver.....	121
6.3.1.4	Intel® VROC (NonVMD NVMe RAID) Driver	121
6.3.2	Getting Started	121
6.3.2.1	RAID Enabled Systems	121
6.3.3	Component Version Detection	121
6.3.4	Using System Report	121
6.3.5	Notification Area.....	123
6.3.5.1	Opening the Application from the Notification Area	124
6.3.5.2	Understanding the Application	125
6.3.5.3	Storage System Status.....	130
6.3.5.4	Volume States.....	131
6.3.5.5	RAID-Ready Mode.....	132

	6.3.5.6	RAID-Ready Migration	132
	6.3.5.7	Intel® VROC Graphical User Interface (GUI)	132
	6.3.5.8	Intel® VROC Graphical User Interface (GUI) Version	133
	6.3.5.9	Intel® VROC GUI Online Support.....	134
6.3.6		Managing Arrays	134
	6.3.6.1	Array Properties.....	134
	6.3.6.2	Adding a Disk to an Array	135
	6.3.6.3	Package Upgrade or Downgrade While Migrating	135
	6.3.6.4	Migration of a System Disk to a RAID Volume	135
	6.3.6.5	Creating a Volume	136
	6.3.6.6	Increasing Volume Size	142
	6.3.6.7	Enabling Disk Data Cache	143
	6.3.6.8	Managing Volumes.....	144
	6.3.6.9	Volume Properties.....	144
	6.3.6.10	Renaming a Volume	145
	6.3.6.11	Rebuilding a Volume.....	145
	6.3.6.12	Resetting Volume to Normal	146
	6.3.6.13	Changing Volume Type	147
	6.3.6.14	Increasing Volume Size	147
	6.3.6.15	Adding a Disk to a Volume	149
	6.3.6.16	Enabling Volume Write-Back Cache	150
	6.3.6.17	Initializing a Volume.....	150
	6.3.6.18	Verifying and Repairing Data	151
	6.3.6.19	RAID Level Migrations	151
	6.3.6.20	Migrating From One RAID Level to Another.....	152
6.3.7		Managing Disks	153
	6.3.7.1	Disk Properties	153
	6.3.7.2	Marking a Disk as Spare	155
	6.3.7.3	Resetting a Disk to Available	156
	6.3.7.4	Resetting a Disk to Normal.....	156
	6.3.7.5	Resetting a Disk to Online/Normal.....	156
	6.3.7.6	Connecting a Disk	157
	6.3.7.7	Managing Ports.....	158
	6.3.7.8	Detecting a Spare Drive	158
6.3.8		Disk Monitor Service	158
6.3.9		NVMe Deallocate/SATA TRIM Command (RAID 0, 1 and 10)	158
6.3.10		Un-Installation	159
	6.3.10.1	Uninstalling the Intel® VROC Software (Except the RAID Driver)	159
	6.3.10.2	Disabling the RAID Driver by Disabling the RAID Controller	159
6.4		Intel® VROC Power Management	160
	6.4.1	Intel® VROC (VMD NVMe RAID) Power Management Hot Plug Limitation	160
6.5		Intel® VROC Miscellaneous Features and Functionality	160
	6.5.1	OS Installation	160
	6.5.2	Using the Device Manager	160
	6.5.3	Disk Write Cache	162
	6.5.4	Write Back Cache	162
	6.5.5	4K Native Sector Sizes Drives	162
	6.5.6	Intel® VROC Microsoft* WHQL Drivers	163
	6.5.7	Intel® VROC Microsoft* UWD Support.....	163
	6.5.7.1	INF Installation Support.....	163
	6.5.7.2	Windows* Application Store Support	163

	6.5.7.3	Stand Alone Intel® VROC GUI	163
	6.5.8	Microsoft* Hyper-V*	165
	6.5.9	Intel® SMART Event Support	165
	6.5.10	Event Log Support	166
	6.5.11	2 Terabyte (TB) Disk Size	166
	6.5.12	2 Terabyte (TB) Volume Size.....	166
	6.5.13	Email Alerting and Notification.....	166
	6.5.13.1	Email Notification UI Visible Enable/Disable	166
	6.5.14	Microsoft* .NET Framework	168
6.6		Intel® VROC Package Updates	169
	6.6.1	Intel® VROC Package Upgrade.....	169
	6.6.2	Intel® VROC Package Downgrade.....	170
7		Intel® VROC (SATA RAID) Key Features	172
	7.1	Intel® VROC (SATA RAID) Pre-OS Features and Functionality.....	172
	7.1.1	Using the BCFS to Differentiate Platform SKUs.....	172
	7.1.1.1	Configuring the Platform’s RAID Related Features.....	173
	7.1.1.2	Intel® VROC (SATA RAID) UEFI Dirty Shutdown Recovery	174
	7.1.2	ATAPI on SATA/sSATA/tSATA Controller	174
	7.1.3	Support for EFI_ATA_PASS_THRU_PROTOCOL.....	174
	7.1.4	EFI_ATA_PASS_THRU_PROTOCOL Supported Functions	175
	7.2	Intel® VROC (SATA RAID) Miscellaneous Features and Functionality	176
	7.2.1	SGPIO on SATA Controller (RAID Mode Only)	176
	7.2.2	Common Storage Management Interface (CSMI)	176
	7.2.3	NCQ (SATA/sSATA/tSATA)	176
	7.2.4	Intel® VROC (SATA RAID) Maximum Number of SATA Drives in a RAID Volume	176
	7.2.5	Intel® VROC (SATA RAID) Maximum Number of sSATA/tSATA drives in a RAID Volume	177
	7.2.6	Intel® VROC (SATA RAID) Maximum Number of RAID Volumes	177
	7.2.7	SCSI Pass-Through	177
	7.2.8	Intel® VROC (SATA RAID) UEFI Support of for EFI_ATA_PASSTHRU	177
	7.2.9	RAID Volume Read Cache	178
	7.2.10	Intel® VROC (SATA RAID) Support for Staggered Spin-Up.....	178
	7.2.11	Intel® VROC (SATA RAID) Driver Downgrade.....	178
8		Intel® VROC (NonVMD NVMe RAID) Key Features	178
	8.1	PCIe Intel® NVMe Device Usage Model.....	178
	8.2	Intel® VROC (NonVMD NVMe RAID) Boot Support	178
	8.3	Intel® VROC (NonVMD NVMe RAID) Installation	178
	8.4	Intel® VROC (NonVMD NVMe RAID) Pass-Through IOCTL Support	178
	8.5	Intel® NVMe SSDs Support List for Intel® VROC (NonVMD NVMe RAID).....	178
	8.6	Intel® VROC (NonVMD NVMe RAID) RAID Volume Support with Multiple Controllers.....	178
	8.7	Intel® VROC (NonVMD NVMe RAID) Feature Limitations	178
Appendix A		Related Documentation	179
	A.1	Relevant Specifications.....	179
Appendix B		Storage System Events	180
Appendix C		External Hardware Capability	185



	C.1	Non-Intel® NVMe Drives	185
	C.2	Expanders and Enclosures	185
Appendix D		Intel® VROC (SATA RAID) Port Bitmap Implementation	186
	D.1	UEFI Driver.....	186
	D.2	EFI_DEVICE_PATH_PROTOCOL.....	186
Appendix E		Troubleshooting	187
	E.1	Failed Volumes.....	187
	E.2	Degraded Volumes.....	189
	E.3	Other Volume States.....	191
	E.4	Disk Events.....	192
	E.5	Drive Information Structure Drive Status Values	194
	E.6	Software Errors	194

Figures

Figure 3-1.	Overview of Intel® VROC Out-Of-Band Management Solution	24
Figure 3-2.	Physical Drive Location Information in UEFI	25
Figure 3-3.	Firmware Management Protocol Support Flow	26
Figure 5-1.	Multiple Intel® VMD Controllers	41
Figure 5-2.	Enabling Intel® VMD	44
Figure 5-3.	Intel® VMD Domains.....	54
Figure 5-4.	Overview of Intel® VROC Out-Of-Band Management Solution	58
Figure 5-5.	Intel® VROC Trial Period Message.....	63
Figure 5-6.	Intel® VROC Trial Period Expired Message	64
Figure 5-7.	Intel® VROC Normal Operation After Upgrade Key Installation.....	65
Figure 6-1.	Intel® VROC RAID Write Hold Closure Example	95
Figure 6-2.	Example of a RAID 10 Configuration	109
Figure 6-3.	Intel® VROC RAID 10 Volume in 'Normal' State in the HII	110
Figure 6-4.	Intel® VROC RAID 10 Volume in 'Failed' State in the HII.....	110
Figure 6-5.	Intel® VROC RAID 10 Volume in 'Failed' State after Reboot in the HII... ..	111
Figure 6-6.	Intel® VROC RAID Recovery Options in the HII.....	111
Figure 6-7.	Intel® VROC 'Reset to Degraded' Option in the HII.....	112
Figure 6-8.	Intel® VROC Recovery Confirmation Option in the HII	113
Figure 6-9.	Intel® VROC RAID 10 Volume in Rebuilding Process in the HII.....	113
Figure 6-10.	Firmware Management Protocol Support Flow	114
Figure 6-11.	Physical Drive Location Information in UEFI	118
Figure 6-12.	Intel® VROC GUI System Report Window	122
Figure 6-13.	Intel® VROC Components from the Intel® VROC GUI System Report ..	123
Figure 6-14.	Selecting System Notifications	125
Figure 6-15.	Intel® VROC GUI Storage System Window	126
Figure 6-16.	Navigation Options in the Intel® VROC GUI	127
Figure 6-17.	Available Options in the Home Screen of the Intel® VROC GUI.....	128
Figure 6-18.	Information Section in the Home Screen of the Intel® VROC GUI	128
Figure 6-19.	Available Options in the Preferences Screen of the Intel® VROC GUI... ..	129
Figure 6-20.	Available Options in the Help Screen of the Intel® VROC GUI.....	130
Figure 6-21.	Intel® VROC GUI Version Information.....	133
Figure 6-22.	Online Support Options in the Intel® VROC GUI	134
Figure 6-23.	RAID Volume Types Description in the Intel® VROC GUI	138
Figure 6-24.	Device Manager Screen	161
Figure 6-25.	Driver Details Screen	162
Figure 6-26.	Intel® VROC UWP GUI Package	164

Figure 6-27. Intel® VROC UWP GUI Installation Process 165
 Figure 6-28. Email Message Format 167

Tables

Table 1-1. Terminology 13
 Table 2-1. Product Release Numbering Scheme for Intel® VROC 17
 Table 2-2. Supported Platforms for Intel® VROC (VMD NVMe RAID) 18
 Table 2-3. Supported Chipsets for Intel® VROC (SATA RAID)..... 19
 Table 3-1. Firmware Management Protocol Support..... 26
 Table 5-1. Key Features of Intel® VROC (VMD NVMe RAID) 40
 Table 5-2. Errors Handled by the Intel® VMD Controller 42
 Table 5-3. Reference Documents in the Intel® VROC Package Kit 49
 Table 5-4. Intel® VROC Upgrade Hardware Keys 51
 Table 5-5. Recommended Strip Size for Intel® NVMe SSDs for Optimal Wear Leveling
 59
 Table 5-6. Intel® VROC RAID 5 Default Settings 61
 Table 6-1. Intel® VROC Common Features 67
 Table 6-2. Intel® VROC RAID Volumes Automatic Initialization..... 72
 Table 6-3. Intel® VROC Strip Size Configurations 75
 Table 6-4. Usage Scenarios for Typical Strip Sizes 75
 Table 6-5. LED Management Capabilities 78
 Table 6-6. Intel® VROC Auto Rebuild Functionality 89
 Table 6-7. Verify and Repair Process in Different RAID Levels 91
 Table 6-8. Verify and Repair Process After Dirty Shutdown..... 92
 Table 6-9. UEFI Specifications and Location 100
 Table 6-10. Firmware Management Protocol Support..... 114
 Table 6-11. Values for AttributeSetting for GetImageInfo for NVMe SSDs..... 115
 Table 6-12. HardwareInstance Values 115
 Table 6-13. Values for AttributeSetting for GetImageInfo for SATA SSDs 117
 Table 6-14. Notification Area Icon States 124
 Table 6-15. Possible States for the Storage System in the Intel® VROC GUI 131
 Table 6-16. Possible States for RAID Volumes in the Intel® VROC GUI 131
 Table 6-17. Array Properties 135
 Table 6-18. Supported RAID Volume Combinations on a Single Array 140
 Table 6-19. RAID Volume Status as shown in the Intel® VROC GUI and HII..... 144
 Table 6-20. Busy Volume States 144
 Table 6-21. General Parameters 145
 Table 6-22. RAID Level Migration Examples Supported by Intel® VROC..... 152
 Table 6-23. List of Disk Properties from the Intel® VROC GUI 153
 Table 6-24. Port Properties in the Intel® VROC GUI 158
 Table 6-25. Microsoft* .NET Framework Impact for Intel® VROC GUI..... 168
 Table 6-26: Intel® VROC OS Driver vs. Pre-OS Compatibility Matrix 170
 Table 7-1. Intel® VROC SATA RAID Features 172
 Table 7-2. BIOS Control Feature Set (BCFS) Bits Description 172
 Table 8-1. Intel® NVMe SSDs Supported in Intel® VROC (NonVMD NVMe RAID) 178



Revision History

Version	Description	Date
001	Initial release.	May 2017
002	Intel® RSTe 5.1 release.	May 2017
003	Intel® RSTe 5.2 release.	June 2017
004	Intel® RSTe 5.3 release.	October 2017
005	Updates to include support for Intel® Xeon® Processor D-2100 Product Family.	January 2018
006	Intel® RSTe 5.5 release.	August 2018
007	Updated to reflect the name change to Intel® VROC.	December 2018
008	Intel® VROC 6.1 release.	April 2019
009	Intel® VROC 6.2 release.	October 2019
010	Intel® VROC 6.3 release.	February 2020
011	Intel® VROC 7.0 release.	June 2020
012	Update to supported operating systems and a product operational clarification.	June 2020
013	Intel® VROC 7.5 release.	February 2021
014	Intel® VROC 7.6 release.	May 2021
015	Intel® VROC 7.7 release.	November 2021
016	Intel® VROC 7.8 release.	August 2022
017	Intel® VROC 8.0 release.	November 2022
018	Initial public release.	July 2023
019	Intel® VROC 8.2 release.	August 2023
020	Intel® VROC 8.5 release.	September 2023
021	Intel® VROC 8.6 release.	March 2024
022	Intel® VROC 8.6 PV3 update release.	May 2024
023	Intel® VROC 9.0 release.	June 2024
024	Intel® VROC 9.1 release.	January 2025
025	Intel® VROC 9.2 release.	June 2025
026	Intel® VROC 9.2 PV2 release.	September 2025
027	Intel® VROC 9.3 release.	October 2025





1 Introduction

This document presents the functional requirements and technical details that make up the Intel® Virtual RAID on CPU (Intel® VROC) family of products: Intel® VROC (VMD NVMe RAID), Intel® VROC (SATA RAID) and Intel® VROC (Non VMD NVMe RAID).

The features outlined will include those for the Pre-Operating System (Pre-OS) components (OS independent) and the Windows* based drivers, tools, and Graphical User Interface (GUI).

1.1 Intended Audience

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

- Intel® VROC (VMD NVMe RAID)
- Intel® VROC (SATA RAID)
- Intel® VROC Pre-OS packages
- Intel® VROC GUI and other management tools

1.2 Reference OEM Platform Documentation

Refer to your OEM for a full list of available feature sets. If any of the information in this document conflicts with the support information provided by the platform OEM, the platform documentation and configurations take precedence.

Customers should always contact the place of purchase or system/software manufacturer with support questions about their specific hardware or software configuration.

1.3 Terminology

Table 1-1. Terminology

Term	Definition
AER	Advanced Error Reporting
AHCI	Advanced Host Controller Interface
AIC	Add In Card
API	Application Programming Interface
ATA	Advanced Technology Attachment
ATAEventPI	Advanced Technology Attachment Packet Interface

Term	Definition
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	Gigabyte
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.
Hot Spare Disk	A disk flagged 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	Megabytes
NAI	Notification Area Icon
NCQ	Native Command Queuing
NPEM	Native PCIe Enclosure Management
NTFS	NT File System
NVMe	Non-Volatile Memory Express
ODD	Optical Disk Devices

錯誤! 使用 [常用] 索引標籤將 **Heading 1** 套用到您想要在此處顯示的文字。



Term	Definition
ODM	Original Design Manufacturer
OEM	Original Equipment Manufacturer
On Demand	Intel® Licensing activation model
OROM	Option ROM
OS	Operating System
PCH	Platform Control Hub
PCIe	Peripheral Component Interconnect Express)
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	Red Hat Enterprise Linux*
RSTe	Rapid Storage Technology enterprise
RWH	RAID Write Hole
SATA	Serial Advanced Technology Attachment
sSATA	Secondary Serial ATA
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 Drive – non-volatile memory
UI	User Interface
UEFI	Unified Extensible Firmware Interface
UWD	Universal Windows* Driver
Intel® VMD	Intel® Volume Management Device
VMD Domain	A grouping of drives behind a single Intel® VMD
Intel® VMD (PCH)	Intel® Volume Management Device for the Platform Controller Hub
Volume Roaming	Moving of RAID volumes and all its array members between two enabled platforms to maintaining data availability or different supported operating systems.
Intel® VROC	Intel® Virtual RAID on CPU
DXE	Drive Execution Environment



錯誤! 使用 [常用] 索引標籤將 **Heading 1** 套用到您想
要在此處顯示的文字。

Term	Definition
PEI	Pre-EFI Initialization
PEIM	Pre-EFI Initialization Module





2 Product Overview

The Intel® VROC 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.

- Intel® VROC (VMD NVMe RAID) – This product provides an enterprise RAID solution on platforms that support the Intel® VMD technology. Intel® VMD must be enabled.
- 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.

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

These three products, although are provided in a single product package, can be installed, and used independently. For example (unless otherwise restricted by the platform itself), Intel® VMD is not required to enable and use Intel® VROC (SATA RAID). Likewise, Intel® VROC (VMD NVMe RAID) can be enabled and used without Intel® VROC (SATA RAID).

2.1 Product Release Numbering Scheme

The product release version is divided into 4 sections or numbers following the format AA.BB.CC.DDDD. For example: 5.0.0.1001 prior to Intel VROC version releases 9.2 PV. Starting with release builds 9.2 the release build number will be five digits long instead of four. For Example: 9.2.0.10'000.

Table 2-1. Product Release Numbering Scheme for Intel® VROC

Number	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.



Number	Description
DDDD: Release Build Number	This section represents the build number of the release AA.BB.CC, NOTE: For production releases, the build number generally begins with the number '1' (e.g. AA.BB.CC.1001).
EEEE: Release Build Number (Starting with 9.2)	This section represents the build number of the release AA.BB.CC, NOTE: For production releases starting in version 9.2 PV, the build number will now consist of five digits instead of four (e.g. AA.BB.CC.10001).

NOTES:

1. With the addition of the five-digit Release Build string, there may be an impact to OEMs who have manufacturing scripts that monitor and set for the initial four-digit Release Build string instead starting with Intel VROC 9.2 PV Release.

2.2 Supported Operating Systems per Platform

Please navigate to the following link to view the latest support matrix for operating systems with Intel® VROC. [Intel® Virtual RAID on CPU \(Intel® VROC\) Operating Systems Support List](#)

2.3 Supported Platforms, Chipsets and SKUs

The Intel® VROC package was designed to work with, tested and validated on platforms that are based on the following Intel® Customer Reference Boards (CRBs):

Table 2-2. Supported Platforms for Intel® VROC (VMD NVMe RAID)

CPU	Platform	VMD 1.0 Device ID	VMD 2.0/3.0 /4.0 Device ID	# Of VMD Domains
Intel® Xeon® Scalable Processor Family – W	Intel® Xeon® Scalable Processor family workstation†	201D	N/A	3 per CPU
Intel® Xeon® Scalable Processor Family – SP	Intel® Xeon® Scalable Processor family server and workstation†	201D	28C0	3 per CPU
Intel® Xeon® Processor D Product Family	Intel® Xeon® Processor D Product Family based platform†	201D	28C0	4 per CPU

CPU	Platform	VMD 1.0 Device ID	VMD 2.0/3.0 /4.0 Device ID	# Of VMD Domains
The Intel® Xeon® Processor Scalable Memory Family	Intel® Xeon® Scalable Processor family server†	201D	28C0	4 per CPU (VMD 1.0) 5 per CPU (VMD 2.0/3.0)
Intel® X299 High End Desktop	Intel® Xeon® Scalable Processor family server and workstation†	201D	N/A	3 per CPU

† Unless otherwise specified in the Release Notes.

NOTES:

1. Intel® VROC (VMD NVMe RAID) 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.
2. Intel® VROC (Non-VMD NVMe RAID) is no longer included in the Intel® VROC package. If support for this product is required refer to the Intel® VROC 7.x based family of products.

Table 2-3. Supported Chipsets for Intel® VROC (SATA RAID)

Chipset	Platform	RAID Controller Device ID	# Of Ports
Intel® C610 series chipset	Platforms containing the Intel® C610 series chipset†	2826 (SATA) 2827 (sSATA)	6 SATA 4 sSATA
Intel® C610 series chipset	Platforms Refreshes containing the Intel® C610 series chipset†	2826 (SATA) 2827 (sSATA)	8 SATA 6 sSATA
Intel® C220 series chipset	Platforms containing the Intel® C220 series chipset†	2826 (SATA)	6 SATA
Intel® C230/240/250/260 series chipset	Platforms containing the Intel® C230-260 series chipset†	2826 (SATA)	8 SATA
Intel® C620 series chipset	Platforms containing the Intel® 620 series chipset†	2826 (SATA) 2827 (sSATA)	8 SATA 6 sSATA
Intel® C422 series chipset	Platforms containing the Intel® C422 series chipset†	2826 (SATA)	8 SATA
Intel® C741 series chipset	Platforms containing the Intel® C741 series chipset†	2826 (SATA) 2827 (sSATA) 282F (tSATA)	8 SATA 6 sSATA 6 tSATA

† Unless otherwise specified in the Release Notes.

3 New Features

3.1 New Features Introduced with Intel® VROC 9.3 Release

Intel VROC 9.3 is a maintenance release that does not contain new features.

3.2 New Features Introduced with Intel® VROC 9.2 Release

Intel VROC 9.2 is a maintenance release that does not contain new features.

3.3 New Features Introduced with Intel® VROC 9.1 Release

Intel VROC 9.1 is a maintenance release that does not contain new features.

3.4 New Features Introduced with Intel® VROC 9.0 Release

3.4.1 Added SATA RAID Support on Montage IOH

Intel VROC supports up to three IOH adapters but only one at a time. There is currently no support for splitting the drivers among more than one, so users can have up to three adapters connected but we have drivers only for 1 of them and for the 3 SATA controllers (SATA, sSATA, tSATA).

3.4.2 OOB for NVMe SSDs behind VMD via MCTP

The manageability feature for VMD owned NVMe devices using NVMe Management Interface (NVMe-MI) over Management Component Transport Protocol (MCTP) over PCIe Vendor Defined Messages (VDM) enables robust out-of-band management. The feature allows for efficient monitoring, configuration, and control of NVMe devices via a standardized communication protocol that operates independently of the host operating system.

3.4.3 NPEM Support for CPU-Attached SSD

Original support starting with Intel VROC 7.5 supported NPEM LED management through an NPEM Supported switch. Now with the new support with Intel VROC 9.0 adds the support of NPEM by using PCIe port standard bits to control LED blinking for each drive port of a system.

3.5 New Features Introduced with Intel® VROC 8.6 Release

3.5.1 Windows Server 2025 Enterprise

VROC 8.6.0.3001 Windows release package introduces support for Windows Server 2025 Enterprise version.

3.5.2 Windows 11 24H2

VROC 8.6.0.2003 Windows release package introduces support for Windows 11 24H2 version.

3.5.3 Sanitize Command

VROC 8.6 PV introduces support of NVM Express Pass-Thru Protocol command *Sanitize*. For more detail, refer to section 5.2.5 NVM Express Pass-Thru Protocol Support.

3.5.4 Storage System Event Changes

With the release of VROC 8.6 PV, the storage system event ID 4155 has been replaced with the new event IDs 4162, 4163 and 4164. This change has been implemented to define IO failure conditions more clearly: IO issued to an empty ODD, SATA or NVMe drive. For more information, please see Appendix B - Storage System Events.

3.6 New Features Introduced with Intel® VROC 8.5 Release

3.6.1 Intel® On Demand Licensing Activation

Intel® On Demand activation model is a solution managed primarily by hardware providers to enable license-based activation of select CPU features and software throughout Intel® Xeon® CPU lifecycle. Depending on the hardware provider, licenses can be purchased to enable Intel® VROC (available with 5th Gen Intel Xeon Scalable Processors codenamed Emerald Rapids) during manufacturing or in field as an upgrade.

With release 8.5, Intel(R) VROC will support Intel(R) On Demand as an alternate License activation method alongside the HW key.

3.6.2 VROC RAID 1 Only SKU

Intel® VROC RAID1 Only SKU configuration is enabled when the Intel® VROC RAID1 Only On Demand license is activated on the platform. This SKU has the same behavior



as the Intel® VROC Pass-Thru SKU with the addition of support for RAID 1 capability. When the Intel® VROC RAID1 Only License is activated, and Intel® VMD is enabled, the Intel® VROC (VMD NVMe RAID) UEFI HII will enable RAID 1 management on all supported NVMe drives. This functionality applies only to those PCIe slots that have Intel® VMD enabled.

The message displayed in this mode is: *Intel® VROC (RAID1 Only)*.

3.7 New Features Introduced with Intel® VROC 8.2 Release

3.7.1 Support for Pre-Boot DMA

Intel VROC introduces support for Pre-Boot DMA protection when the setting is enabled in the OEM platform BIOS. DMA protection is a feature that is supported starting with Windows 10. It is supported on 64-bit IA processors with virtualization extensions, including VT-X.

3.7.2 UEFI PEI Module driver

Intel VROC introduces VROC Pre-EFI Initialization Module (PEIM) to support BIOS recovery scenario during the boot.

VROC PEI module for SATA – provides access to RAID volumes and SATA drives connected to Intel SATA controller. GUID:

```
{0xbcc5d5cb, 0x828c, 0x4fa1, {0x8a, 0x2a, 0xbf, 0x31, 0x1c, 0x2e, 0x61, 0x89}}
```

VROC PEI module for VMD – provides access to RAID volumes and NVMe drives connected to multiple Intel VMD controllers. GUID:

```
{0x51f28c3b, 0xcfca, 0x4308, {0xb7, 0xdb, 0xc6, 0x82, 0xac, 0xcd, 0xe0, 0x8a}}
```

Both modules have common code base with PEI interface implementation (PEIM adapter) and RAID logic. The main difference between them is different transport implementations: AHCI and VMD with NVMe). Every VROC PEI module installs two PEIM-to-PEIM interfaces (PPI), defined in PI Specification:

Provided PPI:

EFI_PEI_RECOVERY_BLOCK_IO_PPI (SATA and VMD)

EFI_PEI_RECOVERY_BLOCK_IO2_PPI (SATA and VMD)

EDKII_PEI_STORAGE_SECURITY_CMD_PPI (VMD)

EDKII_PEI_NVM_EXPRESS_PASS_THRU_PPI (VMD)

Required PPI:

SATA Init Done PPI

VMD Host Controller PPI

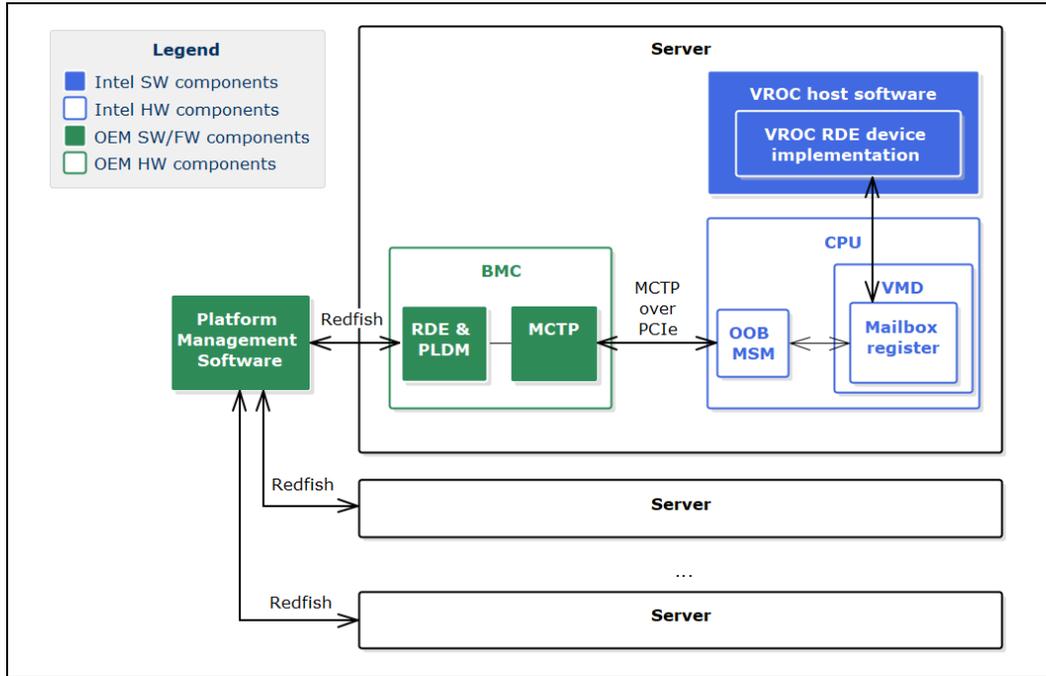
Note: Block I/O PPIs are installed by VROC PEI module only when EFI_PEI_BOOT_IN_RECOVERY_MODE_PPI is already installed prior to VROC PEI module entry point.

3.8 New Features Introduced with Intel® VROC 8.0 Release

3.8.1 Intel® VROC Out-Of-Band Management (Server Platforms Only)

The Intel® VROC Out-of-Band (OOB) management interface supports only NVMe drives and Intel® VROC (VMD NVMe RAID) volumes. SATA drives and Intel® VROC (SATA RAID) volumes can be managed through other Intel® VROC (in-band) interfaces. In the Intel® VROC OOB management solution, the BMC is the entry point for the requests from a management network. It processes the requests (Example: management commands) and sends them to the MCTP endpoint exposed by the Intel® OOB-MSM, via MCTP over PCIe. The Intel® OOB-MSM works as MCTP bridge, it receives requests from BMC and sends it to Intel® VROC OOB host software over Intel® VMD mailbox register. The Intel® VROC OOB host software component periodically polls the Intel® VMD mailbox register for new requests. The Intel® VROC OOB host software performs the requested action and sends the response to the BMC in the opposite way. A high-level view of the solution is presented in the below figure.

Figure 3-1. Overview of Intel® VROC Out-Of-Band Management Solution



3.8.2 LED Management on PCH-VMD

There are two PCIe ports designed and controlled under PCH, Emmitsburg, and this feature with LED management is implemented in the circuit and the on-board secondary CPLD. Designers need to implement the same or similar logic control as this CPLD does to control two LEDs designed on-board to represent the RAID status.

3.8.3 LED Patterns Conflict Precedence

When two or more different LED patterns are applicable to the same drive, the Intel® VROC Windows* driver shall apply the following order in terms of precedence:

1. Locate LED pattern
2. Fault LED pattern
3. RAID rebuild LED pattern

3.8.4 Intel® VROC 8.0 New GUI Design

3.8.4.1 Warning Message Displayed when GUI and Driver Version is Mismatching

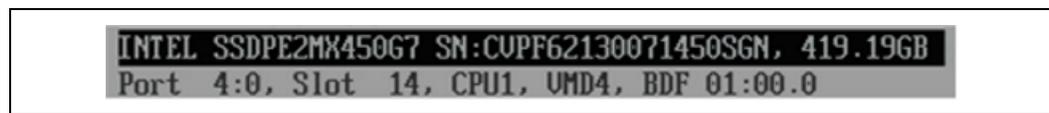
When this warning message condition is detected from one major version or 2 minor versions difference, the Intel® VROC GUI will not allow to manage the storage

controller but can provide a set of functionalities corresponding to the lowest driver version. The Intel® VROC GUI still can start even if there is no driver installed.

3.8.5 Physical Drive Location Information in UEFI

When Intel® VROC product displays how the Intel® VMD controller enumerates the drives, the drive information is showing in VROC HII like in the below figure. The information is not straightforward and not easy to interpret. To present the drive location information in a user-friendly manner, Intel® VROC offers a protocol that will allow customers to report how the drives are connected and how customer would like the information reported in customer platform.

Figure 3-2. Physical Drive Location Information in UEFI



3.8.6 Third SATA Controller (Server Platform Only)

Intel® VROC family of products introduces support for a third SATA controller on the Intel® Platform Controller Hub. Starting with the Intel® C741 chipset, Intel® VROC will provide support for the third SATA controller (tSATA) when the controller is set for RAID mode. Refer to the platform design documentation for how the tSATA ports are laid out.

3.9 New Features Introduced with Intel® VROC 7.8 Release

No new features were introduced in the Intel® VROC 7.8.0.1031 release package.

3.10 New Features Introduced with Intel® VROC 7.7 Release

The Intel® VROC 7.7 release package introduces the following new features: support for UEFI Firmware Management Protocol; support for `EFI_ATA_PASS_THRU_Protocol`; support for the ability for the OEM/ODM to disable Intel® VROC LED Management.

3.10.1 UEFI Firmware Management Protocol Support

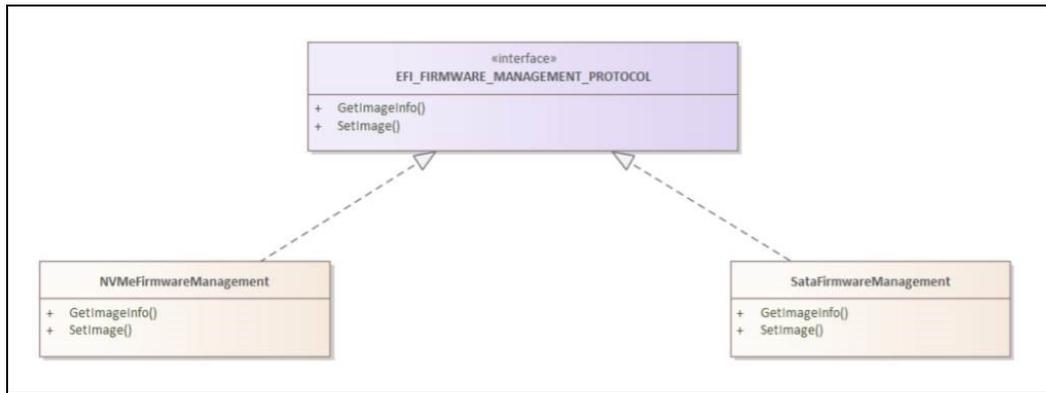
Intel® VROC 7.7 introduces limited support for the *UEFI Firmware Management Protocol (FMP)* as outlined in UEFI Specification version 2.9. The Intel® VROC UEFI drivers support updating the drive firmware (through UEFI FMP) when the drive is managed by Intel® VROC. This includes both Intel® VROC (VMD NVMe RAID) and Intel® VROC (SATA RAID). The Intel® VROC UEFI driver has implemented a subset of protocol functions defined in the UEFI Specification, which is reflected in the below table.

Table 3-1. Firmware Management Protocol Support

EFI Firmware Management Protocol Functions	Intel® VROC UEFI Implementation
GetImageInfo	Supported
SetImage	Supported

The following diagram highlights the process flow to support both NVMe and SATA drives.

Figure 3-3. Firmware Management Protocol Support Flow



3.10.2 Intel® VROC (SATA RAID) Support of EFI_ATA_PASS_THRU_Protocol

The Intel® VROC 7.7 release package introduces limited support for **EFI_ATA_PASS_THRU_PROTOCOL** commands to provide information on the SATA drives managed by Intel® VROC (SATA RAID).

3.10.3 Disable Locate LED within Intel® VROC LED Management

Intel® VROC supports the OEM/ODMs to disable the *Locate LED* functionality within the Intel® VROC LED Management. By enabling the feature, the OEM/ODM can use their own tools to initiate a *Locate LED* functionality within their platform.

3.11 Features Introduced with Intel® VROC 7.6 Release

The Intel® VROC 7.6 release package introduces two new features. These two features are Intel® VROC Pre-OS Failed RAID Volume (limited) Recovery as well as UEFI ATA Passthrough Protocol Support for Intel® VROC (SATA RAID)

3.11.1 Intel® VROC Pre-OS Environment RAID Volume Failure Recovery

Intel® VROC 7.6 introduces a new feature in the VROC HII page that will allow the user to attempt to recover from a failed RAID volume. When a failed RAID volume is encountered during boot, the option is made available in the VROC HII page for which the user can select. This is a multi-step process.

The first step is to select the RAID volume in 'Failed' state and then select the option to force it to 'Degraded' state. This will expose the next option to specify the drive to enable this action.

Once the RAID volume is changed to 'Degraded' state, the standard RAID volume recovery process initiates the rebuilding process.

Note: Intel makes no guarantee of successful recovery from a failed state using this option. This must be treated as a last chance effort and there is no guarantee that there won't be some data loss. Intel always recommends recovering a failed RAID Volume by recreating the RAID volume from scratch and restore the data from the latest platform image backup.

3.11.2 Intel® VROC (SATA RAID) UEFI Support for EFI_ATA_PASSTHRU

Intel® VROC 7.6 (SATA RAID) introduces support for `EFI_ATA_PASS_THRU` protocol by the Intel® VROC (SATA RAID) UEFI driver. The following specific options are supported in Intel® VROC 7.6:

- `PassThru;`
 - `IDENTIFY`
 - `ATA_READ_LOG_EXT`
- `GetNextPort;`
- `GetNextDevice;`
- `BuildDevicePath;`
- `GetDevice;`

3.12 Features Introduced with Intel® VROC 7.5 Release

The Intel® VROC 7.5 release package introduces several new features to support the latest Intel® platforms as well as improve the user experience. The key new features introduced are:

- Intel® VMD 2.0 support that includes:
 - Increase in MSIX vectors to 64.
 - Intel® VMD support for NVMe devices attached to the Platform Controller Hub (PCH).



- Increase in the number of PCIe lanes that can be controlled by VMD from 48 to 64.
- Increase in the number of VMD devices from 3 to 5.

3.12.1 Intel® VMD 2.0 Features

3.12.1.1 Increasing MSIX Vectors to 64

Intel® VROC 7.5 introduces support for customer configurations that can support 64 MSIX Vectors. On platforms that support Intel® VMD 1.0, the MSIX support is limited to 32 MSIX vectors. For Intel® VMD 1.0 platforms, a single VMD domain can support up to 24 NVMe SSDs. This means that those 24 NVMe SSDs will share a single Intel® VMD MSIX vector. As the number of vectors, in newer NVMe devices, increase beyond 32, this can result in a platform performance impact. With the introduction of Intel® VMD 2.0, and Intel® VROC 7.5, this increase to 64 MSIX vectors should help to alleviate this problem. This is because the average dual socket server will have between 48 and 64 cores, which limits the number of MSIX vectors on a given NVMe SSD to 64, for optimal usage.

3.12.1.2 Intel® VMD (PCH) Support

Intel® VROC 7.5 introduced support for the Intel® VROC (VMD NVMe RAID) management of NVMe SSDs connected to the Platform Controller Hub (PCH). Utilizing the Flex-IO capabilities of the PCH, Intel® VMD 2.0 and later technology can now take ownership of slots allowing Intel® VROC (VMD NVMe RAID) to control and manage the NVMe SSDs attached to those slots. This is accomplished by utilizing the Function Level Assignment of the PCH PCI functions. Utilizing the BIOS setup menus, the user will be able to enable Intel® VMD on the designated slots, or PCH functions (depending on the BIOS implementation) and the PCH functions will be reassigned to an Intel® VMD (PCH) controller. To be able to accomplish this, there must be NVMe SSDs present on the slot(s) and the slots used must support Slot Implemented Capabilities. Otherwise, Intel® VMD (PCH) cannot be enabled.

When this feature is fully enabled, it will reassign sSATA ports 2-5 to be PCIe lanes managed by Intel® VMD (PCH). This is translated into PCIe root ports 8-11 and Flex I/O ports 14-17. The general configuration is 2 - X2 PCIe lanes.

For VMD support for PCH on Eagle Stream (EGS) platforms, any root port behind PCH that has a slot, and that slot has the 'slot implemented' bit set can be put behind VMD.

This is the current implementation by our CRB BIOS team where only certain PCIe root ports that have the 'slot implemented' bit set are allowed by the BIOS to be put behind VMD and the rest are greyed out in the BIOS settings on the EGS platform.

Things like AHCI controller on PCH that manages SATA, graphics card, NIC cards etc. do **not** satisfy the above condition and cannot be put behind VMD.

This feature is intended for supporting a simple RAID 1 boot using two NVMe SSDs attached directly to the PCH. However, if the platform supports expanded configurations, using a retimer or switch Add-In-Cards (AIC), full Intel® VROC (VMD NVMe RAID) support can be obtained. Intel® VROC (VMD NVMe RAID) was only

validated with two NVMe SSDs directly attached to the PCH. Any configurations beyond two directly attached NVMe drives is not recommended.

3.12.1.2.1 Intel® VMD (PCH) Pre-OS Support

Intel® VMD (PCH) Pre-OS support is included in the Intel® VROC (VMD NVMe RAID) Pre-OS images, which is part of the Intel® VROC release package. There are no other Pre-OS images required.

3.12.1.2.2 Intel® VMD (PCH) Pass-Thru Boot Support

Intel® VROC 7.5 and later provides Intel® VROC (VMD NVMe RAID) Pre-OS support for NVMe SSDs attached to the PCH when Intel® VMD (PCH) is enabled. This will allow an OS to be installed onto and boot from an NVMe device managed by Intel® VROC (VMD NVMe RAID). No Intel® VROC upgrade key is required to utilize this feature.

3.12.1.2.3 Intel® VMD (PCH) RAID Boot Support

Intel® VROC 7.5 and later provides Intel® VROC (VMD NVMe RAID) Pre-OS support for NVMe SSDs attached to the PCH when Intel® VMD (PCH) is enabled. When an Intel® VROC upgrade key is present and Intel® VMD (PCH) is enabled, the user will be able to use the Intel® VROC (VMD NVMe RAID) Pre-OS HII to setup and manage a RAID volume using the NVMe SSDs attached to the PCH. This will allow an OS to be installed onto and boot from an Intel® VROC (VMD NVMe RAID) volume attached to Intel® VMD (PCH).

3.12.1.2.4 Intel® VMD (PCH) RAID Data Volume Spanning

Once Intel® VMD (PCH) is enabled, Intel® VROC (VMD NVMe RAID) will treat it like any other Intel® VMD controller. Spanning data RAID volumes are supported as they are currently outlined in this document. This applies to the Intel® VROC Pre-OS tools, the Windows* OS GUI and the CLI tool. The option to span VMD controllers is supported, but not recommended.

Note: This is supported but not recommended due to a performance penalty using the PCH.

Boot volume spanning is not supported. All devices used to create a boot volume must reside on the same controller (Intel® VMD controller domain or SATA controller).

3.12.1.2.5 Intel® VMD (PCH) Designations

The Intel® VROC (VMD NVMe RAID) management tools (Pre-OS tools, Windows* GUI and CLI tool) are used to manage devices attached to the Intel® VMD (PCH) controller, the corresponding information displayed will indicate 'PCH' with any device or controller associated with Intel® VMD (PCH).

3.12.1.2.6 Intel® VMD (PCH) Hot Plug Support

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.

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/U.3 form factor only with this release). This includes surprise insertion/removal of 2.5" form factor devices in the Pre-OS and OS runtime environment.

Note: Surprise hot-plug is only supported for U.2/U.3 form factor NVMe drives; other form factor (AIC, M.2) SSDs are not supported.

For hot-plug support in Pre-OS environment, the user needs to go out of VROC HII menu and enter VROC HII again after drive removal or insertion.

3.12.2 Number of Intel® VMD Increased

Intel® VMD 2.0 technology increased the number of PCIe lanes controlled by the Intel® VMD from 48 to 64. This, along with the Intel® VMD (PCH) support, increases the total amount of Intel® VMDs to 5:

- 4 Intel® VMDs off the CPU
- 1 Intel® VMD of the PCH (Intel® VMD (PCH))

3.12.3 Native PCIe Enclosure Management (NPEM)

Intel® VROC 7.5 introduces support for the *Native PCIe Enclosure Management (NPEM)* standard for LED management in a PCIe 4.0 based environment. This capability is discoverable in each switch-downstream-port. If it is discovered to be present, the Intel® VROC LED utility will use NPEM control, capability, and status registers to visually indicate the various drive and volume states.

3.12.4 Limited Self Encrypted Drives

Intel® VROC 7.5 introduces Self Encrypting Drive (SED) key management support. The implemented key management is only in the UEFI environment but allows secure booting with SEDs into all Intel® VROC OS environments. Current support is limited to VROC Premium SKU. For detail of Premium, please refer to section 5.4.1 Intel® VROC RAID Modes of Operation for Upgrade Hardware Keys for RAID Support

3.12.5 Intel® VROC 7.5 GUI UWD Upgrade Limitations

Due to some updates required by the Intel® VROC 7.5 GUI, a platform with Intel® VROC 7.0 based drivers or Intel® VROC 6.x based driver won't be able to upgrade to the Intel® VROC 7.5 GUI, from the Microsoft* App Store, until after the drivers have been updated. This is due to the potential of not being able to properly manage the Intel® VROC RAID volumes already established in the system after the upgrade. Once the platform is upgraded to the Intel® VROC 7.5 drivers, the GUI can be updated.

3.12.6 Intel® VROC 7.5 CLI

Intel® VROC 7.5 CLI tool now supports the creation of a RAID volume size down to two digits after the decimal point.

3.12.7 Unsupported Drives

The Intel® VROC Intel SSD Only upgrade key is not compatible with NVMe SSDs that are not manufactured by Intel. When Intel® VROC is activated using this key, any non-Intel drives are considered unsupported. If an unsupported drive is encountered by the Intel® VROC 7.5 GUI and/or CLI, the displayed information will show the usage as 'Pass Through'.

3.13 Features Introduced with Intel® VROC 7.0 Release

3.13.1 Removal of AHCI Mode Support

With the release of the Intel® VROC 7.0 package, the package will no longer include drivers to support the PCH configured for AHCI mode. All support for this configuration will be via the Microsoft* inbox AHCI driver.

- For fresh clean (OS installation) configurations, the Intel® VROC 7.0 installer will bypass the PCH if it is in AHCI mode.
- For existing platforms that has Intel® VROC AHCI driver installed, running the Intel® VROC 7.0 installer will remove/uninstall the Intel® VROC AHCI driver so the Microsoft* inbox AHCI driver can be used. A message will be displayed to inform the user of the upcoming change.

3.13.2 Removal of Legacy Option ROM Support

With the release of the Intel® VROC 7.0 package, the package will no longer include the Intel® VROC (SATA RAID) Legacy Option ROM Pre-OS components. All Pre-OS support for Intel® VROC (SATA RAID) will only be through the UEFI environment. If Intel® VROC (SATA RAID) support, through the Legacy Option ROM environment, is required continue using the Intel® VROC 6.3 release package.

3.13.3 Removal of Windows* 7 Support

With the release of the Intel® VROC 7.0 package, the package will no longer include support for Windows* 7. All components and support for Windows* 7 (Windows* drivers and Pre-OS components) have been removed. If Windows* 7 support is required, continue using the Intel® VROC 6.3 release package.



4 Intel® VROC Package Components

This section outlines the individual components that are included in the Intel® VROC package that is delivered to customers. The specific details on each component are in the following sections of this document. The Intel® VROC package is comprised of:

- Intel® VROC Command Line Interface (CLI) tool
- Intel® VROC installation utility
- Intel® VROC F6 drivers
- Intel® VROC Pre-OS UEFI packages

These package components comprised of, and support the Intel® VROC (VMD NVMe RAID) and the Intel® VROC (SATA RAID).

4.1 Intel® VROC CLI

The Intel® VROC 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 the Intel® VROC (VMD NVMe RAID) environment and the Intel® VROC (SATA RAID) environment.

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

4.2 Intel® VROC OS Installer

The Intel® VROC 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 Graphical User Interface (GUI) and supporting components. This single installer is used to install the Intel® VROC (VMD NVMe RAID) and the Intel® VROC (SATA RAID).

On platforms with Intel® CPU's that contain Intel® VMD technology (enabled or disabled), the installer will only install Intel® VROC (VMD NVMe RAID) and the Intel® VROC (SATA RAID).

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

Note: If the Intel® VROC OS installer detects the Intel® VROC GUI from the Windows* App Store, it will remove it and replace it by the Intel® VROC GUI contained in the installer.

4.3 Intel® VROC (VMD NVMe RAID) Environment

The Intel® VROC 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).

4.3.1 Intel® VROC (VMD NVMe RAID) Pre-OS Components

VMDVROC_1.efi / VMDVROC_2.efi

These two UEFI device driver files are to be included in the platform BIOS. Both 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:

- 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

This utility can be used by OEMs to test the 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.

- 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 Intel® VROC Release Notes associated with the latest release. By 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 (e.g., **HWKeyCheckVROC.efi**).
- The output from this EFI tool can also be saved as a .txt file (e.g., **HWKeyCheckVROC.efi > HWKeyCheckVROC _Output.txt**).

RCfgVROC.efi

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.

- This tool must match the major release version of the Intel® VROC UEFI driver. By 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 (e.g., **RCfgVROC.efi**)



- The output from this EFI tool can also be saved as a .txt file (e.g., `RCfgVROC.efi /I > RCfgVROC_Output.txt` – this will save all 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

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 by the Intel representative.

- This tool must match the major release version of the Intel® VROC UEFI driver. Attach a Fat32 formatted USB 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 (e.g., `RCmpVROCRS.efi`).
- The output from this EFI tool can also be saved as a .txt file (e.g., `RCmpVROC.efi /I > RCmpVROC_Output.txt` – this will save all the Drive/Volume/Array information in the .txt file).

LedToolVMDVROC.efi

This is a 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.

- This tool must match the major release version of the Intel® VROC (VMD NVMe RAID) UEFI driver. By 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 (e.g., `LedToolVMDVROC.efi /D1 1 /TIMEOUT 2` - this will generate a Locate signal for 2 seconds).
- The tool syntax is described below:

```
LedToolVMDVROC.efi [/?] [/HELP] [/L] [/Dn state-number] [/TIMEOUT time]
```

- `/?`, `/HELP` Displays the 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, it 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 a 2-seconds delay, messages will be reset.

4.3.2 Intel® VROC (VMD NVMe RAID) F6 (or “Load Driver”) Drivers

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

The first 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 the following:

- *iaVROC.free.winX.64bit.AA.BB.CC.DDDD* – used 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 (VMD NVMe RAID) 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_AA.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.

4.4 Intel® VROC (SATA RAID) Environment

The Intel® VROC (SATA RAID) 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).

4.4.1 Intel® VROC (SATA RAID) Pre-OS Components

The VROC (SATA RAID) Pre-OS components are comprised of a UEFI package support. The Intel® VROC (SATA RAID) Pre-OS components are the following.

4.4.1.1 EFI Support for the SATA/sSATA Controller

SataDriver.efi / sSATADriver.efi

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/tSATA drives (create/delete RAID volumes) connected to the SATA controller. Other features included are:

- Seeing and reporting (to the BIOS) any SATA drive attached to the SATA/sSATA/tSATA controller when configured in RAID mode.
- Installing an OS onto a drive (or RAID volume) managed by the SATA/sSATA/tSATA controller when configured in RAID mode.



LedToolSata.efi / LedToolsSata.efi / LedTooltSata.efi

These are UEFI based 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.

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...* By using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of the Intel® VROC 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 (e.g., **LedToolSata.efi**).

- The tool syntax is described below:

```
LedTool.efi [/?] [/HELP] [/D:state_number] [/D0-5:state_number] [/TIMEOUT:time]
```

- **/?** Same as **/HELP**. Displays the help screen. Other options are ignored.
- **/D** Send the specified state to all ports. **/D0-5** option will be ignored.
- **/D0-5** Send the specified state to the selected port. 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**.
- **/TIMEOUT** After sending the states, it waits the specified number of seconds and then turns off all LEDs and performs a reset of the message sending logic.

Note: After a reset, the first sent state will turn on all other LEDs.

- Example:

```
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 / RCfgtSATA.efi

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.

The tool must match the major release version of the Intel® VROC (SATA RAID) UEFI driver. By using a Fat32 formatted USB device, users can download any release of this tool that matches the major release of the 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 (e.g., **RCfgSata.efi**).

The output from this EFI tool can also be saved as a .txt file (e.g., **RCfgSata.efi /I > RCfgSata_Output.txt** – this will save all the Drive/Volume/Array information in the .txt file).

RCmpSata.efi / RCmpsSata.efi / RCmptSata.efi

These are debug utilities to help verify that the SATA/sSATA/tSATA 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 by the Intel representative. This utility must be copied to and executed from a USB key in a UEFI Shell.

The tool must match the major release version of the Intel® VROC (SATA RAID) UEFI driver. By 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 (e.g., `RCmpSata.efi`).

The output from this EFI tool can also be saved as a .txt file (e.g., `RCmpSata.efi > RCmpSata_Output.txt`).

RClrSata.efi / RClrsSata.efi / RClrtSata.efi

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.

4.4.2 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 RAID mode 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 the following:

- *iaStorE.free.winX.64bit.AA.BB.CC.DDDD* – used 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. Although Microsoft Windows* does contain support for the SATA controller in RAID mode, Intel recommends using this driver. There is no inbox driver support for the sSATA/tSATA controller in RAID mode. The Intel SATA/sSATA/tSATA F6 drivers are located in the *VROC_AA.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.

4.4.2.1 Intel® VROC (SATA RAID) F6 (or “Load Driver”) Limitation

Starting with Microsoft Windows* 8, the Microsoft* inbox RAID driver (Intel® 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 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 because 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 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 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 with the installation. The following are the two possibilities:

1. For a RAID volume that is marked as 'Degraded', the installation process will continue 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:
 - Step 1: After the installation process encounters the above-mentioned failure (the installation process does not continue), press F10 to invoke a CMD window.
 - Step 2: If you have not already done so, insert the USB drive into the system. Navigate to your USB drive with the *IntelVROCcli.exe* utility.
 - Step 3: Run command **IntelVROCcli.exe --manage --normal-volume <volumeName>**. This will reset the volume to a normal state.
 - Step 4: Close the CMD window.
 - Step 5: In the Windows* disk selection window, reload the Intel® VROC (SATA RAID) F6 driver.
 - Step 6: Once completed, Windows* should allow installation on the RAID volume.

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





5 Intel® VROC Key Features

The Intel® VROC family of products 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).

Table 5-1. Key Features of 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-Days Trial Period • 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) • Intel® VROC Out of Band Management (OOB)
Utilities	<ul style="list-style-type: none"> • Install/Uninstall Utility • Configuration and Management Utilities 	

5.1 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 the 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 with Intel® VMD 2.0 provides up to 64 PCIe lanes which is subdivided into 4 domains. Plus 1 domain for the PCH (only platforms with the Intel® Xeon 5). Those with Intel® VMD 1.0 provides up to 48 PCIe lanes which is subdivided into 3 domains. Each VMD domain manages x16 lanes. The Intel® VMD domains can be turned on/off on x4 or x2 (x2 starting with VROC 8.6 PV)

granularity and supports either NVMe SSD device or PCIe switch device when enabled. Intel® VMD is also designed to support the following types of NVMe Solid State Drives:

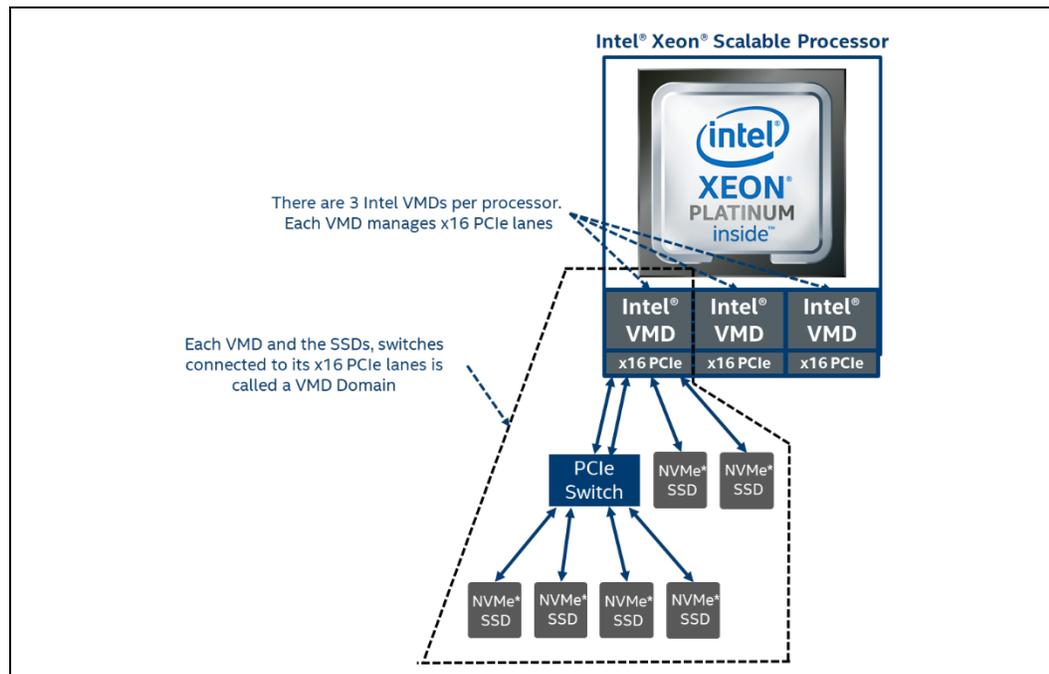
- PCI express Add in Cards (AIC)
- M.2 form factor drives
- U.2 form factor drives

5.1.1 Multiple Intel® VMD Controllers

Each Intel® Xeon® Scalable processor family CPU contains at least of 3 Intel® VMD controllers. Check the Intel® Xeon® Scalable processor family CPU documentation, for a given SKU, for the specific Intel® VMD controller count.

Note: The number of Intel® VMDs seen is dependent on the processor being used.

Figure 5-1. Multiple Intel® VMD Controllers



Note: The ability to turn the Intel® VMD domains on/off is intended for supporting configurations that require utilizing non-NVMe SSD devices (e.g. a Graphics Add-In-Card). This feature is not intended, nor does the VROC product support, a configuration where the Intel® VMD controller has one or more domains disabled (off) with NVMe SSDs connected (e.g. used as a system boot drive) while the remaining domains are enabled (on) and managed by VROC. The VROC family of products requires that Intel® VMD be fully enabled when only NVMe SSDs are attached. And that the on/off domain granularity be used only when non-NVMe SSDs are being used.



5.1.2 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 utilize and/or customize the LED management behavior in this product, see the section [LED Management](#) in this document.

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

5.1.3 Intel® VMD Error Management

Intel® VMD technology handles *Advanced Error Reporting (AER)* and functions as a crash dump and hibernate target within the 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 ecosystem.

The Intel® VMD controller will handle the errors as follows. The list of errors handled are listed in the table below.

- 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 the 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. Gather this information and provide it to your Intel® VROC representative for analysis.

Note: If IIO MCA is enabled on the platform, this will supersede all error reporting, even when VMD is enabled. As a result, VMD error reporting may not properly handle errors seen on VMD root ports. To help ensure that these errors are properly handled, consult your platform BIOS user guide to set IIO MCA to disabled.

Table 5-2. Errors Handled by the Intel® VMD Controller

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)

Advanced Uncorrectable Error	Severity Register
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)

5.1.4 Intel® VMD Enabling

These instructions will explain 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.

Step 1: Immediately following POST, select the option that will allow the user to access the BIOS setup menu. This example uses F2.

Step 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>.

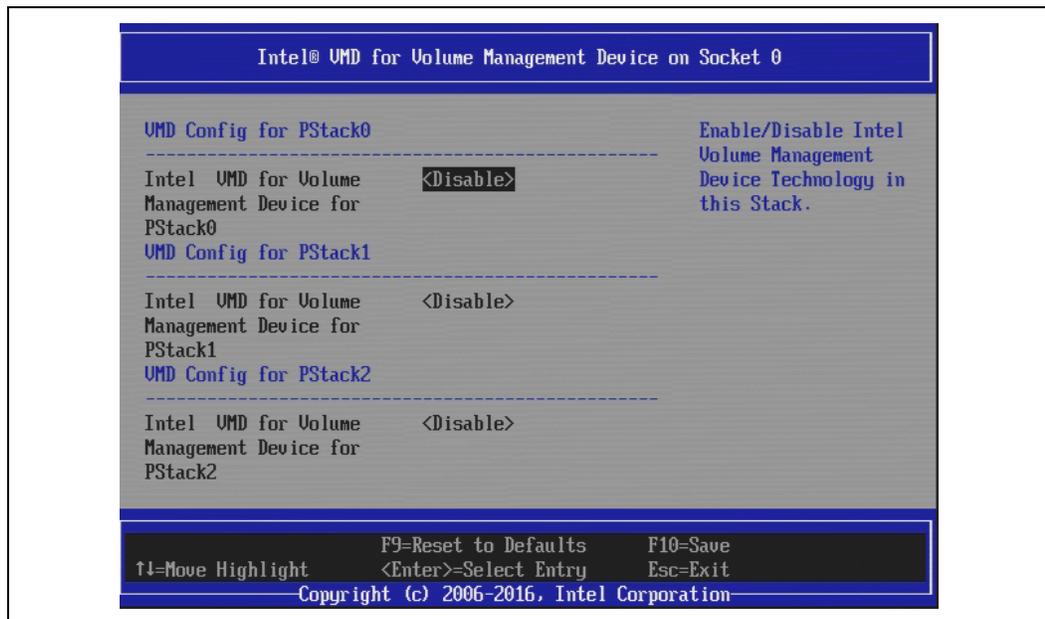
Step 3: Using the arrow keys, move the cursor to **Socket Configuration** and press <Enter>.

Step 4: Using the arrow keys, move the cursor to **Intel® VMD technology** and press <Enter>.

Step 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).

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

Figure 5-2. Enabling Intel® VMD



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

Note: See the user’s platform BIOS manufacturer documentation for a complete list of options that can be configured.

5.1.5 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; other form factor (AIC, M.2) SSDs do not 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). This includes surprise insertion/removal of 2.5” form factor devices in the OS runtime environment.

5.2 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 is 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:

- Initialize and enumerate NVMe devices managed by the Intel® VMD domains (when enabled) and report that information to the BIOS.
- 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.
- 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.
- Interface to and boot from NVMe devices managed by the Intel® VMD.
- Interface to and boot from Intel® VROC (VMD NVMe RAID) volumes.

5.2.1 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, Refer to the section [Intel® VROC Unified Extensible Firmware Interface \(UEFI\) Drivers](#).

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

The Intel® VROC 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, 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.

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

The Intel® VROC (VMD NVMe RAID) UEFI driver supports a maximum number of 96 NVMe drives across all Intel® VMD controllers within the supported platform.



5.2.4 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 GUI in the Windows* operating system for creating RAID volumes after the operating system is up and running.

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

The following assumptions have been made:

- It is known how to enter 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:

Step 1: Enter the BIOS configuration setup menu to access the Intel® VROC (VMD NVMe RAID) HII user interface.

Step 2: Navigate to and select **Intel® Virtual RAID on CPU**.

Step 3: Navigate to and select **Create RAID** Volume.

Step 4: Type in a volume name and press the <Enter> key or press the <Enter> key to accept the default name. A unique name for each RAID volume is recommended.

Step 5: 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.

Step 6: 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 it, use the down arrow key, and press <Enter>.

Step 7: Using the arrow keys, highlight the drives one by one by selecting the < > bracket on the line next to the 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 the drive as part of the RAID.

Step 8: Repeat the previous step for each drive required in the RAID.

Step 9: 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.

Step 10: 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 700 GB is 716800 (700 X 1024 = 716800).

Step 11: 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%.

Step 12: 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.

Step 13: 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.

Step 14: 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.

5.2.5 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_NVME_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 devices, applications will send the *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 the 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 the NVMe device in pass through mode and RAID mode.



- **Get Log Page** - To read values of SMART and general health information, the application sends the `ADMIN_GET_LOG_PAGE` command. Intel® VROC (VMD NVMe RAID) will pass the `ADMIN_GET_LOG_PAGE` command to the 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 the 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 the appropriate NVMe device. The `ADMIN_FIRMWARE_IMAGE_DOWNLOAD` command is used to download the firmware image to the controller. The new firmware image will not start to run right after the `ADMIN_FIRMWARE_IMAGE_DOWNLOAD` command. To select which firmware version will be executed after the NVMe device reset, the `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 the NVMe device in pass through mode and RAID mode.
- **Sanitize** - The *Sanitize* command is used to start a sanitize operation or to recover from a previously failed sanitize operation. It is supported for pass-through NVMe drives only and is defined in NVM Express base specification. For RAID Volumes, to sanitize a RAID member drive, the user must first delete the RAID volume, and then send the *Sanitize* command to the drive. Delete RAID volume functionality is provided by VROC Standard Storage Interface Protocol.

5.2.6 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, see the 1.04 version of the Intel® VROC (VMD NVMe RAID) IOCTL document included in the Intel® VROC package kit.

5.2.7 Intel® VROC (VMD NVMe RAID) Private Device/Member Info Protocol

The Intel® VROC (VMD NVMe RAID) package will provide a 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 Intel® VROC package kit.

5.2.7.1 Reference Documents

Below is a list of documents included in the Intel® VROC package kit.

Table 5-3. Reference Documents in the Intel® VROC Package Kit

Document	Description	Document 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_Final	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 Intel® VROC package 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 Intel® VROC package 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 Intel® VROC package kit.



5.3 Intel® VROC Monitoring and Management

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

5.3.1 Intel® VROC Graphical User Interface (GUI)

The Intel® VROC 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 GUI, see section [Intel® VROC RAID Features and Functionality](#).

5.4 Intel® VROC Features and Functionality

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

5.4.1 Intel® VROC RAID Modes of Operation for Upgrade Hardware Keys for RAID Support

The Intel® VROC (VMD NVMe RAID) 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.

The below table lists all four operational SKUs and their main features.

Table 5-4. Intel® VROC Upgrade Hardware Keys

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 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.</p> <p>In this mode, user will be able to create RAID 0 volumes on NVMe SSDs managed by the Intel® VMD. However, issues encountered are not supported.</p> <p>The message displayed in this mode is: <i>Intel® VROC (VMD NVMe RAID) (in pass-thru mode).</i></p>	Not Needed	<ul style="list-style-type: none"> • All Intel® NVMe SSDs supported in Pass-Thru mode only (no RAID) • Non-Intel® NVMe SSDs supported in Pass-Thru mode only (no RAID) (see the Intel® Virtual RAID on CPU (Intel® VROC) Supported Configurations 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: <i>Intel® VROC (Standard).</i></p>	Standard Key	<ul style="list-style-type: none"> • Pass-Thru SKU features • RAID 0, 1, 10 (Bootable and Data) • LED Management • Hot Plug Support

<p>Intel® VROC Premium</p>	<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: <i>Intel® VROC (Premium).</i></p>	<p>Premium Key</p>	<ul style="list-style-type: none"> • Standard SKU features • RAID 5 (Bootable and Data) • RAID 5 Write Hole Closure • LED Management • Hot Plug Support
<p>Intel® VROC Intel-SSD-Only</p>	<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 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 the Intel® Virtual RAID on CPU (Intel® VROC) Supported Configurations) will be supported in pass-thru mode only.</p> <p>The message displayed in this mode is: <i>Intel® VROC (in pass-thru mode).</i></p>	<p>Intel-SSD-Only Key</p>	<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 the Intel® Virtual RAID on CPU (Intel® VROC) Supported Configurations for list of supported drives) • LED Management • Hot Plug Support

<p>Intel® VROC RAID 1 Only</p>	<p>The Intel® VROC RAID1 Only SKU configuration is enabled when the Intel® VROC RAID1 Only OnDemand license is activated on the platform. This SKU has the same behavior as the Intel® VROC Pass-Thru SKU with the addition of support for RAID 1 capability. When the Intel® VROC RAID1 Only License is activated, and Intel® VMD is enabled, the Intel® VROC (VMD NVMe RAID) UEFI HII will enable RAID 1 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: <i>Intel® VROC (RAID1 Only).</i></p>	<p>Intel® On Demand Licensing activation Only</p>	<ul style="list-style-type: none"> • Pass-thru and RAID 1 support (Bootable and Data) • • LED Management • Hot Plug Support
---------------------------------------	--	---	---

Note: 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.

These 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. 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 the required hardware (a physical hardware header) to support the Intel® VROC RAID upgrade key. Refer to the [Intel® Virtual RAID on CPU \(Intel® VROC\) Upgrade Key Installation Guide](#) for additional information.

5.4.2 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.

5.4.3 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

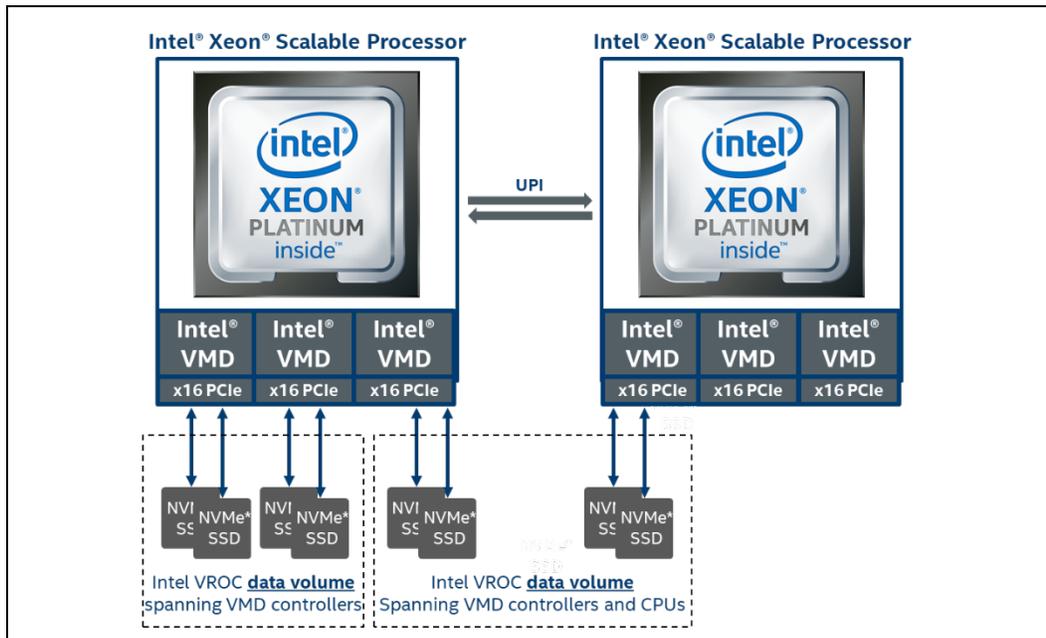


different Intel® VMD domains and can be used for storing data. To enable this option, a selection must be made. This selection field can be found in the Intel® VROC Pre-OS environment (HII and UEFI tools), Windows* OS GUI and in the CLI tool. Note that when this selection is made, a warning message appears to outline the risks of enabling this feature.

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

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

Figure 5-3. Intel® VMD Domains



5.4.3.1 Intel® VROC (VMD NVMe RAID) on Pass-Thru Boot Drives

Intel® VROC 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.

5.4.3.2 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 the 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.
- Intel® VMD has been properly enabled within the BIOS.
- The appropriate Intel® VROC RAID upgrade key (Standard/Premium/Intel-SSD-Only) 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, follow these steps:

- Step 1: Enter the BIOS configuration setup menu to access the Intel® VROC (VMD NVMe RAID) HII user interface.
- Step 2: Navigate to and select **Intel® Virtual RAID on CPU**.
- Step 3: Navigate to and select **Create RAID Volume**.
- Step 4: Type in a volume name and press the <Enter> key or press the <Enter> key to accept the default name. A unique name for each RAID volume is recommended.
- Step 5: Select the RAID level by pressing the <Enter> key to access the menu and scroll through the available values, then press the <Enter> key to select the desired RAID type.
- Step 6: Ensure that the user's vendor supports spanned RAID volumes. To enable it, toggle the status from the default status *off*, to *on*.
- Step 7: 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.
- Step 8: Using the arrow keys to highlight the drives one by one by selecting the < > bracket on the line next to the 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 the drive as part of the RAID.
- Step 9: Repeat the previous step for each drive required in the RAID.
- Step 10: 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 default place this setting at 128k and cannot be modified.
- Step 11: Select the volume capacity and press the <Enter> key. The default value indicates the maximum volume capacity using the selected disks and calculating it at a total of 95% of the value due to disk coercion. To enter a value smaller than the maximum based on the 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 X 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.

Step 12: Navigate to and select **Create Volume**, then press <Enter>.

Step 13: The following message appears: *"You have selected NVMe drives that are connected to multiple Intel® VMD controllers. 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"*. 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.

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%.

Step 14: The user will be presented with a screen reporting on the status of the create process.

- The RAID volume configuration can be verified by looking at the RAID volume information screen. If the RAID volume is 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.

Step 15: 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.

5.4.4 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 Intel® VROC release notes included in the release package.

5.4.5 Intel® VROC (VMD NVMe RAID) Maximum Number of Drives Supported in a RAID Volume

Intel® VROC 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).

5.4.6 Intel® VROC (VMD NVMe RAID) Maximum Number of Switch Levels Supported

Intel® VROC (VMD NVMe RAID) will support the ability to cascade up to two levels of NVMe switches on a single Intel® VMD domain.

5.4.7 Intel® VROC (VMD NVMe RAID) Maximum Number of Arrays Supported

Intel® VROC (VMD NVMe RAID) 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.

5.4.8 Intel® VROC (VMD NVMe RAID) Maximum RAID Members Supported

Intel® VROC (VMD NVMe RAID) will support the ability to create 2 RAID volumes per array on NVMe drives connected to the Intel® VMD controllers.

5.4.9 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 the drives of the RAID volume, the RAID volume may enter into a failed state because of the above-mentioned delay. After all the drives are inserted, if the RAID volume remains in a failed state, an attempt to recover the RAID volume by using the Intel® VROC GUI to reset the volume to normal state should be followed. 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.

5.4.10 Hot Plug Limitation in Windows*

If multiple NVMe SSDs in a system need to be replaced via the hot plug process, the user must give enough time between each process activity. Release one drive at a time. There also must be 45 – 90 seconds (or until the device re-appears / disappears in Device Manager) between each hot plug events. One hot plug event would be the

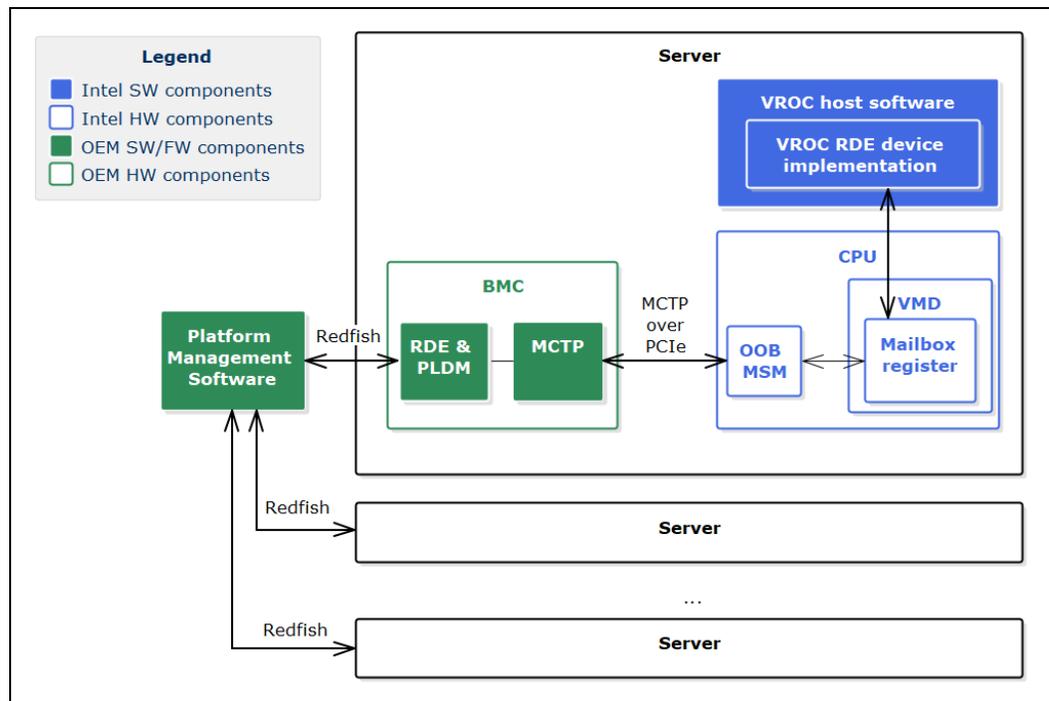
hot-removal of the drive. The next hot plug event would be the hot-insertion of the new drive. simultaneous hot plug events of multiple NVMe drives are not supported.

Surprise Hot Plug Limitations - Due to Microsoft Windows* time restrictions for resuming from S3 and S4, and Intel® VMD device identification requirements, hot plug (insertion and/or removal) of Intel® VROC (VMD NVMe RAID) managed drive, during an S3/S4 low power event, is not supported.

5.4.11 Intel® VROC Out-Of-Band Management (Server Platforms Only)

The Intel® VROC Out-of-Band management interface supports only NVMe drives and Intel® VROC (VMD NVMe RAID) volumes. SATA drives and Intel® VROC (SATA RAID) volumes can be managed through other VROC (in-band) interfaces. In the Intel® VROC OOB management solution, the BMC is the entry point for the requests from a management network. It processes the requests (e.g., management command) and sends them to MCTP endpoint exposed by the Intel® OOB-MSM, via MCTP over PCIe. The Intel® OOB-MSM works as an MCTP bridge, it receives requests from the BMC and sends it to the Intel® VROC OOB host software over the Intel® VMD mailbox register. The Intel® VROC OOB host software component periodically polls the Intel® VMD mailbox register for new requests. The Intel® VROC OOB host software performs the requested action and sends the response to the BMC in the opposite way. A high-level view of the solution is presented in the below figure.

Figure 5-4. Overview of Intel® VROC Out-Of-Band Management Solution



5.4.12 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, see the section on [Various Strip Sizes](#) in this document.

Table 5-5. Recommended Strip Size for Intel® NVMe SSDs for Optimal Wear Leveling

Strip Size \ 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

NOTES:



1. 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.
2. 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.

5.4.12.1 Intel® VROC Optimized Wear Leveling

The Intel® VROC configuration tools (Intel® VROC GUI, Pre-OS HII and CLI tools) will automatically select the strip size to help maximize Intel® SSD wear leveling (when creating a RAID 5 volume).

5.4.12.2 Intel® VROC RAID 5 Default Settings

The following table shows the default values that will be used by the Intel® VROC configuration tools when creating an Intel® VROC RAID 5 volume.

Table 5-6. Intel® VROC RAID 5 Default Settings

Strip Size Drives	Default Value (kiB)
3	64
4	32
5	128
6	64
7	128
8	16
9	128
10	64
11	128
12	32
13	128
14	64
15	128
16	8
17	128
18	64
19	128
20	32
21	128
22	64
23	128
24	16

5.4.13 Intel® VROC (VMD NVMe RAID) 90-Day Trial Period

Intel® VROC (VMD NVMe RAID) 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® VROC 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 an Intel® VROC (VMD NVMe RAID) volume is created.



5.4.13.1 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 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).
- There is no option to disable the Intel® VROC (VMD NVMe RAID) trial period.

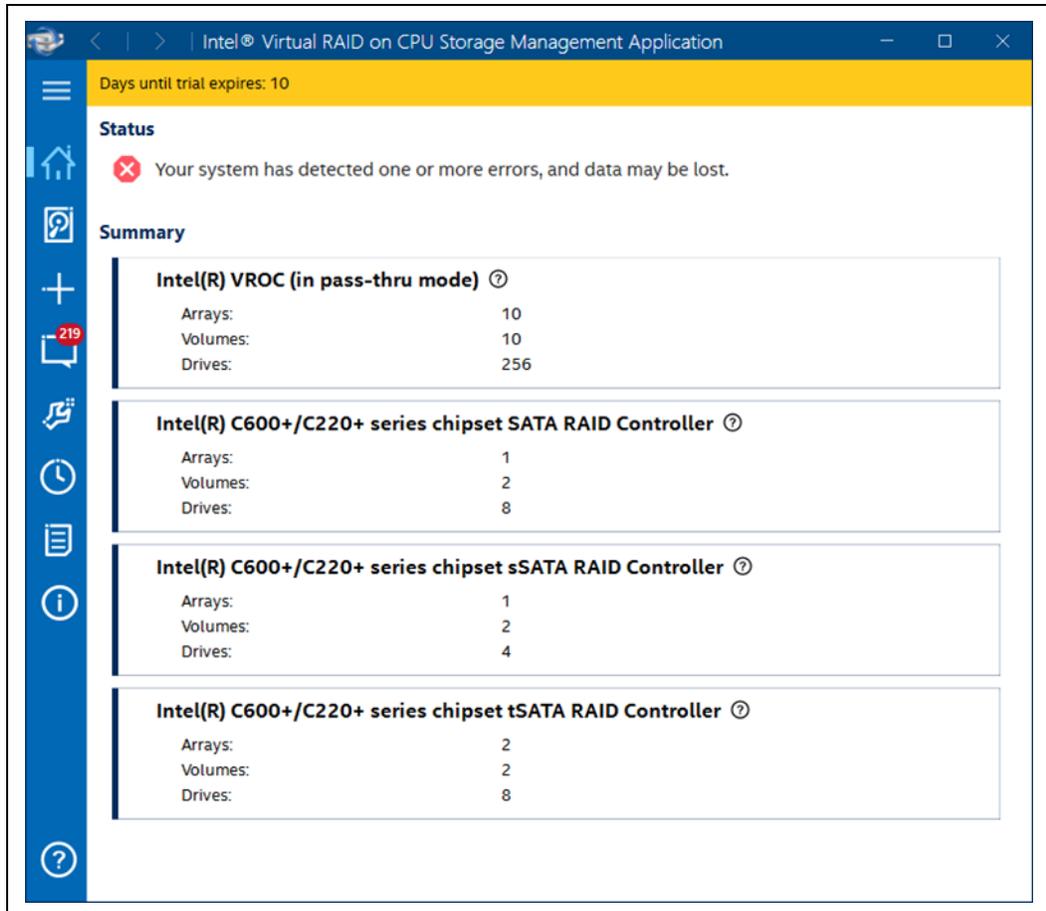
5.4.13.2 Intel® VROC (VMD NVMe RAID) Trial Period Usage

Since the 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 GUI will show no indication of a trial period being activated. Once Intel® VROC 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 GUI as follows:

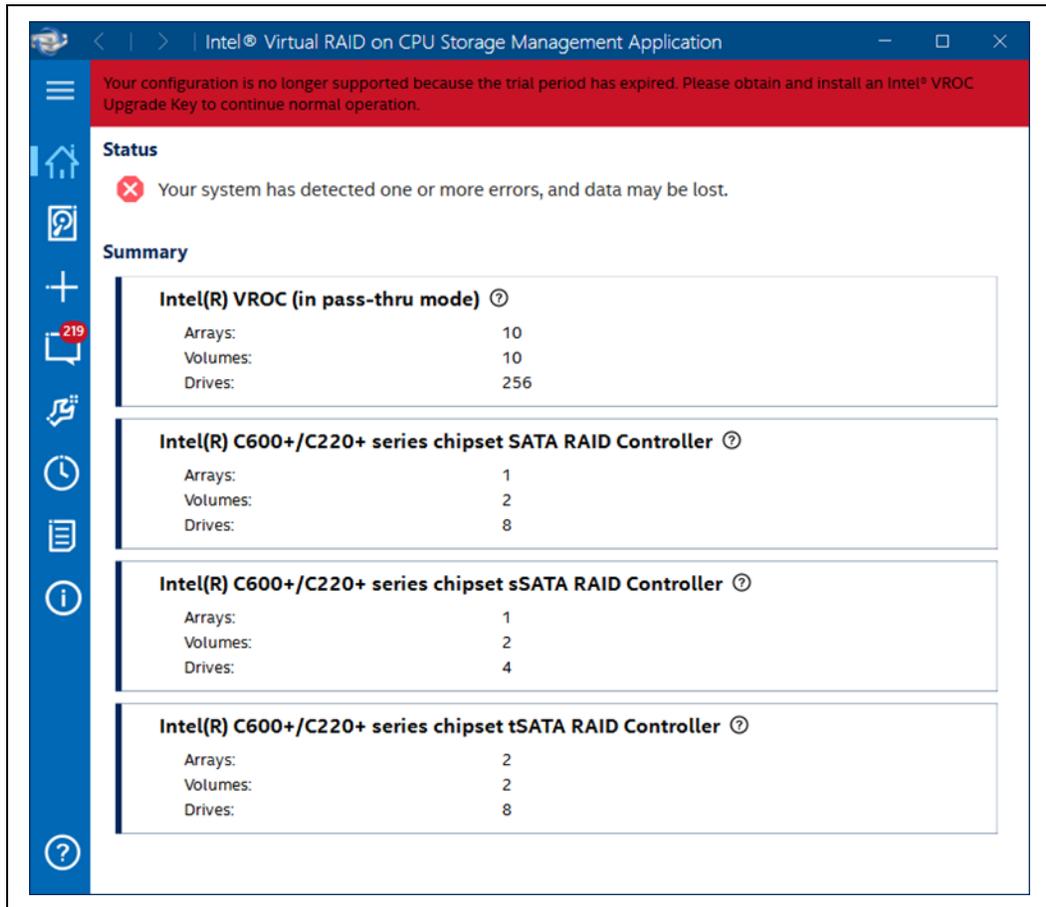
Figure 5-5. Intel® VROC Trial Period Message



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 the 90 days period, the trial period will expire and be displayed in the Intel® VROC GUI as follows:

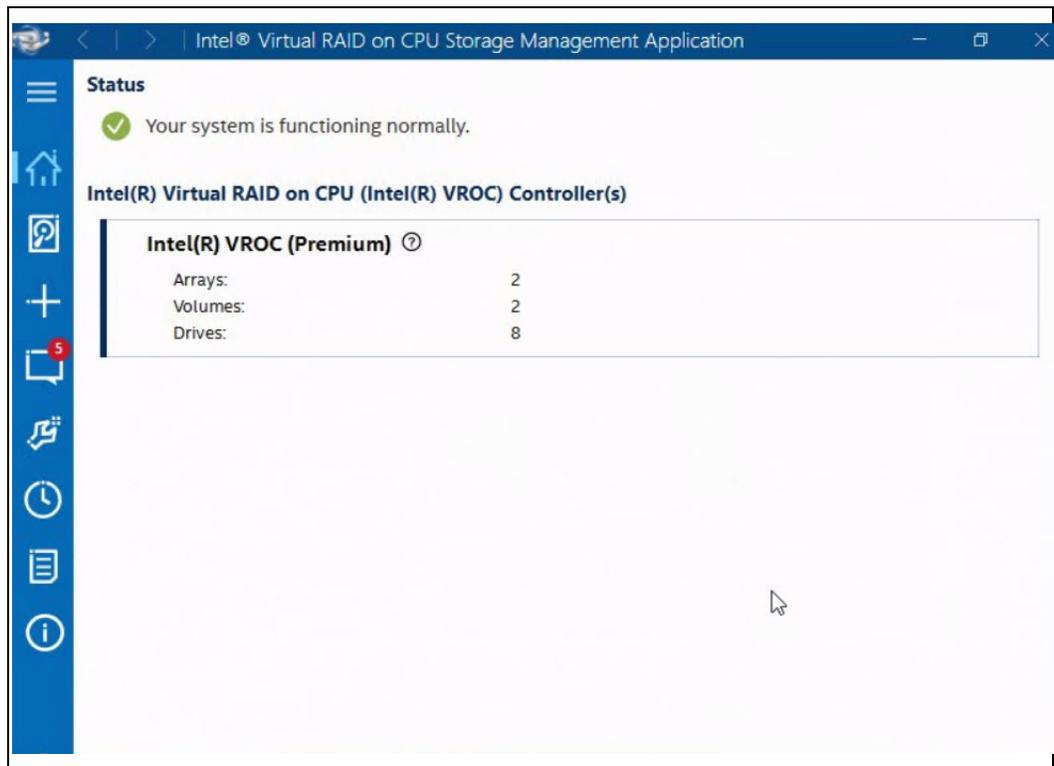
Figure 5-6. Intel® VROC Trial Period Expired Message



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.

Figure 5-7. Intel® VROC Normal Operation After Upgrade Key Installation



5.4.14 Intel® VROC (VMD NVMe RAID) Pre-Load Installation

The Intel® VROC package supports the ability to preload the Intel® VROC (VMD NVMe RAID) OS drivers when the Intel® VMD is disabled. This allows for the Intel® VROC (VMD NVMe RAID) driver to be loaded into the platform, by the Intel® VROC installation application, when the Intel® VMD is disabled. Later, when the Intel® VMD is enabled, the Intel® VROC (VMD NVMe RAID) driver can be installed/updated without having to rerun the Intel® VROC installation application.

5.4.14.1 Intel® VMD Enabled/Disabled when a System Boot Drive Attached

When operating with a system boot drive attached via an Intel® Xeon Scalable CPU PCIe slot, with the Intel® VMD disabled; enabling that Intel® VMD controller requires reinstalling the Operating System to function properly. Not doing so may result in a system failure. The same may apply if the Intel® VMD controller is being switched from an enabled state to disabled.

5.4.15 Intel® VROC Manual Tuning of CPU Utilization

The Intel® VROC package supports the ability for the user to adjust how much CPU utilization can be used by the Intel® VROC RAID engine. When Intel® VROC is installed, the registry key that controls this feature is installed. The default setting of



this registry key is 10. The range of values supported is 1 to 16 threads. If the user selects a value outside of that range, the Intel® VROC driver will automatically revert to the default value of 10. The Windows* registry key is the following:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\iaVROC\Parameters\Device\Threads
```

This registry setting means that the Intel® VROC RAID engine cannot exceed more than that number of CPU threads (in the default case, 10). That means that if further action is required, the Intel® VROC RAID engine will wait until one of those threads completes and is freed up. The default value was chosen because it is shown to be the most efficient in performance testing. However, every workload and setup are different. This allows OEMs/ODMs (and end users) to tune this value for their usage.

Note: Intel does not recommend changing this registry key unless the user has a full understanding of 1) how to correctly modify a registry key; 2) why the value is being changed; and 3) its corresponding impact on the platform operation. If the number of threads is needed to be changed it can be done in the registry to a maximum of 32.



6 Intel® VROC Common Features

These are the common features between the Intel® VROC (VMD NVMe RAID) and Intel® VROC (SATA RAID) packages.

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

Table 6-1. Intel® VROC Common Features

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* • Online 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 • Physical Drive Location
Utilities	<ul style="list-style-type: none"> • Install/Uninstall Utility • Configuration and Management Utilities 	

6.1 Intel® VROC RAID Features and Functionality

This section outlines those features and functionality supported with the Intel® VROC family of products.

6.1.1 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) /tSATA(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 *the 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 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 $N-1 \times$ the smallest size of the drives, where $N \geq 3$ and:

- \leq max number of drives supported by SATA or sSATA or tSATA 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 of 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 and more efficient requests, resulting in full stripe writes from the cache to the RAID 5 volume.

For example, a three-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 RAID 5 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.

6.1.2 Matrix RAID

The Intel® VROC 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 the remaining available space.

Note: When utilizing a Matrix RAID configuration, avoid mixing a non-redundant RAID volumes (RAID 0) with redundant RAID volumes (RAID 1/5/10). Any failed disk will fail the RAID 0 volume and may result in the array not functioning properly.

6.1.3 Creating a RAID Volume

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

- The Intel® VROC Pre-OS environment
 - Intel® VROC Pre-OS 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)

Note: Starting with Intel VROC 8.5 PV release, volume names must conform to the below character set in all operating environments:

- Uppercase A – Z
- Lowercase a – z
- Numbers 0 – 9
- Period (.)
- Underscore (_)
- Hyphen (-)

Additionally, the hyphen (-) character cannot be the first character of any volume name.

If the user tries to create a volume using invalid characters, the operation will be blocked by Intel VROC and a message will be displayed which indicates which characters from the desired name are not compliant.

Volumes created before Intel VROC 8.5 PV will not be affected by this change, even if they contain illegal characters in their name.



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

The configuration tools listed in section [Intel® VROC \(VMD NVMe RAID\) Environment](#) 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, and RAID 10 levels. Strip size is not applicable (N/A) for RAID 1.

The following command will create a RAID volume using all 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 will display usage help for all support command line parameters:

```
C:\>Rcfgxxxx.efi /?  
C:\>IntelVROCCLI.exe --help
```

6.1.3.2 Using the Intel® Virtual RAID on CPU GUI

This example assumes that Intel® VROC has been installed and the drives are connected to the CPU via the Intel® VMD.

The following example will step through the process of creating a two-drive RAID 1 data volume:

Step 1: Start the Intel® Virtual RAID on CPU GUI application (run as administrator):
Start Menu->All Programs->Intel->Intel Virtual RAID on CPU

Step 2: From the **Home** screen, click the **Create Volume** button.

Step 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.

Step 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 one: **Volume_0000**. Users can modify the RAID volume name and assign a unique name for each RAID volume.

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 example, 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* information under the help file in the GUI.

Step 5: In the confirmation 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:

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

Step 6: 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. Click **OK** to close this dialog box.

Step 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 (at the right) will refresh to show the status, properties, and available options of the newly created RAID volume.

Note: In the case where the RAID volume consists of hard drives or hard disk drives (HDD), both disk data cache and volume write back cache can be enabled to speed up data access, especially during RAID volume rebuilding. For disk data cache, please refer to section 6.3.6.7 and for Volume write-back cache details, please refer to section 6.3.6.16.

6.1.4 Instant Initialization

The Intel® VROC 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 product will initialize the following Intel® VROC RAID volumes upon resuming from a dirty shutdown:

- RAID 1 volumes that has not previously been initialized.
- RAID 5 volumes that have not previously been initialized with the *Array Cache* feature enabled.

Note: For best overall RAID 5 performance, Intel recommends that RAID 5 volumes be initialized prior to running any performance testing.

6.1.4.1 Intel® VROC Automatic Initialization

The following table outlines the scenarios that will result in the RAID volume initialization starting automatically as a result of a platform dirty shutdown:

Table 6-2. Intel® VROC RAID Volumes Automatic Initialization

RAID Level	RAID Array Cache	RAID Volume Initialized	RAID Volume Status after Dirty Shutdown
RAID 0	On	N/A	Normal
	Off	N/A	Normal
RAID 1	On	No	Initializing
	Off	No	Initializing
	On	Yes	Verify & Repair
	Off	Yes	Verify & Repair
RAID 5	On	No	Initializing
	Off	No	Normal
	On	Yes	Verify & Repair
	Off	Yes	Normal
RAID 10	On	No	Initializing
	Off	No	Initializing
	On	Yes	Verify & Repair
	Off	Yes	Verify & Repair

6.1.4.2 Intel® VROC RAID Volume Initialization

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

Step 1: Start the Intel® Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 2: From the **Volumes** pane, click the array or volume to which the user wants to modify.

Step 3: The volume properties now display on the right. 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.

6.1.5 Deleting a RAID Volume

RAID volumes can be deleted in three different ways. The method most widely used by end-users is by using the Intel® VROC 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, any existing data on all disks that are part of the selected volume will be permanently lost. We recommend backing up all valuable data before continuing.

6.1.5.1 Using the Intel® VROC GUI

To delete a RAID volume using the Intel® VROC GUI follow these steps:

Step 1: As administrator, start the Intel® Virtual RAID on CPU GUI application:
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 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.

Step 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.



6.1.5.2 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 from an Intel® Customer Reference Board (CRB).

Step 1: Enter the setup and configuration mode of the BIOS.

Step 2: Navigate to **EDKII Toolkit** and press the <Enter> key.

Step 3: Navigate to **Intel® Virtual RAID on CPU** and press the <Enter> key.
Alternate destinations may be **Intel® VROC sSATA/tSATA Controller** or **Intel® VROC SATA Controller**.

Step 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.

Step 5: On the page that shows all the RAID volume information, use the arrow keys to navigate to select the **Delete** option and press the <Enter> key.

Step 6: Use the arrow keys to navigate to select **Yes** and press the <Enter> key.

6.1.5.3 Using the Intel® VROC CLI Tool

Use the 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 ports:

```
C:\> IntelVROCCli.exe -D -n OEMRAID0
```

6.1.6 Degraded RAID Volumes

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

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

6.1.7 Intel® VROC GUI Support of Various Controllers

The Intel® VROC 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/tSATA, Intel® VMD controllers and supported drives attached to them. The Intel® VROC GUI management tool will display each different controller and show all disks attached to that controller. The Intel® VROC GUI will allow creating, deleting, and monitoring devices, RAID arrays and volumes on drives connected to the SATA, sSATA, tSATA and Intel® VMD controllers.

6.1.8 Various Strip Sizes

The Intel® VROC family of products will provide the ability to change strip size on existing volumes (migration required). Intel® VROC 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 is in Kilobytes.

The following table lists the available strip size configurations.

Table 6-3. Intel® VROC Strip Size Configurations

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

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 (N/A) for RAID 1 due to its redundant configuration. The default value for all the RAID types listed in the above table is the recommended strip size based on the system configuration and the volume type selected; changing the pre-selection for RAID 0, 5 and 10 is best suited for advanced users.

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

Table 6-4. Usage Scenarios for 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, 5, 10
8 KB	Best for databases (fast read transfer rate with faster write transfer rate than with 4KB strips).	RAID 0, 5, 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, 10

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

6.1.8.1 Setting the Strip Size When Creating a Volume

Step 1: Under **Home**, click **Create**.

Step 2: Select the volume type, and then click **Next**.

Step 3: Make the required disk selection, and then select a new data strip size from the drop-down list in the **Advanced** section.

Step 4: Complete the volume creation process as described in previous sections.

6.1.8.2 Setting the Strip Size When Changing Volume Type

Step 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.

Step 2: Click **Change type**.

Step 3: Make the necessary volume type and disk selections, and then select a new data strip size.

Step 4: Click **OK** to change the type of the existing volume. The **Manage** page refreshes and reports the new volume configuration.

6.1.9 Disk Coercion

Intel® VROC 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 UI's (Pre-OS, Windows* GUI and CLI tool) will provide an option to manually override this to be able to use all available disk space.

6.1.10 Minimum RAID Volume Size Recommendation

Intel® VROC family of products recommends a minimum RAID volume size of 100 Megabytes. Creating a RAID volume smaller than 100 Megabytes may result in unexpected behavior.

Note: With such small sized volumes, actions resulting in rebuild or migration may not be reflected in the VROC GUI. Such actions can take just milliseconds and the refresh rate of the GUI to collect volume/array/drive information is every second so it may not be able to catch all the notifications. This issue only effects the VROC GUI display of

notifications. The operation of the volume is not affected and the actions perform correctly.

6.1.11 Flush Unit Access (FUA)

The Intel® VROC 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 making the change to know how to open the Windows* Registry Editor, how to modify the entries, and properly save the updates.

6.1.12 RAID Volume Roaming

The Intel® VROC family of products will support the ability to move all 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.

6.1.13 RAID Volume Roaming between Linux* and Windows*

The Intel® VROC 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.

6.1.14 Support of Pass-Through Drive

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



6.1.15 Failed Drive Reinsertion

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

6.1.16 LED Management

The Intel® VROC 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.

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

- NVMe LED management via Intel® VROC.
- SATA LED management via the PCH controller (SATA/sSATA/tSATA) in RAID mode.

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

The LEDs management capabilities are outlined as follows:

Table 6-5. LED Management Capabilities

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

<p style="text-align: center;">Drive Fails</p>	<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 <p>The drive Bad Block management table is full</p>	<p>The Fail pattern will be illuminated as follows:</p> <p>From the time a failed drive is identified until it is:</p> <ul style="list-style-type: none"> - physically removed <p>OR</p> <ul style="list-style-type: none"> - the RAID volume, that contains the failed drive, is either deleted or physically removed. <p>From the time a</p> <ul style="list-style-type: none"> - non-Failed drive that is part of a RAID volume is removed <p>OR</p> <ul style="list-style-type: none"> - 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 <p>OR</p> <ul style="list-style-type: none"> - the platform is rebooted 	<p style="text-align: center;">Option 1</p>
<p style="text-align: center;">Managed Unplug</p>	<p>When a managed hot unplug is being administered, the status LED of that drive will blink in the defined Locate pattern until the drive is physically ejected.</p>	<p style="text-align: center;">None</p>	<p style="text-align: center;">Enabled</p>
<p style="text-align: center;">RAID Volume Initialization or Verify and Repair Process</p>	<p>When a RAID volume is in either the <i>Initialization</i> or <i>Verify and Repair</i> processes, the status LEDs will blink in the defined Rebuild pattern on all drives in the RAID volume until the process completes.</p>	<ol style="list-style-type: none"> 1. Enabled – status LEDs will blink in the Rebuild Pattern 2. Disabled – status LEDs will not blink 	<p style="text-align: center;">Enabled</p>



<p>RAID Volume Rebuilding Options</p>	<p>When a RAID volume is in a Rebuild state, the status LED(s) will blink in the defined Rebuild pattern on an either the specific drive being rebuilt or on the entire RAID volume rebuilt.</p>	<ol style="list-style-type: none"> 1. Disabled (only on one drive) 2. Enabled (on all drives) 	<p>Disabled (only on one drive)</p>
<p>RAID Volume is Migrating</p>	<p>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.</p>	<ol style="list-style-type: none"> 1. Disabled (No Status LED Blinking) 2. Enabled (Blinks Status LEDs) 	<p>Enabled</p>

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

Note that the registry key line items below are an example for the SATA controller. For sSATA controller, please change to “\iaStorB\”. For tSATA controller, please change to “\iaStorT\”. And for VROC VMD controllers supporting NVMe devices, please change to “\iaVROC\”. The settings for each keys below are the same for all controllers.

- 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
```

- This 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 registry key set to a value of 0x00000001 will enable this configuration option and will blink all the status LEDs of the RAID volume being rebuilt.
- The registry key that is used to customize the configuration option for the RAID volume that is migrating 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 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 LED Management does not include drive activity LED management.

6.1.16.1 Failed Drive Indicator

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

6.1.16.2 Activate LED

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

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

Step 1: Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 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**.

Step 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/tSATA controller, LED management is only supported when these controllers are in RAID mode.



6.1.16.3 Rebuild LED

The Intel® VROC 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.

6.1.16.4 LED Patterns Conflict Precedence

When two or more different LED patterns are applicable to the same drive, the Intel® VROC Windows* driver shall apply the following order in terms of precedence:

1. Locate LED pattern.
2. Fault LED pattern.
3. RAID rebuild LED pattern.

6.1.16.5 LED Management on PCH-VMD

There are 2 PCIe ports designed and controlled under the Server platform PCH, since Emmitsburg.

6.1.16.6 UEFI LED Management

The Intel® VROC 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 UEFI HII. *Rebuild* and *Fail* blink patterns will occur automatically based on the condition of the RAID volume state or member drives state.

Note: When performing the *Locate* functionality in the UEFI environment, the process is to manually activate and then manually disable. The LED will continue to blink until the process is manually stopped.

6.1.17 Modify Volume Name

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

6.1.18 Hot Plug (Surprise and Managed)

The Intel® VROC family of products will support the ability to hot plug (remove and replace) disk drives on properly configured and operational platforms whether or not the 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 or tSATA controllers.

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

Step 1: Open the Intel® VROC GUI:

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

Step 2: In the **Devices** pane, select the NVMe drive to be removed.

Step 3: In the **Disk Properties** pane, look for **Location type** and click on **Remove Disk**.

Step 4: A message window will pop up asking *Are you sure you want to remove this disk?* Click **Yes**.

Intel® VROC 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. Consult the platform design documentation for additional information.

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.

6.1.18.1 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, or having backup drives ready in case a failure occurs, enabling *Auto Rebuild* to assist in the recovery.

6.1.18.1.1 Enabling Surprise Hot Plug for Intel® VROC

The following steps are an example of how to enable surprise hot plug in the Intel® Customer Reference Board (CRB) platform BIOS:

Step 1: From the boot options menu, select the option that will allow the user to enter the BIOS setup.

Step 2: With the **EDKII Menu** highlighted, press <Enter>.

Step 3: Navigate to **Socket Configuration**, then press <Enter>.



Step 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. Consult the user's platform documentation because each BIOS is different, and the steps taken to enable this feature may be different.

6.1.18.1.2 Enabling Surprise Hot Plug for PCH

The following steps are an example of how to enable surprise hot plug in the Intel® Customer Reference Board (CRB) platform BIOS:

Step 1: From the boot options menu, select the option that will allow the user to enter the BIOS setup.

Step 2: With the **EDKII Menu** highlighted, press <Enter>.

Step 3: Navigate to **Platform Configuration**, then press <Enter>.

Step 4: Navigate to **PCH Configuration**, then press <Enter>.

Step 5: Navigate to **PCH SATA** or **PCH sSATA** or **PCH tSATA 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. Consult the user's platform documentation because each BIOS is different, and the steps taken to enable this feature may be different.

6.1.18.2 Intel® VROC Pre-OS Hot Plug Support

The Intel® VROC family of products will support the ability to perform drive hot plugs in the Pre-OS environment when the drives are in a pass-through configuration (not in a RAID volume). The Intel® VROC HII does not dynamically update, and the screen must be refreshed (by going up a menu level and coming back in) to see the hot plug transition.

Note: Intel® VROC does not currently support Pre-OS Hot Plug on drives within a RAID volume.

6.1.19 Hot Spare Disk

The Intel® VROC 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.

6.1.19.1 Global Hot Spare

The Intel® VROC 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 sSATA/tSATA controller will be available for any RAID array on the sSATA/tSATA controller.

Note: 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. The one exception is a boot volume. The spare must be in the same domain. Limitation: A spanned RAID volume can use hot spare disk inside one of its VMD controller. A non-spanned RAID volume will not auto use hot spare disk not in the same VMD controller.

6.1.19.2 Mark a Disk as Hot 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. As a result, make sure the size if the drive is commensurate to the RAID array it targeted for.

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

- From the Intel® VROC GUI
- From the Intel® VROC Pre-OS environment (UEFI HII) user interface support
- From the Intel® VROC CLI tool

6.1.19.2.1 Mark a Disk as Spare in the Intel® VROC GUI

The Intel® VROC 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:

Step 1: Start the **Intel Virtual RAID on CPU GUI** application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 2: In the **Devices** pane that displays the drives, select the drive to be marked as a spare.



Step 3: In the **Disk Properties** pane click on **Mark as Spare**.

6.1.19.2.2 Mark a Disk as Spare in the UEFI HII

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

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. To identify a drive as a hot spare in the UEFI HII use the following steps:

Step 1: Enter the setup and configuration mode of the BIOS.

Step 2: Navigate to **EDKII Toolkit** and press the <Enter> key.

Step 3: Navigate to **Intel® Virtual RAID on CPU** and press the <Enter> key. Alternate destinations may be **Intel® VROC sSATA/tSATA Controller** or **Intel® VROC SATA Controller**. The non-RAID disk will be listed below the RAID volume in a category below.

Step 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**.

Step 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.*

Step 6: To proceed, highlight **Yes**, and press <Enter>.

6.1.19.2.3 Return Hot Spare Disk Back to Available

The Intel® VROC 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 the Intel® VROC GUI

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

Step 1: Start the **Intel Virtual RAID on CPU GUI** application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 2: In the **Devices** pane that displays the drives, select the drive that is marked as a spare.

Step 3: In the **Disk Properties** pane click on **Return to available**.

Return Hot Spare Disk Back to Normal in the UEFI HII

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

Step 1: Enter the setup and configuration mode of the BIOS.

Step 2: Navigate to **EDKII Toolkit** and press the <Enter> key.

Step 3: Navigate to **Intel® Virtual RAID on CPU** and press the <Enter> key. Alternate destinations may be **Intel® VROC sSATA/tSATA Controller** or **Intel® VROC SATA Controller**. The non-RAID disk will be listed below the RAID volume in a category below.

Step 4: Highlight the target drive that is designated as a hot spare and press the <Enter> key. This screen presents the **Physical Disk Information** page. Select **Reset** to available and press the <Enter> key.

Step 5: To proceed, highlight **Yes**, and press <Enter>.

6.1.20 Drive Requirements

This section outlines the requirements and expectations for drives used in an Intel® VROC RAID array/volume.

6.1.20.1 Clean Drive Status

Intel® VROC 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 rebuild process.

6.1.20.2 Minimum Disk Size Requirements

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

6.1.20.3 Disk Addition: Minimum Required Disk Size

The Intel® VROC 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.

6.1.20.4 Targeted Disk Replacement: Minimum Required Disk Size

The Intel® VROC 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.



6.1.21 Intel® VROC Rebuild

The Intel® VROC family of products supports RAID rebuild when a RAID volume is degraded. Intel® VROC supports two types of RAID volume rebuilds: *Automatic* and *Manual*.

6.1.21.1 Intel® VROC Rebuild Components

The components of the Intel® VROC family of products that are directly involved in the rebuild process includes the UEFI driver and the OS driver.

6.1.21.1.1 UEFI Rebuild

As part of the RAID volume rebuild process, the Intel® VROC UEFI driver only manages the Intel® VROC metadata and displays the information to the user accordingly. While the system is booting, if the Intel® VROC UEFI driver determines that the RAID volume is either in (or needs to be placed into) a rebuild state, it will make the necessary modifications and displayed that information to the user. The Intel® VROC UEFI driver does not, however, perform the rebuild process.

6.1.21.1.2 OS Rebuild

The Intel® VROC OS driver will update the metadata when necessary, and will perform the actual rebuild process (comparing, computing, and performing the data and metadata updates).

6.1.21.2 Auto Rebuild

The Intel® VROC 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 enough 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 the member drive 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:

Table 6-6. Intel® VROC Auto Rebuild Functionality

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

6.1.21.2.1 Auto Rebuild to a Re-Inserted RAID Member

The Intel® VROC family of products will support the ability to automatically initiate a 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.

6.1.21.2.2 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 will claim the spare drive for that RAID volume and automatically begin to rebuild the RAID volume.

6.1.21.2.3 Auto Rebuild on Hot Insert

The Intel® VROC 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:
 - Drive removed,
 - 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).

Note: If the drive inserted meets the size requirements (and Rebuild on Hot Insert is enabled) and the new drive is part of another RAID volume (contains RAID metadata) the contents of that drive will be removed, the drive will become part of the array and the RAID volume rebuild will begin.



The Intel® VROC GUI shall give the user the ability to enable and disable this feature following these steps:

Step 1: Start the **Intel Virtual RAID on CPU GUI** application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.

Step 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).

Step 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 [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.

6.1.21.2.4 Error Threshold Monitoring/Handling

The Intel® VROC 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).

6.1.21.3 Manually Invoked Rebuild

The Intel® VROC 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:
 - Drive being replaced,
 - 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).

6.1.21.3.1 Rebuild Resumption

The Intel® VROC family of products will support continuing a RAID volume rebuild after a system shutdown, sleep state or hibernation state.

6.1.21.3.2 Auto RAID 0 Rebuild

The Intel® VROC family of products will support the ability to perform an automatic rebuild on a RAID 0 volume when a SMART event is encountered, and a global hot

spare drive has been defined. This feature is automatically enabled when a global hot spare has been defined for a RAID array that contains a RAID 0 volume(s).

Note: Global hot spare rules apply (drive size and controller/domain).

With an appropriate disk defined as a global hot spare, a RAID 0 volume will automatically begin a RAID volume rebuild when one of the RAID 0 member drives encounters a SMART event. The spare drive will be brought into the RAID volume to replace the drive with the SMART event and the content will be copied to the new drive.

6.1.22 Verify and Repair

The Intel® VROC 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. It then reports the number of inconsistencies, or the 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:

Table 6-7. Verify and Repair Process in Different RAID Levels

RAID Level	Verify	Verify and 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 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.



6.1.23 Starting Verify and Repair Process

The following section outlines the different methods with which the *Verify and Repair* process either starts automatically or is manually initiated.

6.1.23.1 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. The following table outlines the scenarios that will result in the RAID volume starting the *Verify and Repair* process automatically as a result of a platform dirty shutdown:

Table 6-8. Verify and Repair Process After Dirty Shutdown

RAID Level	RAID Array Cache	RAID Volume Initialized	RAID Volume Status After Dirty Shutdown
RAID 0	On	N/A	Normal
	Off	N/A	Normal
RAID 1	On	No	Initializing
	Off	No	Initializing
	On	Yes	Verify and Repair
	Off	Yes	Verify and Repair
RAID 5	On	No	Initializing
	Off	No	Normal
	On	Yes	Verify and Repair
	Off	Yes	Normal
RAID 10	On	No	Initializing
	Off	No	Initializing
	On	Yes	Verify and Repair
	Off	Yes	Verify and Repair

6.1.23.2 Manual Initiation

The Intel® VROC 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 the RAID volume data verification and repair process:

- Step 1: Start the **Intel Virtual RAID on CPU GUI** application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel Virtual RAID on CPU.
- Step 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.

Step 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*.

The verification progress is shown in 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.

6.1.24 Read Patrol

The Intel® VROC 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, sSATA and tSATA 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/Os are being processed.

Note: *Read Patrol* does not extend to drives marked as hot spares.

The Intel® VROC 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 Intel® VROC driver will record the event but not store all errors across all RAID volumes being scrubbed. The Intel® VROC driver will display the following read patrol error information: *Unrecovered LBAs and its corresponding RAID volume and physical disk*.

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

6.1.25 Check Pointing

The Intel® VROC 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.

6.1.26 Bad Block Management

The Intel® VROC 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.

6.1.26.1 Marking a Bad Block

When Intel® VROC is performing a READ operation and encounters a Bad Block condition on one of the member drives, it will mark the same sector on the other RAID member drive(s) as a bad block sector and reflect it in the Intel® VROC metadata.

6.1.26.2 Full Bad Block Table Rebuild

When Intel® VROC is performing a RAID volume rebuild and encounters a full *Bad Block Table*, it will stop rebuilding the RAID volume at that point and will leave the RAID volume in a Degraded state.

6.1.26.3 Degraded Read

When Intel® VROC is performing a READ operation and encounters a Bad Block condition, that READ command response will indicate that a failure has occurred.

6.1.26.4 Read Patrol Bad Block Management

Intel® VROC (SATA RAID) will support *Bad Block Management* while the system is performing Intel® VROC (SATA RAID) Read Patrol activities. The Intel® VROC (SATA RAID) Read Patrol feature is disabled by default and must be manually enabled via the Intel® VROC GUI. During the Intel® VROC (SATA RAID) Read Patrol process, if a Bad Block is encountered, an attempt will be made to correct the Bad Block condition.

6.1.26.5 Multiple Bad Blocks Intel® VROC (SATA RAID) Read Patrol

When the Intel® VROC encounters multiple Bad Blocks in the corresponding RAID volume member (corresponding LBA's), the RAID volume state will not be changed if encountered during one of the following processes:

- Intel® VROC (SATA RAID) Read Patrol
- Intel® VROC Initialization
- Intel® VROC Verify and Repair

6.1.26.6 Bad Block Management Logging

Intel® VROC will log Bad Block conditions found in the Windows* Event Log. Review the [Disk Events](#) for those events that will be logged.

6.1.26.7 Bad Block Management Recovery

Intel® VROC will attempt to recover from a Bad Block condition whenever possible. This includes when Bad Block Management is performed during Read Patrol, RAID volume initialization and the Verify and Repair processes.

6.1.26.8 Intel® VROC GUI – Bad Block Reporting

Intel® VROC GUI will report any Bad Block activity logged through the Intel® VROC GUI.

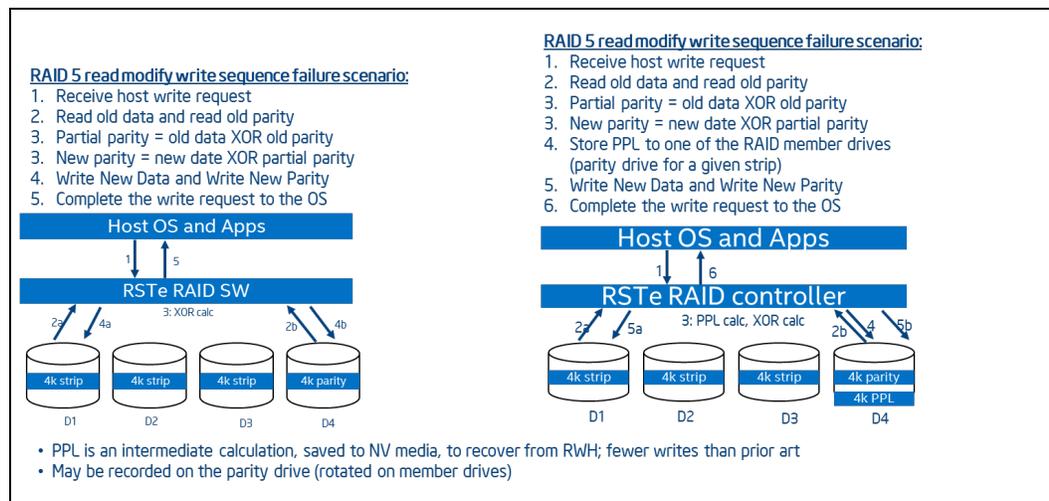
6.1.27 Intel® VROC RAID Write Hole Closure

The Intel® VROC 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.

Figure 6-1. Intel® VROC RAID Write Hold Closure Example



RWH is a potential condition when using RAID 5 that during the time that a RAID strip is being written, a power loss is encountered. With hardware RAID and battery backed DRAM; logging could be used to recover. With Intel® VROC and RWH, a journaling SSD can be added to preserve the Partial Parity and reduce the potential data loss issue.



6.1.27.1 RAID Write Hole Closure (RWH) Details

The Intel® VROC family of products will support the option to configure this feature in the Intel® VROC 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.

6.1.27.1.1 RAID Write Hole Closure (RWH) in Windows*

The Intel® VROC 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 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.

6.1.27.1.2 RAID Write Hole Closure (RWH) in Pre-OS

The Intel® VROC UEFI drivers will support the ability to recover from the RAID 5 volume invalid state caused by an 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*.

6.1.27.1.3 RWH Backwards Compatibility

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

6.1.27.2 RWH Recovery

The Intel® VROC drivers enter an RWH recovery state when a RAID 5 volume encounters an 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 an 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 disabled. The *Verify and Repair* process is covered by the RWH recovery process.

6.1.28 Controller Default Values

The Intel® VROC family of products will 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 package, to define the default behavior. The second method is to use the Windows* registry key option.

6.1.28.1 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|tSATA|VMD
```

Where the **controllerName** will be the name of the controller as shown when using the **-I** option of CLI command. Either **--controller controllerName** or **--controllerMode SATA|sSATA|tSATA|VMD** can be used. 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%, 0x00000000 (disabled) | 0x00000001(enabled)
```



6.1.28.2 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|tSATA
```

Where the **controllerName** will be the name of the controller as shown when using the **-I** option of CLI command. Either **--controller controllerName** or **--controllerMode SATA|sSATA|tSATA|VMD** can be used. 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)
```

6.1.29 Driver Dirty Shutdown (DS) Recovery

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

6.1.29.1 Driver DS Recovery Caused by Hot-Plug

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

6.1.29.2 Driver DS Recovery Caused by Driver Start

The Intel® VROC 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.

6.1.29.3 Driver DS Recovery Caused by Enabling Drives

The Intel® VROC drivers will permanently recover from the RAID 5 volume invalid state caused by a *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.

6.1.29.4 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.

6.1.29.5 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 RAID 5 I/O request.

6.1.29.6 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.

6.1.29.7 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.

6.1.29.8 Dirty Stripe Journaling

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

6.1.29.9 Partial Parity Logging (PPL)

The Intel® VROC 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.

6.1.29.10 Reading/Changing Controller Parameters

The Intel® VROC Command Line Interface (CLI) can be used to read and/or change the following two controller parameters:

Rebuild on Hot Insert (ROHI):

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

Read Patrol:

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

Both the Read Patrol and ROHI states are stored in a registry in the same location as registry for LEDs management as *ReadPatrol* and *RebuildOnHotInsert*.

Note: Parameter changes only applies to the OS environment and will not be seen or honored in the Pre-OS environment. The Pre-OS will treat these parameters as disabled.



6.2 Intel® VROC Pre-OS Features and Functionality

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

6.2.1 Intel® VROC Unified Extensible Firmware Interface (UEFI) Drivers

The Intel® VROC family of products will 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 UEFI images included support for Intel® VMD and SATA/sSATA/tSATA controllers set to RAID mode.

6.2.1.1 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 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.

Intel® VROC 7.6 Pre-OS is also compliant with version 2.7.

Table 6-9. 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 Version 1.28	https://edk2-docs.gitbooks.io/edk-ii-build-specification/content/#edk-ii-build-specification

6.2.1.2 Intel® VROC 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® Virtual RAID on CPU* or *Intel® VROC SATA Controller* or *Intel® VROC sSATA Controller* or *Intel® VROC tSATA 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 UEFI driver **FORMSET_GUID** is:

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

6.2.2 UEFI System BIOS Requirements for Platform Compatibility with Intel® VROC UEFI

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

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

The Intel® VROC 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))

6.2.2.2 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

6.2.2.3 Protocols Provided by the Intel® VROC UEFI Drivers

The Intel® VROC 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 Pass Through drives, RAID members, Spare, Journaling and SED drives (SED drive's authentication encryption key are managed by VROC SED driver. For more details, please refer to document #729830)
 - When installing this protocol for RAID member drives (SATA, and NVMe drives), VROC driver will detect drive's state changes - all possible changes between Locked, Unlocked and Unencrypted states. And a re-enumeration execution after state change to avoid RAID failure (or degrading) when the RAID members are unlocked one-by-one.
 - Two new properties: *Encryption Ability* (None/ Other/ SED), and *Encryption Status* (Unencrypted/ Locked/ Unlocked (displayed only if Encryption Ability is Other or SED)) are added for every drive in all VROC user interfaces (UEFI HII, RCfg, Windows application GUI, CLI and OOB). For drives in Locked state, block all actions (Ex. Mark as spare, Create volume, and other drive/volume actions) that uses this drive.

—The following changes are also required: VROC/VMD UEFI drivers to implement Level 0 Discovery and ATA Identify discovery in the drive's enumeration flow and after each *SSCP.ReceiveData* function call (to refresh the state)

—VROC Windows driver: implement *Level 0 Discovery* and *ATA Identify Discovery* in the drive's enumeration flow.

- **EFI_EXT_SCSI_PASS_THRU_PROTOCOL:**

—All SCSI commands are supported (for ATAPI devices)

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

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

6.2.2.4.1 Platform UEFI BIOS

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

6.2.2.4.2 Download and Integrate the Intel® VROC UEFI Package

Step 1: Download the latest kit from the Intel® VIP (Validation Internet Portal) website.

Step 2: From the kit, select the Pre-OS zip file which will contain all the Intel® VROC and Intel® VROC (SATA RAID) Pre-OS images and tools.

Step 3: Go to the *efi_standalone_vroc_rs* (Intel® VROC) or *efi_sata*, *efi_ssata* and *efi_tsata* (Intel® VROC (SATA RAID)) directories to find the UEFI driver binary file.

Step 4: 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.
- tSATADriver.efi – Include this binary when building the BIOS image to support Intel® VROC (SATA RAID) tSATA Controller.
- tSATADriver.depex – Include this file when building the BIOS image to support Intel® VROC (SATA RAID) tSATA Controller.

Step 5: Use the proper integration tools based on the binary file selected above.

6.2.2.4.3 Verify Compliance

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

6.2.2.5 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.

6.2.2.5.1 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-8BD8-44A1-9D06-08039DFDEC63
- *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.*

6.2.2.5.2 EXT_SCSI_PASSTHRU_PROTOCOL Test

`EFI_ATA_PASS_THRU_PROTOCOL.BuildDevicePath` - call `BuildDevicePath` with `NULL DevicePath`.

- *Test Index:* 5.7.8.2.1

6.2.2.5.3 HII Test\HII Config Access Protocol Test

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

- *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:* It's 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:* It's 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.

6.2.3 UEFI Health Protocol

The Intel® VROC family of products will provide support for the UEFI Health Protocol as defined in the UEFI specifications listed above. Intel® VROC will provide support for `EFI_DRIVER_HEALTH_PROTOCOL.GetHealthStatus()`. The types of health status messages that Intel® VROC will report will match those messages that are displayed in the Intel® VROC 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.

6.2.4 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.

6.2.4.1 Using the UEFI Shell

When the platform BIOS is configured to boot from the 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 the 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"
```

6.2.4.2 Intel® VROC (SATA RAID) Device Information Display in UEFI

The Intel® VROC family of products 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 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 to, the UEFI driver now reports out the physical port value instead of the enumerated value. The data is displayed as a set of three values in the following order: **X-Y-Z**.

- **X** – 0: pass thru 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)

6.2.4.3 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:

- Step 1: Press the key indicated for entering into the setup and select **Boot Options** when the menu appears.
- Step 2: Navigate to **Intel® VROC (SATA RAID) Controller**, and press <Enter>.
- Step 3: Highlight **Create RAID Volume**, and press <Enter>.
- Step 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 unique name, and press <Enter> to save the new name.
- Step 5: Highlight the **RAID Level**. Default setting is RAID 0. To change this to another RAID configuration, press <Enter> and select the desired RAID level, then press <Enter> to save the user's selection.
- Step 6: 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.
- Step 7: Repeat Step 6 with each drive needed until all desired drives have been included within the RAID.
- Step 8: 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.
- Step 9: 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.
- Step 10: To finalize the 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%.

Step 11: The user will see the created volume on the following screen. 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: Step 11 is extremely important, failure to save at this point will cause the user's RAID volume to not be saved. The user must save 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 with the system installation.

6.2.5 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.

6.2.5.1 Pre-OS Driver Installation Using the "Load Driver" Method

Follow these steps to install the Pre-OS driver during the Windows* OS installation:

Step 1: During the operating system installation, after selecting the location to install the Windows* OS click on the **Load Driver** button to install the desired Intel® VROC drivers.

Step 2: When prompted, insert the media with the Intel® VROC driver files and press <Enter>.

Step 3: The user can find the media and browse to the folder where the files are located.

Step 4: Follow the steps to load the driver and return to the installation.

Step 5: Continue the installation.

6.2.5.2 Selectable Boot Volume

The Intel® VROC 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.

6.2.6 Pre-OS RAID Failure Recovery Process

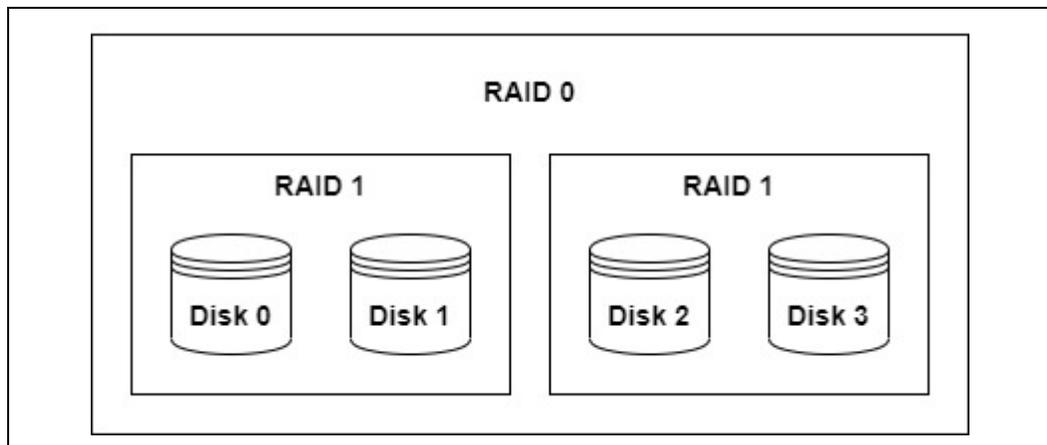
The Intel® VROC Pre-OS UEFI HII provides support for the ability to attempt to recover a failed Intel® VROC RAID volume. When a failed RAID volume is encountered during boot, the option is made available in the Intel® VROC HII page for which the user can select. This is a multi-step process.

Note: Intel makes no guarantee of successful recovery from a failed state using this option. This must be treated as a last chance effort and there is no guarantee that there won't be some data loss. Intel always recommends recovering a failed RAID volume by recreating the RAID volume from scratch and restore the data from the latest platform image backup.

This feature is accessed through the Intel® VROC HII when the platform boots with a failed Intel® VROC RAID volume present. It is important to note that for this to work, there must be enough Intel® VROC RAID member drives present to place the failed RAID volume back into a 'Degraded' state. The process, to attempt to recover the failed RAID volume, begins by booting the platform into the BIOS setup and then navigating into the Intel® VROC HII. By selecting the failed volume, an option will be available to force the volume back into a 'Degraded' (or 'Normal') state. This will expose the next option to specify the Intel® VROC RAID volume member drives to enable this action. Once the necessary drives are selected, the RAID volume will be forced back into a degraded state. From that state the RAID volume can be rebuilt.

The following example shows how a failed Intel® VROC RAID volume is recovered. In this example, the system has four drives in the configuration shown in the chart below. The RAID volume configuration is a four-drive RAID 10.

Figure 6-2. Example of a RAID 10 Configuration

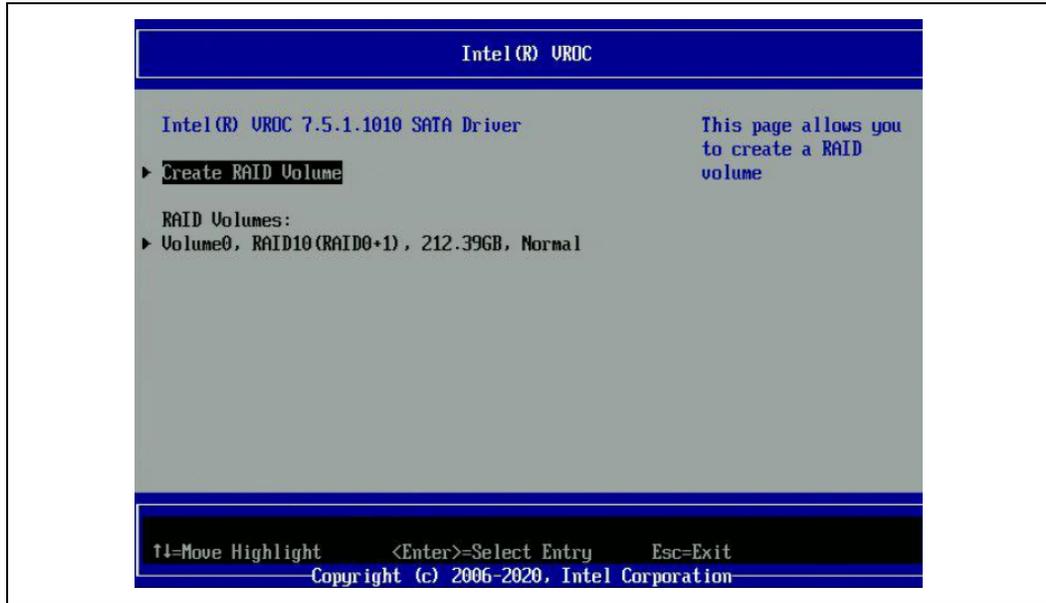


In this example, Disk 0 is removed (putting the RAID volume into a 'Degraded' state) followed by Disk 1 to simulate a platform hardware failure scenario. This platform behavior will cause the RAID 10 volume into a failed state. The example steps through the process of reinserting the missing disks and the steps to recover the failed volume.

Note: This process does not guarantee that all the data can be recovered.

Step 1: In this first screenshot from the Intel® VROC HII, shows the Intel® VROC RAID 10 volume in a 'Normal' state.

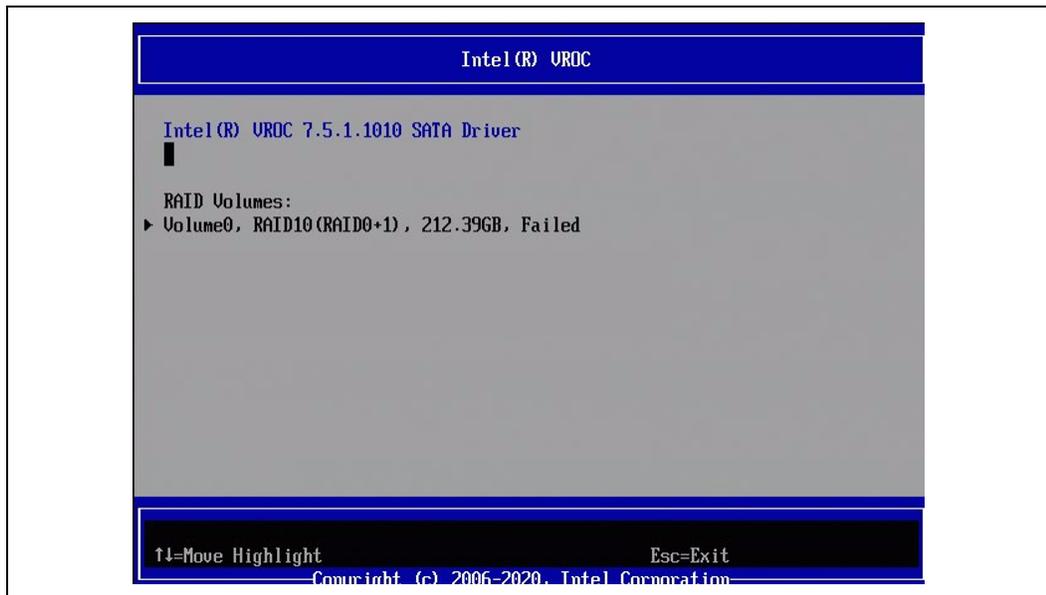
Figure 6-3. Intel® VROC RAID 10 Volume in 'Normal' State in the HII



Step 2: Disk 0 and Disk 1 are removed causing the RAID 10 to go from 'Normal' to 'Degraded' to a 'Failed' state.

Note: In the Intel® VROC HII menu, the 'Volume 0' is listed as 'Failed.'

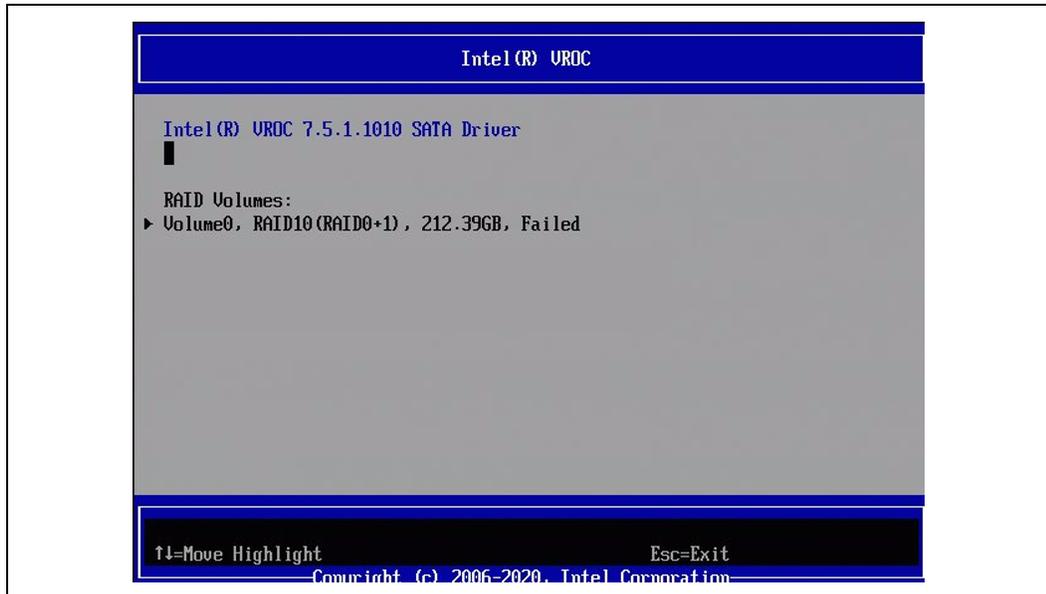
Figure 6-4. Intel® VROC RAID 10 Volume in 'Failed' State in the HII



Step 3: After powering down, the drives are reattached to the system, in the same position that they were before removal. The system is powered back on and

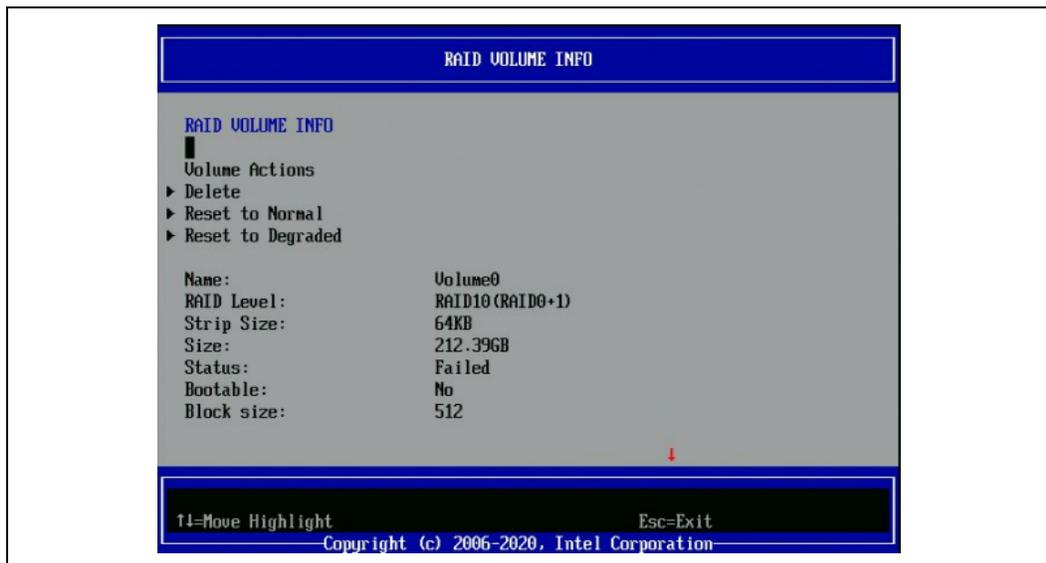
then booted into the Intel® VROC HII, where the RAID 10 is shown to be in a 'Failed' state.

Figure 6-5. Intel® VROC RAID 10 Volume in 'Failed' State after Reboot in the HII



Step 4: To begin the recovery process, select the failed RAID volume. An option to recover will be displayed: *Delete, Reset to Normal, Reset to Degraded*.

Figure 6-6. Intel® VROC RAID Recovery Options in the HII



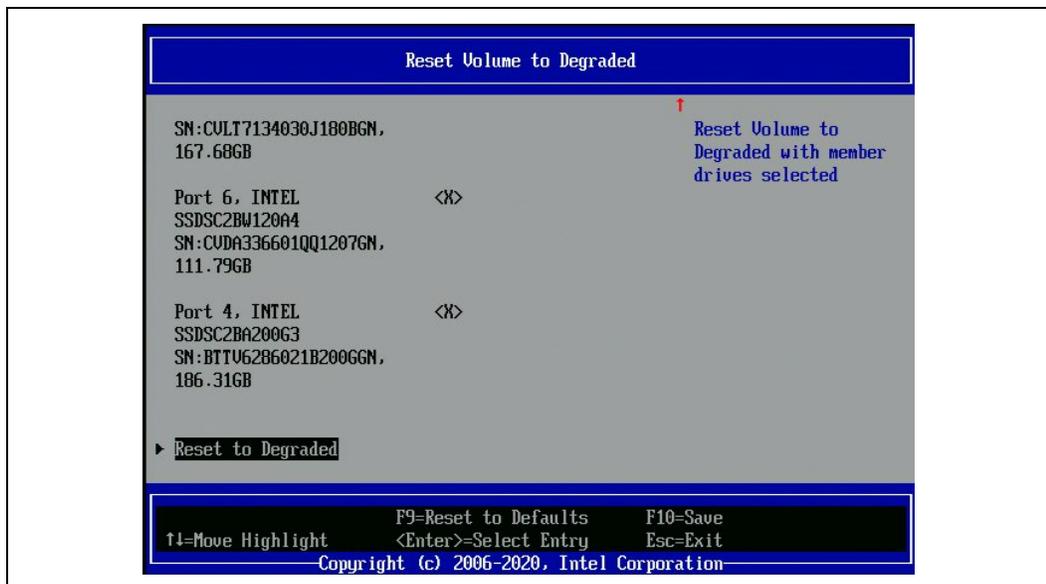
Step 5: Select *Reset to Degraded*. This will open the *Reset to Degraded* menu. In the *Reset to Degraded* menu, select the drives that were not removed from the system (in this case, Disk 2 and Disk 3), and then select the drive that was

the last drive present when the volume was last in a 'Degraded' state. In this example that is Disk 1.

- Note:** Selecting the most recently removed drive can minimize data loss in case of the following event:
- a. Drive 0 is lost from the RAID 10, due to a hardware error, causing the RAID 10 to become 'Degraded.' In the 'Degraded' state you may still write to the RAID; however, it is no longer fully fault tolerant.
 - b. Writing to the 'Degraded' RAID 10 can continue, causing a sync difference between the 'Degraded' RAID 10 and the metadata on Disk 0.
 - c. When Disk 1 is lost (again due to further hardware failures), the RAID 10 will transition to a 'Failed' state and all I/O will stop.
 - d. By forcing the RAID 10 back to a 'Degraded' state, using Disk 1, Disk 2 and Disk 3, theoretically the data on the (now 'Degraded') RAID volume should match that from before Disk 1 was lost.
 - e. This can allow the 'Degraded' RAID 10 to be rebuilt to Disk 0.
 - f. In this event, to restore the RAID 10 to the latest state while minimizing the chance of data loss (with Disk 1 in its original position in the system), it is important to select Disk 1 to sync with Disk 2 and Disk 3.

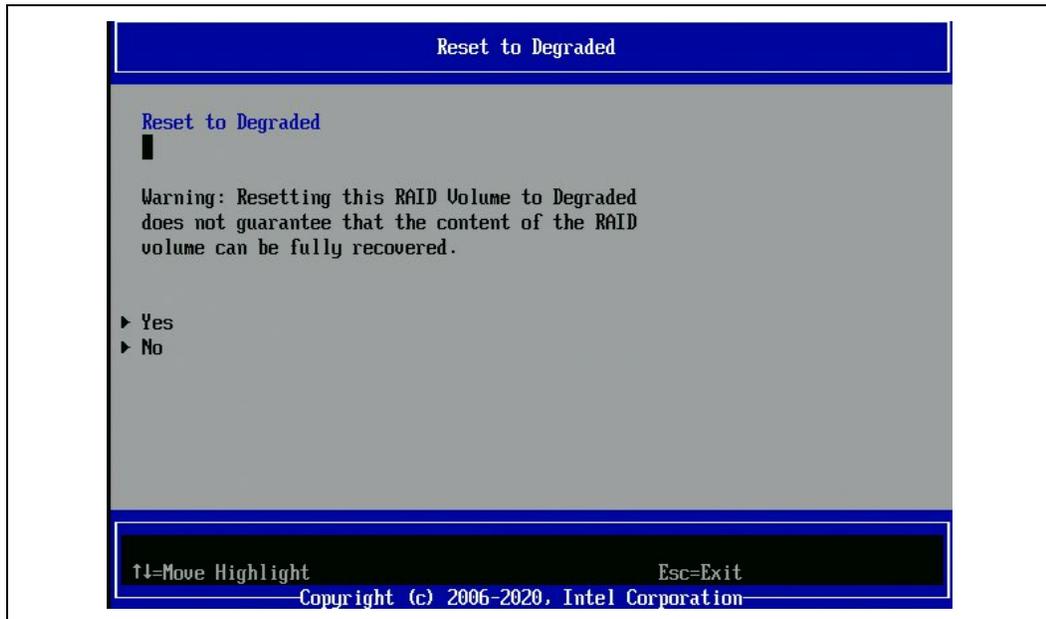
Step 6: Then select the *Reset to Degraded* option at the bottom. This will sync the data between the three drives selected.

Figure 6-7. Intel® VROC 'Reset to Degraded' Option in the HII



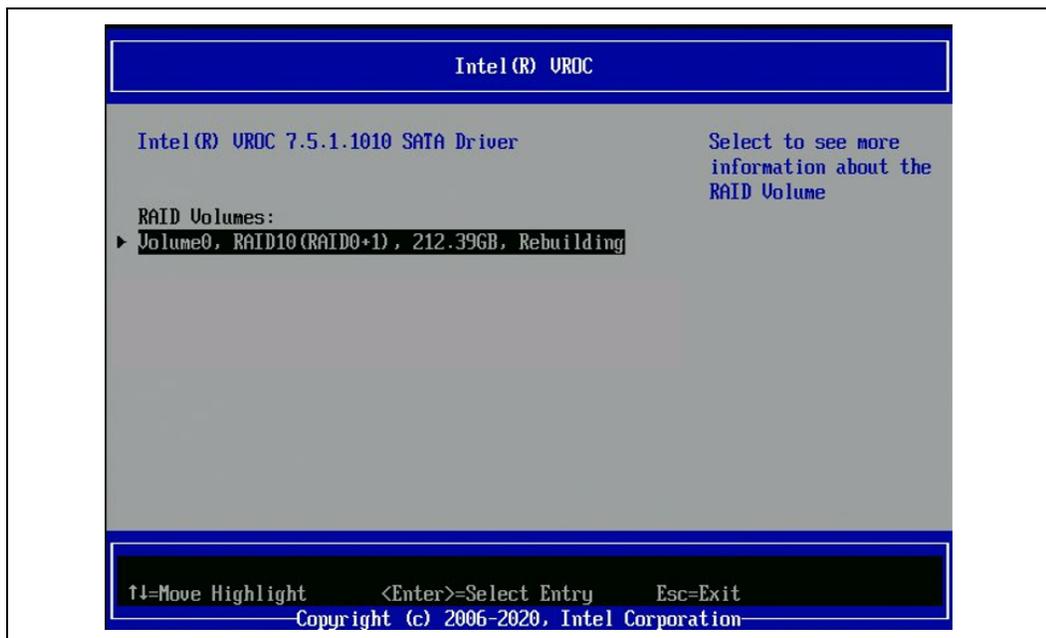
Step 7: On the confirmation screen. Select <Yes>.

Figure 6-8. Intel® VROC Recovery Confirmation Option in the HII



Step 8: The RAID 10 can now be rebuilt. In this example, Disk 0 is also present and in its original location. In the 'Degraded' state, Intel® VROC will see that all four drives from the RAID 10 are present and healthy, hence the RAID volume will begin the rebuilding process automatically.

Figure 6-9. Intel® VROC RAID 10 Volume in Rebuilding Process in the HII



6.2.7 UEFI Firmware Management Protocol Support

Intel® VROC 7.7 provides limited support for the *UEFI Firmware Management Protocol (FMP)* as outlined in UEFI Specification version 2.9. The Intel® VROC UEFI drivers will provide support for updating drive firmware (through UEFI FMP) when the drive is managed by Intel® VROC. This includes both Intel® VROC (VMD NVMe RAID) and Intel® VROC (SATA RAID). The Intel® VROC UEFI driver has implemented a subset of protocol functions defined in the UEFI Specification, to support this feature, which is reflected in the below table.

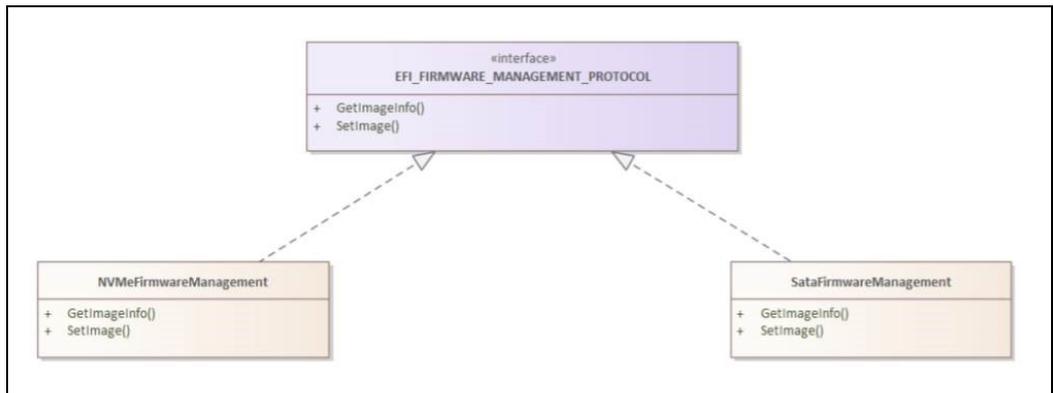
Table 6-10. Firmware Management Protocol Support

EFI Firmware Management Protocol Functions	Intel® VROC UEFI implementation
GetImageInfo	Supported
SetImage	Supported

Intel® VROC will return an `EFI_UNSUPPORTED` status for UEFI FMP commands not listed above.

The following diagram highlights the process flow to support both NVMe and SATA drives.

Figure 6-10. Firmware Management Protocol Support Flow



6.2.7.1 GetImageInfo for NVMe SSDs

The Intel® VROC (VMD NVMe RAID) UEFI driver will respond to a `GetImageInfo` command with `EFI_FIRMWARE_IMAGE_DESCRIPTOR(s)` filled in the following way:

- `ImageIndex` will indicate the Firmware Slot Index (1...N), defined in the NVMe specification.
- `ImageTypeId` will be `25D4CA70-05E8-44CB-B13D-ABA064717432`.
- `ImageId` will be the `ImageIndex`.
- `ImageIdName` string will be `Drive model: <ModelNumber>, Firmware Slot: <FirmwareSlotIndex>`.
- `Version` will be `0`.

- **VersionName** will be the **Firmware Revision for Slot** [**<ImageId>**], defined in the NVMe specification.
- **Size** will be 1.
- **AttributesSupported**, the following attributes will be set to:
 - **IMAGE_ATTRIBUTE_IN_USE**
 - **IMAGE_ATTRIBUTE_IMAGE_UPDATABLE**
 - Rest of attributes shall be cleared.
- **AttributesSetting** will be set according to the table below:

Table 6-11. Values for AttributeSetting for GetImageInfo for NVMe SSDs

Attribute	Value
IMAGE_ATTRIBUTE_IMAGE_UPDATABLE	If this is the first firmware slot and FRMW bit 0 is set, then attribute cleared to 0, otherwise attribute set to 1.
IMAGE_ATTRIBUTE_RESET_REQUIRED	Cleared
IMAGE_ATTRIBUTE_AUTHENTICATION_REQUIRED	Cleared
IMAGE_ATTRIBUTE_IN_USE	If this firmware slot is active (pointed by AFI 2:0 in NVMe specification), then set, otherwise cleared.
IMAGE_ATTRIBUTE_UEFI_IMAGE	Cleared
IMAGE_ATTRIBUTE_DEPENDENCY	Cleared

- **Compatibilities** will be set to 0.
- **LowestSupportedImageVersion** will be set to 3.
- **LastAttemptVersion** will be set to 0.
- **LastAttemptStatus** will be set to the last returned status by the **SetImage** function if called, otherwise it will be set to **LAST_ATTEMPT_STATUS_SUCCESS**.
- **HardwareInstance** will be constructed in the following way:

Table 6-12. HardwareInstance Values

Bits	Value (Physical Location of NVMe Drive)
63:32	0
31:24	CPU Socket Number
23:16	VMD Controller Number
15:8	Root Port Offset
7:0	PCIe Slot Number

Intel® VROC (VMD NVMe RAID) UEFI driver will return as many **EFI_FIRMWARE_IMAGE_DESCRIPTOR**s as the number of NVMe firmware slots with non-zero firmware revisions.



Intel® VROC (VMD NVMe RAID) UEFI driver will return **PackageVersion** equal to **0xFFFFFFFF** and **PackageVersionName** equal to **NULL** pointer.

6.2.7.2 Set Image for NVMe SSDs

When Intel® VROC (VMD NVMe RAID) UEFI driver receives a **SetImage** command, it will update the NVMe SSDs with the image provided.

6.2.7.2.1 Firmware Activation

When Intel® VROC (VMD NVMe RAID) receives a **SetImage** function for an NVMe SSD that does not support firmware activation without reset, Intel® VROC will send the firmware image to the NVMe SSD; activate the firmware and then reset the NVMe SSD (*Controller Level Reset*).

When Intel® VROC (VMD NVMe RAID) receives a **SetImage** function for an NVMe SSD that does support firmware activation without reset, Intel® VROC will send the firmware image to the NVMe drive and activate the new firmware immediately.

6.2.7.2.2 Firmware Update Status

Intel® VROC (VMD NVMe RAID) will report the firmware progress update in multiple stages:

- Progress reported proportionally to the number of image chunks already sent to the drive up to 90%.
- 100% progress reported after firmware activation.

6.2.7.3 GetImageInfo for SATA SSDs

Intel® VROC (SATA RAID) UEFI driver will respond to a **GetImageInfo** command. It will respond with **EFI_FIRMWARE_IMAGE_DESCRIPTOR(s)** filled in the following way:

- **ImageIndex** will be set to **1**.
- **ImageTypeId** will be **9C9A55B0-A5CD-456F-9AC3-76765D68733D**.
- **ImageId** will be the **ImageIndex**.
- **ImageIdName** string will be **Drive model: <ModelNumber>**.
- **Version** will be **0**.
- **VersionName** will be the firmware revision in **IDENTIFY DEVICE DATA**, defined in ATA specification.
- **Size** will be **0**.
- **AttributesSupported**, the following attributes will be set to:
 - **IMAGE_ATTRIBUTE_IN_USE**
 - **IMAGE_ATTRIBUTE_IMAGE_UPDATABLE**
 - Rest of attributes shall be cleared.
- **AttributesSetting** will be set according to the table below:

Table 6-13. Values for AttributeSetting for GetImageInfo for SATA SSDs

Attribute	Value
IMAGE_ATTRIBUTE_IMAGE_UPDATABLE	If DOWNLOAD MICROCODE or DOWNLOAD MICROCODE DMA is supported, then attribute set to 1, otherwise attribute cleared to 0.
IMAGE_ATTRIBUTE_RESET_REQUIRED	Cleared
IMAGE_ATTRIBUTE_AUTHENTICATION_REQUIRED	Cleared
IMAGE_ATTRIBUTE_IN_USE	Set
IMAGE_ATTRIBUTE_UEFI_IMAGE	Cleared
IMAGE_ATTRIBUTE_DEPENDENCY	Cleared

- **Compatibilities** will be set to 0.
- **LowestSupportedImageVersion** will be set to 3.
- **LastAttemptVersion** will be set to 0.
- **LastAttemptStatus** will be set to the last returned status by the **SetImage** function if has been called, otherwise it shall be equal to **LAST_ATTEMPT_STATUS_SUCCESS**.
- **HardwareInstance** will be set to the AHCI port number to which the SATA drive is connected.

The Intel® VROC (SATA RAID) UEFI driver will return one **EFI_FIRMWARE_IMAGE_DESCRIPTOR**.

Intel® VROC (SATA RAID) UEFI driver will return **PackageVersion** equal to 0xFFFFFFFF and **PackageVersionName** equal to NULL pointer.

6.2.7.4 SetImage for SATA SSDs

When Intel® VROC (SATA RAID) receives a **SetImage** command, Intel® VROC will update the SATA drive with the image provided. Intel® VROC will select **ATA DOWNLOAD MICROCODE** subcommand based on the drive capabilities, the **VendorCode** argument and the **_ATA_DM_SUBCOMMAND_SELECTION_POLICY** to update the firmware.

If the **SetImage** command contains invalid parameters, Intel® VROC will return the **EFI_INVALID_PARAMETER** status.

6.2.7.5 Firmware Update Progress

Intel® VROC (SATA RAID) will report the firmware progress update in multiple stages, reported proportionally to the number of image chunks already sent to the drive, up to 100%, when the **ATA DOWNLOAD MICROCODE** subcommand is 03h.

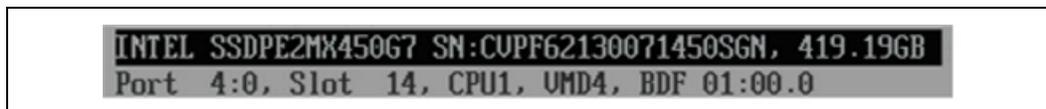
When the **ATA DOWNLOAD MICROCODE** subcommands are 0Eh and 0Fh, Intel® VROC (SATA RAID) will report the firmware update progress in multiple stages:

- Progress reported proportionally to the number of image chunks already sent to the drive up to 90%.
- 100% progress reported after sending subcommand 0Fh.

6.2.8 Physical Drive Location Information in UEFI

When the Intel® VROC product displays how Intel® VMD controller enumerates the drives, the drive information is showing in VROC HII like in the below figure. The information is not straightforward and not easy to interpret. To present the drive location information in a user-friendly manner, Intel® VROC offers a protocol that will allow customers to report how the drives are connected and how customers would like the information reported in the customer platform.

Figure 6-11. Physical Drive Location Information in UEFI



6.2.9 Multiple Namespaces in NVMe SSDs

Intel® VROC (VMD NVMe RAID) UEFI driver will process NVMe SSDs that support multiple namespaces, by reporting the first namespace as the one that Intel® VROC (VMD NVMe RAID) will recognize.

6.2.10 Disable Intel® VROC LED Management

The Intel® VROC provides support for the OEM/ODMs to disable the Locate LED functionality within the Intel® VROC LED Management. By enabling the feature, the OEM/ODM can use their own tools to initiate a Locate LED functionality within their platform. The mechanism that is used to accomplish this feature through the EFI variable is called *VrocUefiConfig*. When this feature is enabled, the Locate LED option will be hidden in the Intel® VROC UEFI HII menus.

6.2.10.1 VrocUefiConfig Spec to Disable Intel® VROC LED Management

The description of the *VrocUefiConfig* interface to be able to disable Intel® VROC LED Management is as follows:

```
#define VROC_UEFI_CONFIG_STR L"VrocUefiConfig"
#define VROC_UEFI_CONFIG_SATA_HII_DISABLE_LOCATE_LED 0x1
#define VROC_UEFI_CONFIG_VMD_HII_DISABLE_LOCATE_LED 0x2
typedef UINT64 VrocUefiConfig;
```

To disable Locate LED from HII, customers must set this variable via **SetVariable()** (UEFI_Spec_2_7.pdf, section 8.2):

```
typedef
EFI_STATUS
SetVariable (
    IN CHAR16 *VariableName,
    IN EFI_GUID *VendorGuid,
    IN UINT32 Attributes,
    IN UINTN DataSize,
    IN VOID *Data
);
```

SetVariable() example call which disables Locate LED from both Intel® VROC SATA and Intel® VROC VMD HII:

```
EFI_GUID IntelVrocGuid = { 0x193dfefa, 0xa445, 0x4302, { 0x99, 0xd8, 0xef, 0x3a, 0xad,
0x1a, 0x04, 0xc6 } };
VrocUefiConfig = VROC_UEFI_CONFIG_SATA_HII_DISABLE_LOCATE_LED |
VROC_UEFI_CONFIG_VMD_HII_DISABLE_LOCATE_LED;
SetVariable(
    VROC_UEFI_CONFIG_STR,
    &IntelVrocGuid,
    EFI_VARIABLE_BOOTSERVICE_ACCESS,
    sizeof(vrocUefiConfig),
    &vrocUefiConfig
);
```

6.2.10.1.1 Calling SetVariable()

The Intel® VROC UEFI driver HII will read the *VrocUefiConfig* EFI variable only once, when entering the *Disk Info Form* for the first time. The Disk Info Form is the form where Locate LED can be located. That means that **SetVariable()** does not have to be called before loading the Intel® VROC UEFI drivers, because all drivers will be loaded and initialized prior to the user entering the Intel® VROC HII. For this, there are no additional depex files needed and **SetVariable()** may be called any time during BIOS initialization.

6.2.10.1.2 VrocUefiConfig Not Set or Set to 0

When the *VrocUefiConfig* is not set or set to a value of 0, the Intel® VROC UEFI drivers work as usual and Intel® VROC will manage the Locate LED process.

6.2.10.1.3 Intel® VROC Tools

The Intel® VROC RAID Config tools (RCfg) will report an error message when *VrocUefiConfig* states that Intel® VROC LED Management is not to be supported.

6.3 Intel® VROC Monitoring and Management

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

6.3.1 Intel® VROC GUI

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

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

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



6.3.1.1 Intel® VROC Installation

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

When the Intel® VROC 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 installer will cancel the 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 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 (SATA RAID) driver for SATA
- Intel® VROC (SATA RAID) driver for sSATA
- Intel® VROC (SATA RAID) driver for tSATA
- Intel® VROC System Service
- Intel® VROC Notification Area Icon

6.3.1.1.1 Default Installation Path

The Intel® VROC installer software will install the Intel® VROC solution to a default folder unless specified otherwise.

The default path shall be C:\Program Files (x64)\Intel\Intel(R) Virtual RAID on CPU.

6.3.1.2 Intel® VROC (SATA RAID) Driver

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

6.3.1.3 Intel® VROC (VMD NVMe RAID) Driver

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

6.3.2 Getting Started

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

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

6.3.2.1 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.

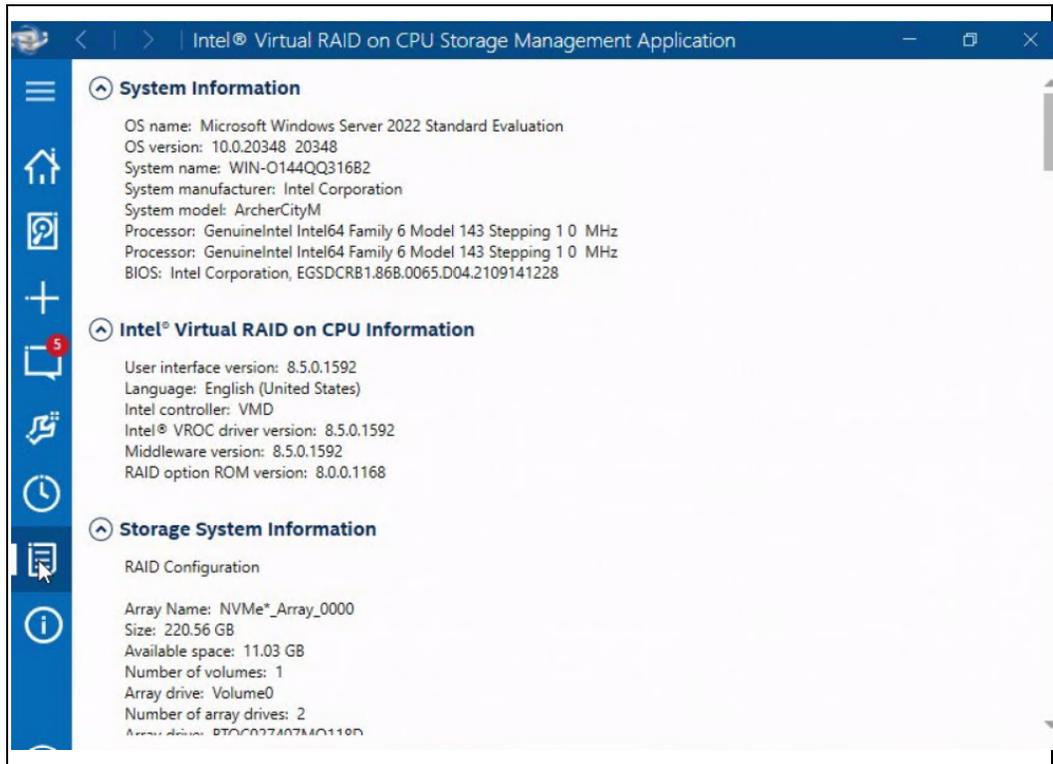
6.3.3 Component Version Detection

At launch, the Intel® VROC GUI 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 will be used to determine the features visible in the GUI.

6.3.4 Using System Report

The System Report from the Intel® VROC GUI will provide the detailed information about the system. It contains platform information, Intel® 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.

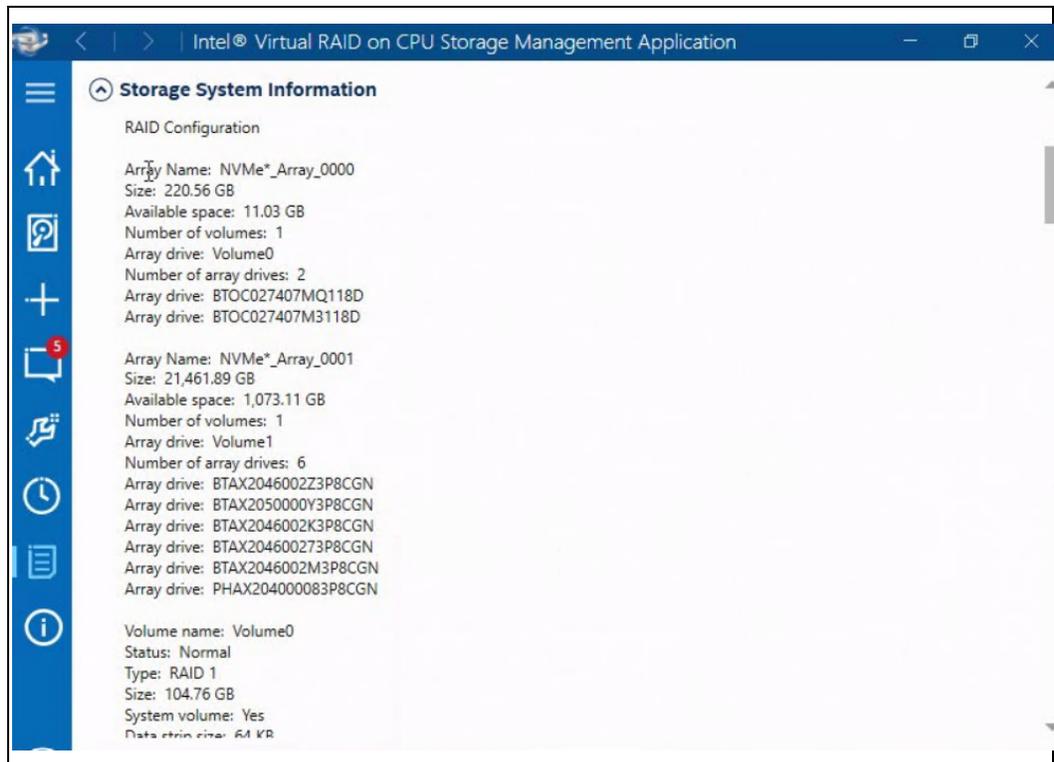
Figure 6-12. Intel® VROC GUI System Report Window



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 the **Save** button and follow the instructions.

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

Figure 6-13. Intel® VROC Components from the Intel® VROC GUI System Report



6.3.5 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 the Intel® VROC 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 package is installed. Both administrators and standard users can change the notification area settings using the application or directly from the notification area. Setting changes are applied on a per user basis, and do not affect other users’ settings.

Note: The Intel® VROC user interface icon will not automatically display after the product is initially installed. After the Intel® VROC user interface is opened for the first-time following the installation, the Intel® VROC notification tray icon will display. As a result, there is no notifications of potential issues encountered with Intel® VROC managed devices until after the interface is opened following the installation. After the Intel® VROC user interface is opened for the first time, the system will then be configured to automatically display any notifications each time the interface is opened.

6.3.5.1 Opening the Application from the Notification Area

- To open the Intel® VROC application from the notification area, follow these steps:

Step 1: Right-click the icon.

Step 2: Click **Open Application**.

The notification area icon can be in the following states described in the below table.

Table 6-14. Notification Area Icon 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.

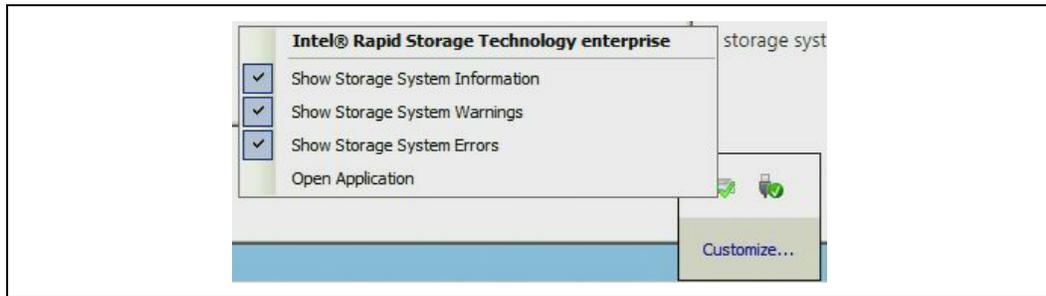
6.3.5.1.1 Selecting System Notifications

To select system notifications, follow these steps:

Step 1: Right-click on the Intel® VROC notification tray icon.

Step 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.

Figure 6-14. Selecting System Notifications



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

6.3.5.1.2 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.

6.3.5.2 Understanding the Application

The Intel® VROC 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.

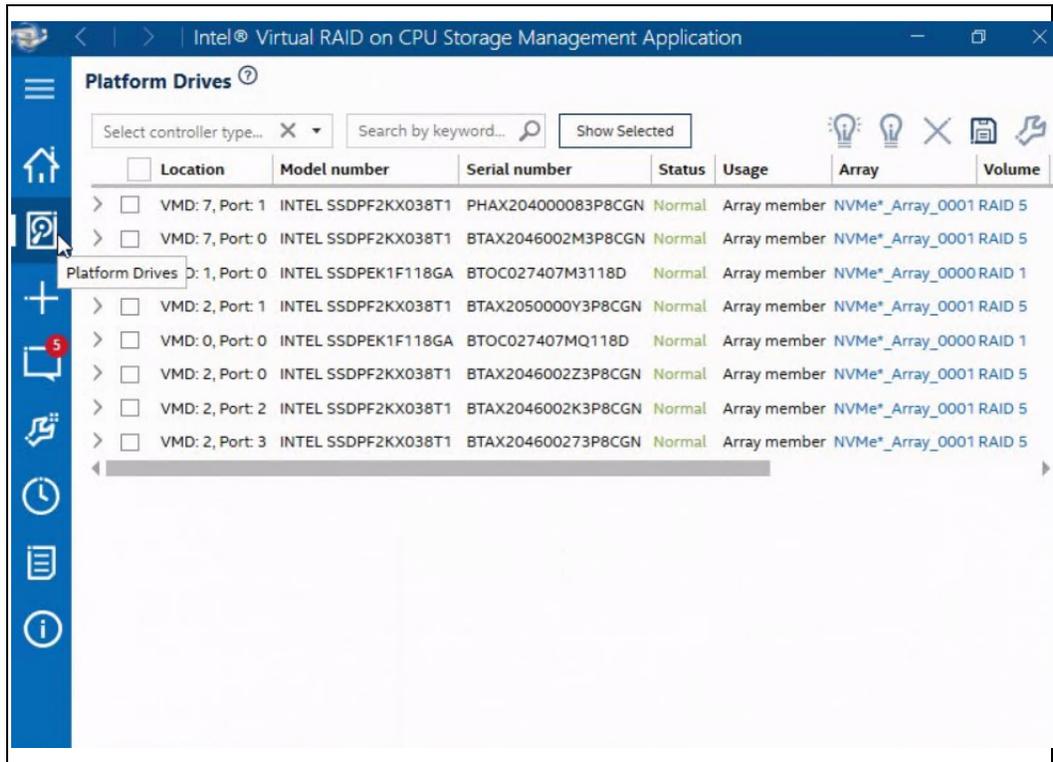
6.3.5.2.1 Storage System Configuration

The Intel® VROC storage system combines hardware capabilities with Intel® VROC 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 listed in the controller information are 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 shows the properties information of what may be selected (i.e., Volume_0000).

Figure 6-15. Intel® VROC GUI Storage System Window



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:

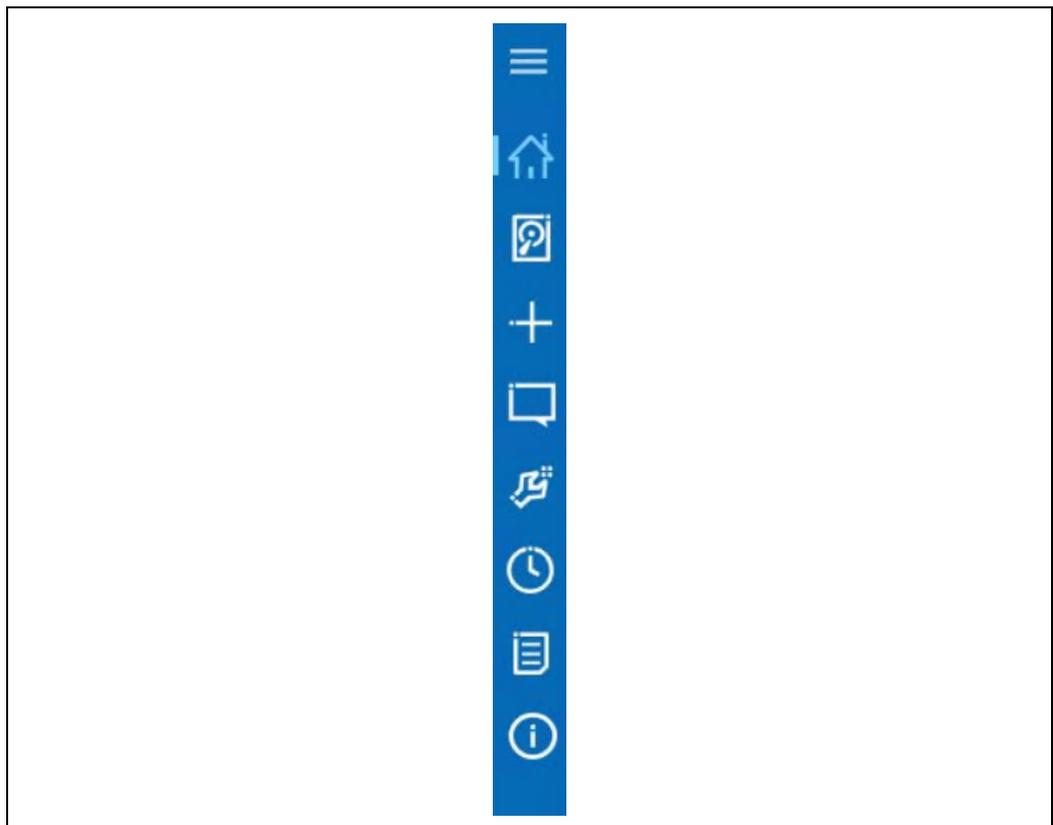
- There is a combination of SSDs and HDDs used.
- There are at least two drives that have a size difference of more than 10%.

6.3.5.2.2 Navigation

The Intel® VROC 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.

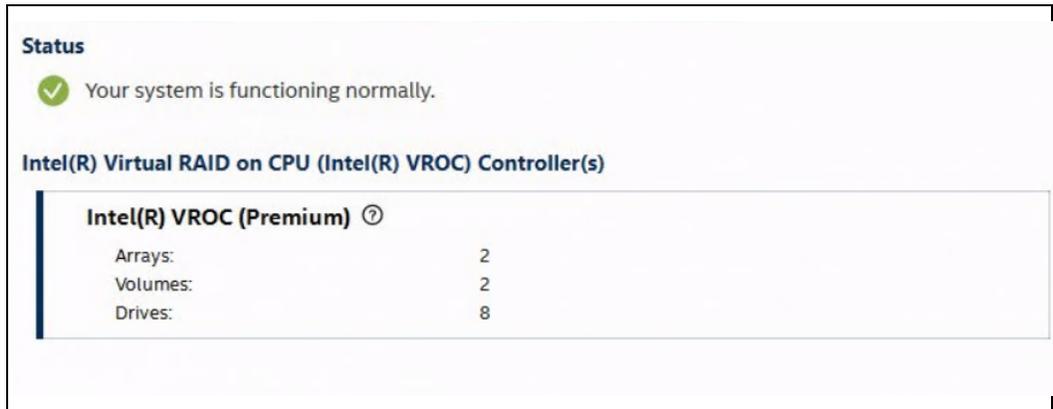
Figure 6-16. Navigation Options in the Intel® VROC GUI



Home

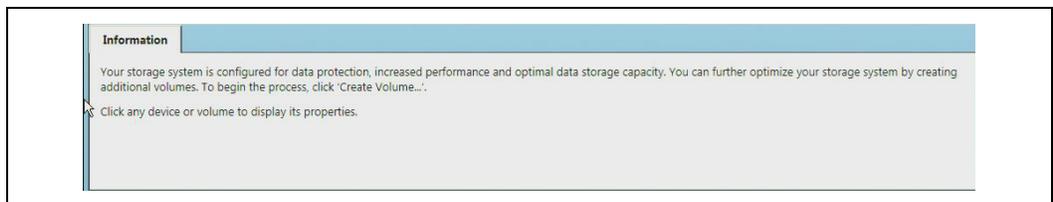
Home is the main area of operation, where the user can select to **Scan** for hardware changes or **Create Volume**.

Figure 6-17. Available Options in the Home Screen of the Intel® VROC GUI



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

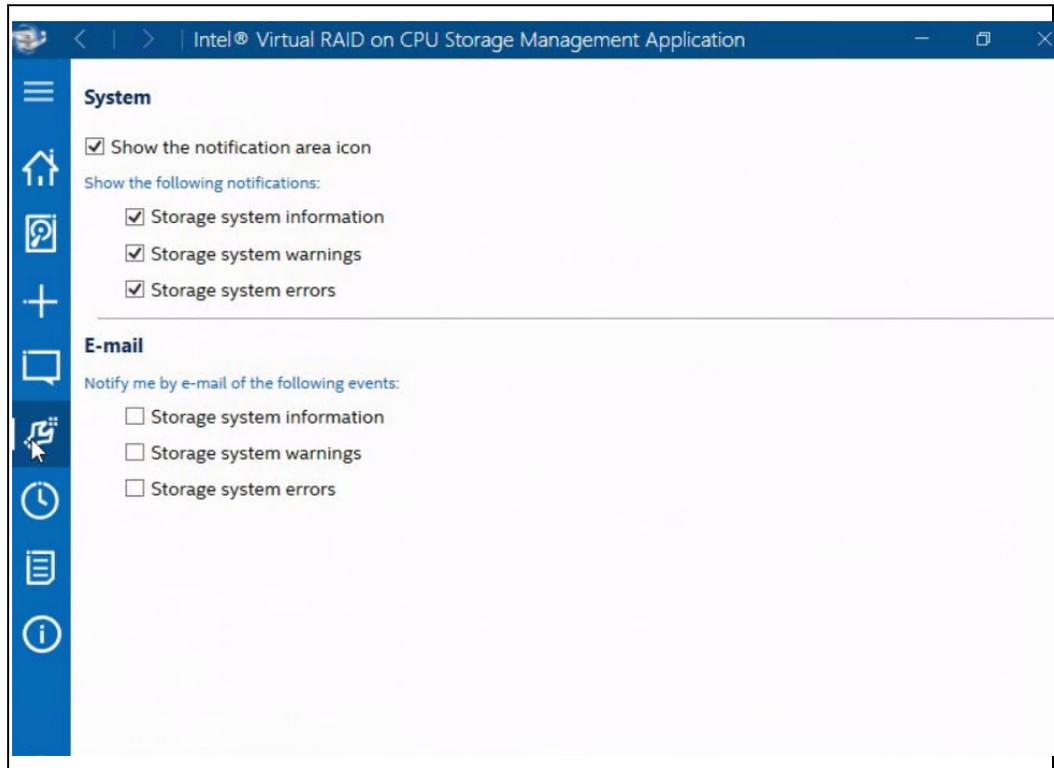
Figure 6-18. Information Section in the Home Screen of the Intel® VROC GUI



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.

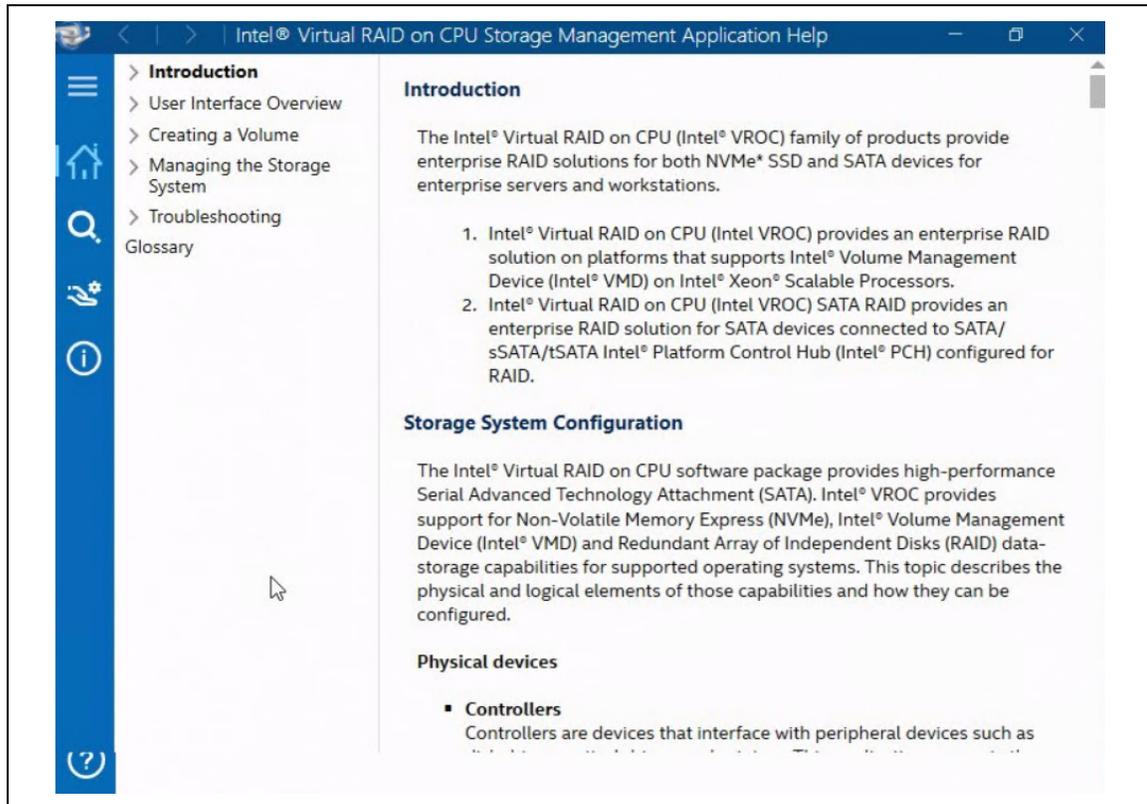
Figure 6-19. Available Options in the Preferences Screen of the Intel® VROC GUI



Help:

The *Help* area provides detailed information on using the Intel® VROC GUI and key system information.

Figure 6-20. Available Options in the Help Screen of the Intel® VROC GUI



6.3.5.3 Storage System Status

Anytime the Intel® VROC 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 to enhance or repair the user’s storage system.

6.3.5.3.1 Understanding the Status Area

To get the full benefits of what Intel® VROC GUI has to offer, it is critical to maintain a healthy storage system. The application helps the user track and report 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 states described in the table below.

Table 6-15. Possible States for the Storage System in the Intel® VROC GUI

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 <i>Create</i> 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 <i>Manage</i> 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 <i>Manage</i> subsection displays any disk or volume states reported by the storage system that may require the user's attention 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 <i>Manage</i> subsection displays any disk or volume states reported by the storage system that require the user's immediate attention 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.

6.3.5.4 Volume States

The below table shows the possible states a RAID volume may be in.

Table 6-16. Possible States for RAID Volumes in the Intel® VROC GUI

Volume Type	Normal	Degraded	Failed
RAID 0		Not applicable	

Volume Type	Normal	Degraded	Failed
Single-Disk (Cache)			
Single-Disk (Data)		Not applicable	
RAID 1			
RAID 5			
RAID 10			

Note: Refer to the [Troubleshooting](#) section for more information to recover RAID volumes in *Degraded* or *Failed* state.

6.3.5.5 RAID-Ready Mode

A RAID-Ready system refers to a system that has been configured to support the Intel® VROC family of products. The system BIOS has the appropriate pre-boot drivers and has been configured for RAID mode. Intel® VROC 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.

6.3.5.6 RAID-Ready Migration

- Pass-through to 2-XX drive RAID 0 (Intel® VROC (VMD NVMe RAID)).
- Pass-through to 2-8 drive RAID 0 (Intel® VROC (SATA RAID)).
- Pass-through to a 2 drive RAID 1.
- Pass-through to a 3 to XX drive RAID 5 (Intel® VROC (VMD NVMe RAID)).
- Pass-through to a 3 to 8 drive RAID 5 (Intel® VROC (SATA RAID)).
- Pass-through to a 4 drive RAID 10.

6.3.5.7 Intel® VROC Graphical User Interface (GUI)

The Intel® VROC 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.

6.3.5.7.1 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.

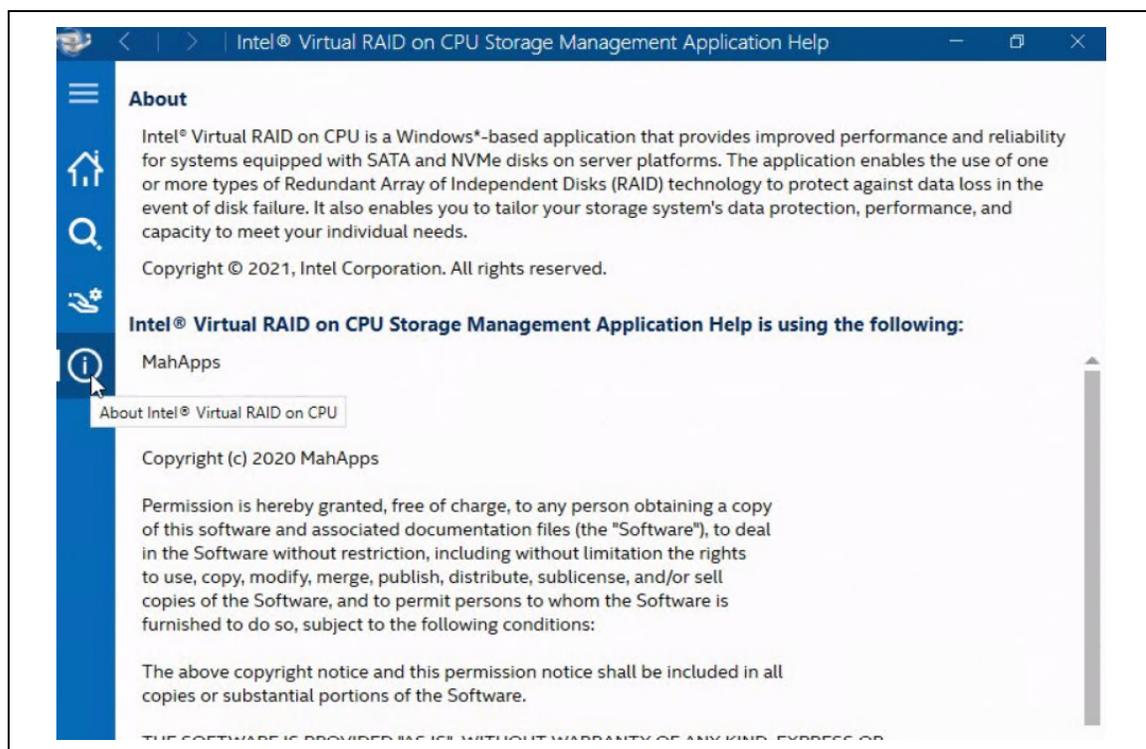
6.3.5.7.2 Intel® VROC Enabled Mode GUI Support

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.

6.3.5.8 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**.

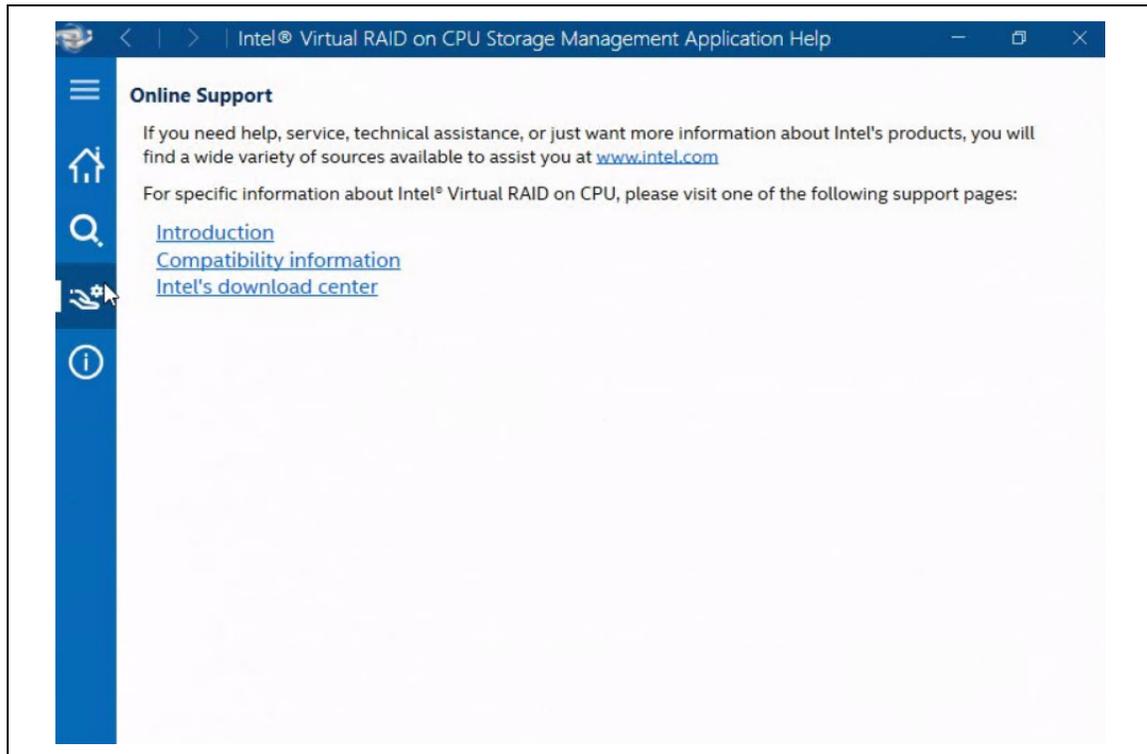
Figure 6-21. Intel® VROC GUI Version Information



6.3.5.9 Intel® VROC GUI Online Support

In the Intel® VROC GUI **Help** menu, the submenu **Online Support** when selected will display a pop-up window with the support URLs as shown in the figure below:

Figure 6-22. Online Support Options in the Intel® VROC GUI



6.3.6 Managing Arrays

The user must be logged on as an 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.

6.3.6.1 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.

Table 6-17. Array Properties

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.

6.3.6.2 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:
 - It must be on the same controller as the existing array (SATA, sSATA, tSATA or Intel® VMD). It cannot be used to cross controllers (SATA to sSATA/tSATA or SATA to Intel® VMD).
 - 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.

6.3.6.3 Package Upgrade or Downgrade While Migrating

When performing a migration, do not initiate an Intel® VROC upgrade or downgrade until the migration completes. If a driver upgrade or downgrade is performed while a migration is in progress, the sub sequential system restart could result in the platform not recognizing the volume or the data on it.

6.3.6.4 Migration of a System Disk to a RAID Volume

With the Intel® VROC release package, performing a migration from a system boot disk to a RAID volume requires a system reboot as the final step in completing the migration. Under this condition the migration process will start, but the progress status will remain at 0% and data copying will not begin until the platform has been rebooted. This is to ensure that once the migration begins, other power management state changes will not negatively impact the operation of the system.

For Intel® VROC versions prior to 6.2, we recommended rebooting the system after a system disk to RAID volume migration is initiated. Otherwise, the migration must be allowed to complete before the system can successfully perform power management



state changes. If a power state change occurs prior to a system reboot or before the migration completes, a system failure may be encountered.

6.3.6.5 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.
- Intel® NVMe SSDs attached to a CPU (that does not contain the Intel® VMD).
- SATA drives attached to the SATA/sSATA/tSATA 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.

6.3.6.5.1 Volume Requirements

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

- The user is logged on as an administrator.
- The computer is RAID ready.
- Two or more supported disks, including the operating system disk are connected, and in a normal state, and unlocked (in the case of a password-protected disks).
- Each disk that is 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 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.

- If the first selection is a system disk, any additional disks selected must be of equal or greater size so that all the system files can be migrated to the new volume.
- 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.
- 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.

- The disks used to create a volume must have the same type of connection (NVMe/SATA).

6.3.6.5.2 Creation Process

Now that we have reviewed the volume requirements, this section will guide 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 below (*Completing the Process*) for details.

Creating a Custom Volume

To create a custom volume, follow the below steps:

Step 1: Click **Create**.

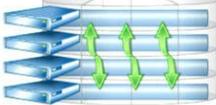
Step 2: Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.

Step 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%.

Figure 6-23. RAID Volume Types Description in the Intel® VROC GUI

 <p>RAID 1 : Real-time data protection</p> <p>Combines two disks where each disk stores an exact copy of the data to appear as a mirror of each other.</p> <p>Disks required: 2</p> <p>Advantage: Full data redundancy and excellent fault-tolerance; increased read transfer rate.</p> <p>Disadvantage: Storage capacity is only as large as the smallest disk; slight decrease in write transfer rate.</p> <p>Application: Typically used in workstations and servers to store critical data. Available in specific mobile configurations.</p>	 <p>RAID 0 : Optimized disk performance</p> <p>Combines two or more disks and breaks down data into units that are spread across the array disks.</p> <p>Disks required:</p> <ul style="list-style-type: none"> SATA controllers: 2 to maximum number of ports in a single SATA Controller Intel VMD controllers: 2 to 24 <p>Advantage: Increased data access and storage performance; no loss in data capacity.</p> <p>Disadvantage: No data redundancy (if one disk fails, all data on the volume is lost).</p> <p>Application: Typically used in desktops and workstations to store high-performance, temporary data and software.</p>
 <p>RAID 5 : Efficient data hosting and protection</p> <p>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.</p> <p>Disks required:</p> <ul style="list-style-type: none"> SATA controllers: 3 to maximum number of ports in a single SATA Controller Intel VMD controllers: 3 to 24 <p>Advantage: Data redundancy; improved storage performance and capacity; high fault-tolerance and read performance.</p> <p>Disadvantage: Time-consuming to rebuild and decreased performance during the process.</p> <p>Application: Good choice for large amounts of critical data, such as file and application servers and Internet and Intranet servers.</p>	 <p>RAID 10 Balanced performance and data protection</p> <p>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.</p> <p>Disks required: 4</p> <p>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.</p> <p>Disadvantage: 4 disks are required, resulting in increased cost.</p> <p>Application: High-performance applications and high-load database servers requiring data protection, such as video editing.</p>

Configuring the Volume

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

RAID Volume

Step 1: Type a new volume name if the default name change is desired.

Step 2: Select the required number of disks.

Step 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.

Step 4: Click **Next**. This button will not be active until all the required selections have been made.

Advanced configuration settings:

Step 1: Select the array allocation by using the slider.

Step 2: Select a data strip size.

Step 3: Enable or disable the volume write-back cache.

Step 4: Select the check box to initialize the volume. The user can choose to perform this action later.

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 the 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 has configured the volume with the disk selection and other settings, the user will review the projected configuration and, if satisfied, will 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. Intel strongly recommends 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.

Step 1: Under **Confirm**, review the selected configuration.

Step 2: Click **Create Volume** if the user wants to create the volume using the selected configuration. The process starts immediately.

Step 3: Click **OK** to confirm.

Step 4: Once completed, a dialog displays to notify the user that the volume was successfully created. 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:

Step 5: Once the migration status reports 100% complete, restart the user's computer for the operating system to recognize the new volume size.

Step 6: Create a new partition or extend the existing partition to utilize the new volume space using Windows* Disk Management.

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



6.3.6.5.3 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.

Step 1: Click **Create** or **Create a custom volume** under *Status*.

Step 2: Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.

Step 3: Click **Next**.

Step 4: Select **Yes** to add the volume to an existing array.

Step 5: Make any necessary changes in the *Advanced* section.

Step 6: Click **Next**.

Step 7: Review the selected configuration. Click **Back** or an option in the left pane if the user wants to make changes.

Step 8: Click **Finish** to start the creation process.

Table 6-18. 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

Combine	With
	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 if the volume requirements are met.

Step 1: Click **Create** or **Create a custom volume** under *Status*.

Step 2: Select the volume type. Selecting a volume type in the list updates the graphical representation to provide a detailed description of that type.

Step 3: Click **Next**.

Step 4: Select **No** to add a volume to a new array.

Step 5: Select the required number of disks.

Step 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.

Step 7: Make any necessary changes in the *Advanced* section.

Step 8: Review the selected configuration. Click **Back** or an option in the left pane if the user wants to make changes.

Step 9: Click **Next**.

Step 10: Click **Finish** to start the creation process.

RAID Volume Creation with Data Preservation

Intel® VROC 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 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)).
- 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)).

6.3.6.6 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 32MB 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 the available space on that array. If two volumes are present on a single array and capacity expansion is possible, only the space available at the end of the second volume will be used to increase the volume size.

This option is only available if:

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

Increasing the Volume Size from Manage Array

Step 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.

Step 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.

Step 3: Click Yes to confirm. Caution: Once the data migration starts, the operation cannot be canceled.

Step 4: Once the migration has completed, restart the user's computer for changes to take effect and to verify that the product operates properly. Then use Windows* Disk Management to increase the partition size on the volumes or add another partition.

Note: Until the system is restarted, continuing to use the products management tools may result in unknown errors being encountered.

Increasing the Volume Size from Manage Volume

Step 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.

Step 2: Click Increase size.

Step 3: Click Yes to confirm. Caution: Once the data migration starts, the operation cannot be canceled.

Step 4: Once the migration has completed, restart the user's computer for changes to take effect and to verify that the product operates properly. Next, use Windows* Disk Management to increase the partition size on the volumes, or to add another partition.

Note: Until the system is restarted, continuing to use the products management tools may result in unknown errors being encountered.

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

6.3.6.7 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).

Step 1: Under Home, in the storage system view, click the array the user wants to manage. The element properties should now be displayed on the right.

Step 2: In the Advanced section, click Enable or Disable depending on the option available.

Step 3: Click **Yes** to confirm. The page refreshes and displays the new setting.



6.3.6.8 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.

6.3.6.9 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.

Table 6-19. RAID Volume Status as shown in the Intel® VROC GUI and HII

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.
Degraded	Indicates that one array disk is missing or has failed. A RAID 0 volume cannot be in this state because of the striping configuration.
Failed	<ul style="list-style-type: none">• RAID 0 volume: Indicates that one or more array disks are missing or have failed.• RAID 1 volume: Indicates that both array disks are missing or have failed.• RAID 5 or 10 volume: Indicates that two or more array disks are missing or have failed.
Unknown	Indicates that an unknown error was detected.

Table 6-20. Busy Volume States

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 a 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.

Table 6-21. General Parameters

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.

6.3.6.10 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.

Step 1: 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.

Step 2: Click Rename.

Step 3: Type a new volume name, and then click **OK**.

Note: Volume names must conform to the list of allowed characters. Please see section [6.1.3](#) for more information.

6.3.6.11 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 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 data before continuing.



6.3.6.11.1 Rebuilding from 'Status' (Manually)

Step 1: 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.

Step 2: Click **Rebuild to another disk** next to the volume the user wants to rebuild.

Step 3: In the **Rebuild Volume** dialog, select the disk that will replace the failed disk. Only compatible disks in a normal state will be displayed. Refer to [Volume Requirements](#) for more information.

Step 4: 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.

6.3.6.11.2 Rebuilding from 'Manage' (Manually)

Step 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.

Step 2: Click **Rebuild to another disk**, and then follow the procedure described above.

6.3.6.11.3 Auto Rebuild to Missing Disk

The Intel® VROC 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.

6.3.6.11.4 Rebuild Resumption

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

6.3.6.12 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.

Step 1: Under **Home**, in the disk properties pane, click **Reset volume to normal**.

Step 2: Click **Yes** to confirm.

Step 3: The page refreshes and the volume appears as normal.

6.3.6.13 Changing Volume Type

The user can choose to change the type of an existing volume based on the user's storage system needs.

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 SSDs connected to the CPU. Changing volume types does not require re-installation of the operating system.

Step 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.

Step 2: Click **Change type**.

Step 3: In the **Change Volume Type** dialog, type a new name if the user wants to change the default name.

Step 4: Select the new volume type, and then click **OK**. Caution: Once the data migration starts, the operation cannot be canceled.

Step 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 system will 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.

6.3.6.14 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 32MB 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 the 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 system will 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.

6.3.6.14.1 Increasing the Volume Size from Manage Array

Step 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.

Step 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.

Step 3: Click **Yes** to confirm. Caution: Once the data migration starts, the operation cannot be canceled.

Step 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: Until the system is restarted, continuing to use the product management tools may result in unknown errors being encountered.

6.3.6.14.2 Increasing the Volume Size from Manage Volume

Step 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.

Step 2: Click **Increase size**.

Step 3: Click **Yes** to confirm. Caution: Once the data migration starts, the operation cannot be canceled.

Step 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 on **My Computer**, select **Manage**, then in the console tree select **Disk Management**.

Until the system is restarted, continuing to use the product management tools may result in unknown errors being encountered.

6.3.6.15 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 the section [Connecting a Disk](#) 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 system will 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 section [Adding a Disk to an Array](#) for more information.

Step 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.

Step 2: Click **Add Disk**.

Step 3: Select the disk the user wants to use to increase the array capacity.

Step 4: Click **Add Disk**.

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

Step 5: 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 Management, click **Start**, right click on **My Computer**, select **Manage**, then in the console tree select **Disk Management**.

Until the system is restarted, continuing to use the product management tools may result in unknown errors being encountered.

6.3.6.16 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 Uninterruptable Power Supply (UPS).

Enabling the volume write-back cache is done with these steps:

Step 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.

Step 2: In the *Advanced* section, click **Enable**, and then click **Yes** to confirm. 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 is done with these steps:

Step 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.

Step 2: In the *Advanced* section, click **Disable**, and then click **Yes** to confirm. The page refreshes and the write-back cache is now disabled.

6.3.6.17 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 is done with these steps:

Step 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.

Step 2: Click **Initialize**.

Step 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 is done with these steps:

Step 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.

Step 2: Click **Verify**.

Step 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.

Step 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® VROC icon, or in the application under the *Home* pane.

Warning: The initialization 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.

6.3.6.18 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 to keep valuable data and the overall storage system healthy.

Step 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.

Step 2: Click **Verify**.

Step 3: Select the check box if the user wants errors found to be repaired automatically during the verification process.

Step 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.

6.3.6.19 RAID Level Migrations

The RAID level migration feature in the Intel® VROC products will enable the ability to convert the content of a drive (behind Intel® VMD or attached to the SATA and/or sSATA and/or tSATA 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 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.

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

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

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

Table 6-22. RAID Level Migration Examples Supported by Intel® VROC

Change Type From	To (SATA/sSATA/tSATA)	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** section.

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

6.3.6.20 Migrating From One RAID Level to Another

RAID level migration allows an existing RAID configuration to be migrated to another RAID configuration.

Note: Not all migrations are supported on all chipsets. The support varies depending on the chipset and the ports supported on the chipset.

For the migration options to be accessible, the minimum required drives for the RAID level must be met.

Follow the procedure below to perform a RAID migration:

Step 1: Start the Intel Virtual RAID on CPU GUI application (run as administrator):
Start Menu -> All Programs -> Intel -> Intel(R) Virtual RAID on CPU.

Step 2: From the **Volumes** pane, click the array or volume to which the user wants to modify. The volume properties now display on the right.

Step 3: Click **Change type**.

Step 4: In the **Change Volume Type** dialog, type a new name if the user wants to change the default name.

Step 5: Select the new volume type, and then click **OK**.

Step 6: The *Manage* page refreshes and reports the new volume type.

Step 7: After the migration starts, the user can view the migration progress under status.

When the status field indicates the volume as *Normal*, the migration is complete.

6.3.7 Managing Disks

The user must be logged on as 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.

6.3.7.1 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.

Table 6-23. List of Disk Properties from the Intel® VROC GUI

Parameter	Value	As Seen in the HII
Port	Reports the port number to which the disk or device is attached.	SATA: Port VMD: Root port number; Root port offset; Slot Number
Port Location	Reports whether the port is internal or external.	Information not available
Usage	Array Member: a disk that has been grouped with other disks to form an array containing RAID volumes.	Information not available - but status <i>RAID Member</i> indicates that state.
	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.	Information not available - but status <i>Spare</i> indicates that state.
	Available: the disk is physically connected to the computer, healthy, and available to be used in an array or as a spare disk.	Information not available - but status <i>non-RAID</i> indicates that state.

Parameter	Value	As Seen in the HII
	WARNING: Assigning an available disk to an array or marking it as a spare will permanently delete any existing data on that disk.	
	Unknown: the disk is available but contains metadata that cannot be displayed in the operating system. Even though the disk is reported as normal, the user will need to clear and reset the disk to make the disk available.	Information not available - but status <i>Unknown</i> indicates that state.
Status	Normal: the disk is present, functioning as expected.	Normal
	At risk: an impending error condition was detected on the disk, and it is now at risk of failure.	Information not available
	Missing: the disk is not present or physically connected to the computer.	Missing
	Failed: the disk has failed to properly complete read and write operations in a timely manner, and it has exceeded its recoverable error threshold.	Error Occurred
	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 <i>offline</i> because it is no longer an active member of the RAID volume. The Intel® VROC GUI (that contains all three of these drives) will not do anything with this drive until RAID metadata is cleared. If Drive-1 is removed and taken to another system, the Intel® VROC GUI will show the drive as <i>online</i> as part of a degraded RAID 1 volume.	Offline
Size	Reports the total capacity of the disk in Megabytes (MB) in the disk properties and in Gigabytes (GB) in the storage system view.	Size
Serial Number	Reports the manufacturer's serial number for the disk.	Serial Number
Model	Reports the model number of the disk.	Model Number
Firmware	Reports the version of the firmware found in the disk.	Information not available

Parameter	Value	As Seen in the HII
System Disk	Reports whether the disk contains system files that are required to start and run the operating system.	Information not available
Protection	Reports whether the disk is protected with a password.	Information not available
Disk Data Cache	Reports whether the disk supports this feature.	Information not available
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.	Information not available
Physical Sector Size	Reports the size of physical sectors on the disk (bytes).	Information not available
Logical Sector Size	Reports the size of logical sectors on the disk (bytes).	Block Size

6.3.7.2 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.

Step 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.

Step 2: Click **Mark as spare**.

Step 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.

6.3.7.3 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.

Step 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.

Step 2: Click **Reset to available**. The page refreshes and the disk usage is now reported as available.

6.3.7.4 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 the [Troubleshooting](#) section for guidelines on rebuilding a failed or degraded volume.

Step 1: Under **Home**, in the *Device* pane, locate the disk reported as at risk or failed. The user can also perform this action from the **Disk Properties** to the right.

Step 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.

6.3.7.5 Resetting a Disk to Online/Normal

After a disk was marked as *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.

Step 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.

Step 2: Click **Clear metadata** on the status line, next to *Offline*.

Step 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 should now be reported as *Normal*.

This process can also be performed using the Intel® VROC CLI and in the Pre-OS environment.

6.3.7.6 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 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

Step 1: Power off the user's computer.

Step 2: Replace the disk that reports a problem.

Step 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 the [Troubleshooting](#) section 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)

Step 1: Power off the user's computer.



Step 2: Connect the new disk.

Step 3: Turn the user's computer back on.

During the system startup, the application's Pre-OS 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.

6.3.7.7 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 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.

Table 6-24. Port Properties in the Intel® VROC GUI

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

6.3.7.8 Detecting a Spare Drive

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

6.3.8 Disk Monitor Service

Intel® VROC 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.

6.3.9 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 application or user interface control to configure the feature.

Support for the NVMe Deallocate/SATA TRIM command allows the OS to pass information to the Solid-State Drive (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 lifespan of SSDs.

6.3.10 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.

6.3.10.1 Uninstalling the Intel® VROC Software (Except the RAID Driver)

Step 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® VROC software except the RAID driver.

Step 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.

Step 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.

6.3.10.2 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!

Step 1: Enter the 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 the setup.

Step 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 Pre--OS status screen during boot, and the user should no longer see the Intel® VROC controller(s) in Device Manager.

Step 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.

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.

6.4 Intel® VROC Power Management

Intel® VROC 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)
- Hibernate 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.

6.4.1 Intel® VROC (VMD NVMe RAID) Power Management Hot Plug Limitation

Intel® VROC (VMD NVMe RAID) product does not support performing drive hot plug (removal/inserting) while the system is in a reduced power state from either a Sleep or Hibernate power management activity. This limitation does not apply to Intel® VROC (SATA RAID).

6.5 Intel® VROC Miscellaneous Features and Functionality

6.5.1 OS Installation

The Intel® VROC 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 or tSATA controllers.

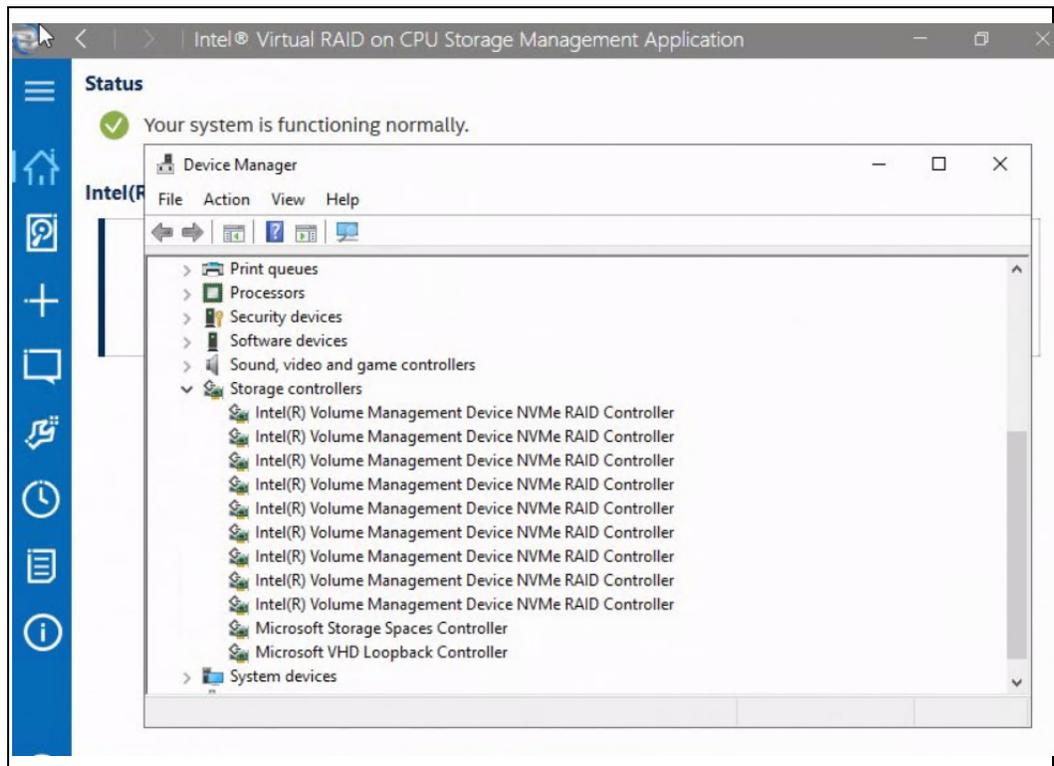
6.5.2 Using the Device Manager

Step 1: Open the *Start Menu* of the user's system and type the words **Device Manager**.

Step 2: Scroll down to the section that reads **Storage Controllers**.

Step 3: Click on **Storage Controllers** to expand the menu. The exact wording will vary some depending on the operating system involved, Windows* 10 shows this as *Microsoft* Storage Spaces Controllers*.

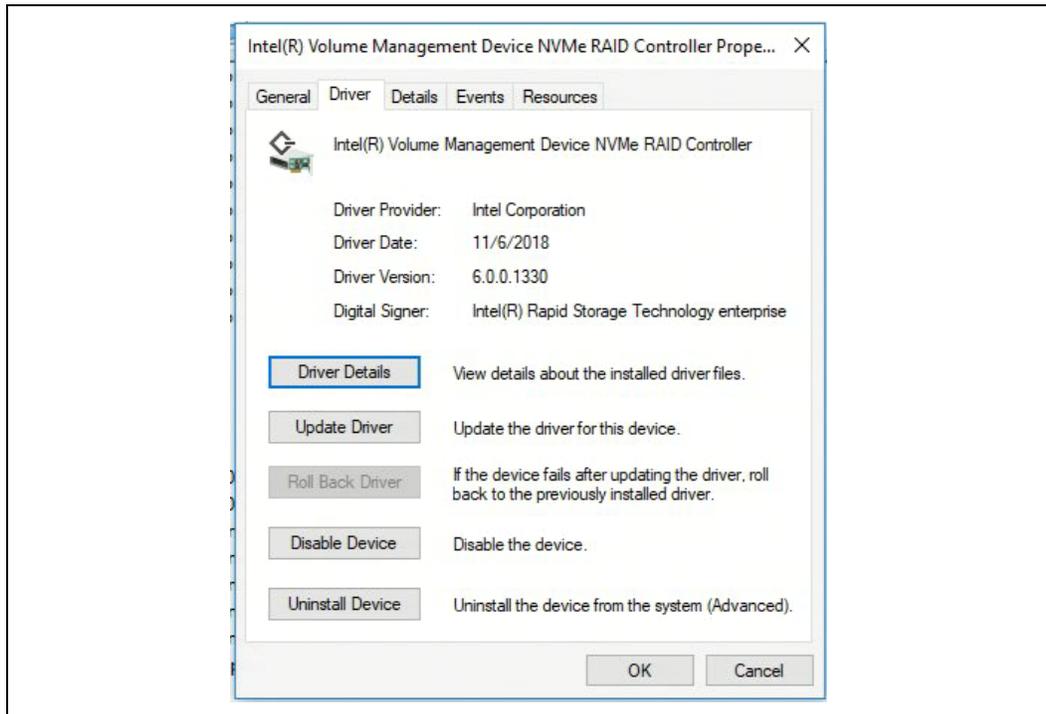
Figure 6-24. Device Manager Screen



Step 4: Right click on **Microsoft* Storage Spaces Controllers** and select **Properties**.

Step 5: Select the tab that reads **Driver**. *Driver Version* denotes the version of the drivers that are installed on the system controller for this system.

Figure 6-25. Driver Details Screen



6.5.3 Disk Write Cache

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

6.5.4 Write Back Cache

Intel® VROC 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.

6.5.5 4K Native Sector Sizes Drives

Intel® VROC family of products will support the following drive sector sizes:

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

Note: The only limitation is that mixing 4K and 512(b) drives into a single RAID volume is not allowed. This includes the creation, expansion, and migration features.

This feature is only supported in UEFI mode.

6.5.6 Intel® VROC Microsoft* WHQL Drivers

The production version of the Intel® VROC drivers shall be released after passing Microsoft* WHQL logo requirements for storage devices under all supported operating systems.

6.5.7 Intel® VROC Microsoft* UWD Support

The Intel® VROC family of products meets Microsoft's UWD *D*, *C* and *H* compliance criteria. Intel® VROC supports this feature only on 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.

6.5.7.1 INF Installation Support

The Intel® VROC 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 click and select install. Once prompted, reboot the system to ensure the installation process takes full affect.

6.5.7.2 Windows* Application Store Support

The Intel® VROC family of products has provided/uploaded the Intel® VROC 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.

6.5.7.3 Stand Alone Intel® VROC GUI

The Intel® VROC family of products supports the ability to load the Intel® VROC GUI without running the installer or going to the Microsoft* Application Store. Within the Intel® VROC package there is a standalone compressed file that can be copied to the target system, unzipped, and used to load/install the Intel® VROC UWD GUI application. As with the application uploaded to the Windows* Application Store, this process requires that all the Intel® VROC drivers have already been installed using the .inf installation process.

When the stand-alone Intel® VROC GUI detects that the installed driver version does not match (newer or older) than the Intel® VROC GUI version, it will display a warning message that the two versions need to be match.

Following is an example of how the Intel® VROC UWP GUI package might be installed manually:

Step 1: Open PowerShell (admin) by right click on start and select PowerShell (admin).

Step 2: When prompted with PowerShell window, navigate to release folder/KIT/UserInterface/UI release folder/UWP using `cd` command. The files in the destination folder should be similar to those in the screenshot below.

Figure 6-26. Intel® VROC UWP GUI Package

Name	Date modified	Type	Size
Add-AppDevPackage.resources	8/6/2019 7:44 PM	File folder	
Add-AppDevPackage	8/1/2019 3:19 PM	Windows PowerS...	34 KB
Intel(R) VROC_6.1.9584.0_x64	8/1/2019 3:19 PM	APPXBUNDLE File	11,703 KB
Intel(R) VROC_6.1.9584.0_x64.appxsym	8/1/2019 3:19 PM	APPXSYSM File	1,771 KB
Intel(R) VROC_6.1.9584.0_x64	8/1/2019 3:19 PM	Security Certificate	2 KB
Microsoft.VCLibs.x64.14.00.Desktop	8/6/2019 5:01 PM	APPX File	6,566 KB

Step 3: Enable UWP app side-loading:

```
REG ADD HKLM\Software\Policies\Microsoft\Windows\Appx /v AllowAllTrustedApps /t
REG_DWORD /d 1 /f
```

Step 4: Install the certificate:

```
CERTUTIL -addstore -f "TrustedPeople" "<Certificate File.cer>"
```

Step 5: Run DISM to install the app:

```
DISM /Online /Add-ProvisionedAppxPackage /PackagePath:"<appxbundle file>"
/DependencyPackagePath:"Microsoft.VCLibs.x64.14.00.Desktop.appx"
/LogPath:dism.log /SkipLicense
```

Step 6: Restart the system.

Step 7: After a restart, launch the Intel® Virtual RAID on CPU application. When prompted with UAC dialog, press **Yes** to launch it.

Note: The files may differ depending on the release version. Substitute filenames inside `<>` with appropriate file names.

When the Intel® VROC installation executable is run, if it detects that the Intel® VROC GUI, from Windows* AppStore, is present in the platform it will be removed and install the GUI package (contained in the installer).

The Intel® VROC Installer GUI will also display a warning message when it detects the version of the driver is mismatched to itself.

For UWP in Intel® VROC 8.0 GUI, the name of GUI will be Intel® VROC management application. Follow the below steps for UWP in Intel® VROC 8.0 GUI:

Step 1: Navigate to the release folder using `cd` command.

Step 2: Enable UWP app side-loading:

```
REG ADD HKLM\Software\Policies\Microsoft\Windows\Appx /v AllowAllTrustedApps /t
REG_DWORD /d 1 /f
```

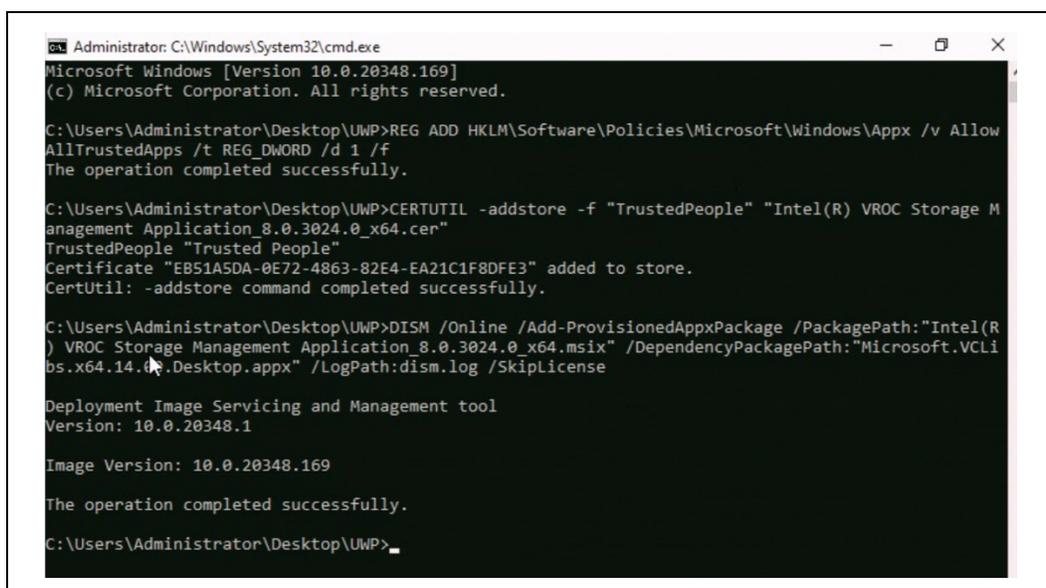
Step 3: Install the certificate:

```
CERTUTIL -addstore -f "TrustedPeople" " Intel(R) VROC Storage Management
Application_8.0.3024.0_x64.cer"
```

Step 4: The above two commands are only required for not signed Production Version (PV) driver. For PV with license file UWP installation, use the below single script command:

```
DISM /Online /Add-ProvisionedAppxPackage /PackagePath:"Intel(R) VROC Storage
Management Application_8.0.3462.0_x64.msix"
/DependencyPackagePath:"Microsoft.VCLibs.140.00.UWPDesktop_14.0.30704.0_x64__8wek
yb3d8bbwe.appx" /LogPath:dism.log /SkipLicense
```

Figure 6-27. Intel® VROC UWP GUI Installation Process



6.5.8 Microsoft* Hyper-V*

Microsoft* Hyper-V* can be installed onto an Intel® VROC (SATA RAID or VMD NVMe RAID) boot volume. Intel® VROC supports Microsoft* Hyper-V* virtual machine management on both Intel® VROC (SATA RAID or VMD NVMe RAID) boot and data volumes.

6.5.9 Intel® SMART Event Support

The Intel® VROC family of products will provide support for SMART alerts for NVMe and SATA disks. A SMART drive *critical warning* 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.



6.5.10 Event Log Support

The Intel® VROC family of products will log member disk error status to an event log stored in RAID metadata. See section [Storage System Events](#) for more details.

6.5.11 2 Terabyte (TB) Disk Size

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

6.5.12 2 Terabyte (TB) Volume Size

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

6.5.13 Email Alerting and Notification

The Intel® VROC family of products will support email notification of certain storage events.

6.5.13.1 Email Notification UI Visible Enable/Disable

Intel® VROC will support email notification of certain storage events. See section [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.

6.5.13.1.1 Configuration

The Intel® VROC 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).

- The user can enable/disable the email notification feature.
- The 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.
- The 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)

The user can configure the email alert/notifications to send test emails to all addresses specified.

6.5.13.1.2 Email Message Format

Message Header:

- Item 1. Return email address: Email address of the originating computer.
- Item 2. Recipient email address: Email address of computer receiving the email notification.
- Item 3. Subject: System formatted subject content with product name, the storage system event level, and the hostname of the originating computer.

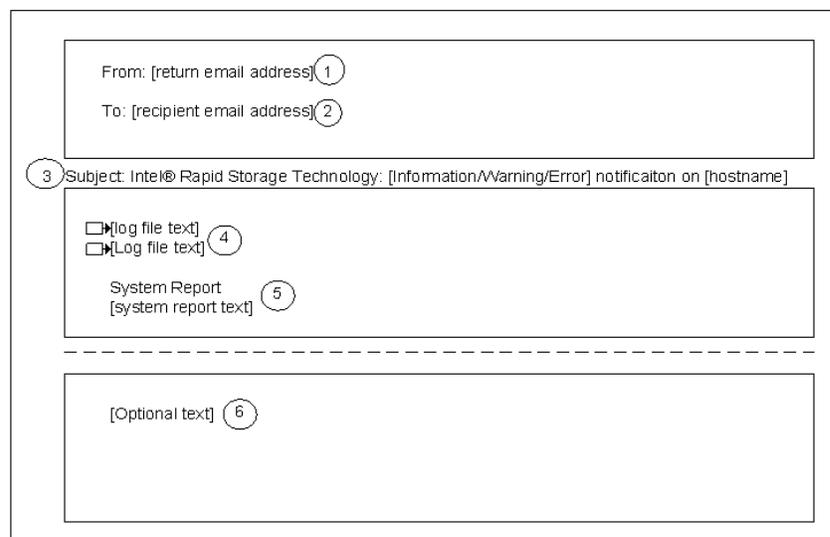
Message Body:

- Item 4. Log file text: Contains the text of the event as it is displayed in the event log.
- Item 5. System report: Contains the system configuration information of the originating computer as seen in the Intel® VROC GUI Preferences page.

Optional Text:

- Item 6. 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).

Figure 6-28. Email Message Format



6.5.13.1.3 Protocol Support

Email alert shall support SMTP host & SMTP port.

Note: Intel® VROC will check the SMTP Host and Port settings but does not check credentials for clients.



6.5.13.1.4 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.

6.5.14 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 on the Windows* operating system installed:

Table 6-25. Microsoft* .NET Framework Impact for Intel® VROC GUI

Intel® VROC GUI	
Server 2k12 R2	Install Latest .NET Framework
Server 2k16	Install Latest .NET Framework
Windows* 2k19	No Impact
Win 10 RS3	Install Latest .NET Framework
Win 10 RS4	No Impact
Win 10 RS5	No Impact
Win 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 instructions to help in this process:

Step 1: Download the latest version of .NET Framework from Microsoft*.

Step 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>).

Step 3: Copy the compressed file to a USB drive.

Step 4: Copy the compressed file from the USB drive to the Download directory of the platform being configured.

Step 5: Uncompressed the file.

Step 6: Run the executable file as administrator.

For more information 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.

6.6 Intel® VROC Package Updates

Intel® VROC family of products supports upgrading to new release packages. Each Intel® VROC release package contains the Pre-OS components, the OS drivers, the product installation application, and tools. Intel recommends that all 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 platform BIOS with the corresponding Intel® VROC Pre-OS 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 drivers and tools are designed to be backwards compatible with older Intel® VROC Pre-OS images (unless otherwise outlined in the products release notes). The Intel® VROC Pre-OS 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 Pre-OS 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.

6.6.1 Intel® VROC Package Upgrade

The Intel® VROC package upgrade is the process of transitioning to the latest/newest version of the Intel® VROC Pre-OS and/or Intel® VROC OS components. The table below shows the compatibility matrix outline of the Intel® VROC Windows* driver component vs. the Intel® VROC Pre-OS environment.



Table 6-26: Intel® VROC OS Driver vs. Pre-OS Compatibility Matrix

Intel VROC Driver Versions for Windows*														
Pre-OS Versions	6.2	6.3	7.0	7.5	7.5	7.8	8.0	8.2	8.5	8.6	9.0	9.1	9.2	9.3
6.3 Pre-OS	NS	S	S	S	S	S	S	S	S	S	S	S	S	S
7.0 Pre-OS	NS	NS	S	S	S	S	S	S	S	S	S	S	S	S
7.5 Pre-OS	NS	NS	NS	S	S	S	S	S	S	S	S	S	S	S
7.6 Pre-OS	NS	NS	NS	NS	S	S	S	S	S	S	S	S	S	S
7.8 Pre-OS	NS	NS	NS	NS	NS	S	S	S	S	S	S	S	S	S
8.0 Pre-OS	NS	NS	NS	NS	NS	NS	S	S	S	S	S	S	S	S
8.2 Pre-OS	NS	S	S	S	S	S	S	S						
8.5 Pre-OS	NS	S	S	S	S	S	S							
8.6 Pre-OS	NS	S	S	S	S	S								
9.0 Pre-OS	NS	S	S	S	S									
9.1 Pre-OS	NS	S	S	S										
9.2 Pre-OS	NS	S	S											
9.3 Pre-OS	NS	S												

Note: “Not Supported” is notated as “NS” within the table, while “Supported” combination of Pre-OS and Windows Driver version are denoted as “S”

Intel® VROC Windows* OS backwards compatibility to Intel® VROC Pre-OS is supported to N-3 versions. Validation will be on N-1 version.

Supported (as used in the above table) means that issues will be accepted and processed with all other supported configurations.

6.6.2 Intel® VROC Package Downgrade

The Intel® VROC package downgrade is the process of transitioning to an older version of the Intel® VROC Pre-OS and/or Intel® VROC OS components. Intel does not recommend downgrading as a general practice. However, there may be certain cases where it may be required. If it is deemed necessary to downgrade, the following order should be followed:

1. Update/downgrade to platform BIOS with a supported Intel® VROC Pre-OS image first (e.g., downgrade from platform BIOS with Intel® VROC 7.5 Pre-OS to platform BIOS with Intel® VROC 7.0 Pre-OS).
2. Update/downgrade the Intel® VROC OS driver package to the desired package (e.g., downgrade from Intel® VROC 7.5 down to Intel® VROC 7.0).





7 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.

Table 7-1. Intel® VROC SATA RAID Features

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

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.

7.1 Intel® VROC (SATA RAID) Pre-OS Features and Functionality

7.1.1 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.

Table 7-2. BIOS Control Feature Set (BCFS) Bits Description

BCFS Bit Number	Bit Meaning	Values		Additional Info
0	Enable/disable RAID 0	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 RAID 1	RAID type disabled	0x0	
		RAID type enabled	0x1	
2	Enable/disable RAID 10	RAID type disabled	0x0	

BCFS Bit Number	Bit Meaning	Values		Additional Info
		RAID type enabled	0x1	
3	Enable/disable RAID 5	RAID type disabled	0x0	
		RAID type enabled	0x1	
4	RESERVED			
5	RESERVED			
6	Enable/disable unlock of HDD in OS	OS can unlock HDDs	0x1	
		OS cannot unlock HDDs	0x0	
7	Enable/disable LED SGPIO	Feature disabled	0x0	
		Feature enabled	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			
10 - 12	RESERVED			
13 - 14	RESERVED			
15	RESERVED			

Note: This document does not cover details on how to setup a system BIOS. For that level of information, contact the user's platform's BIOS vendor or the user's Intel field representative to put the user in contact with the appropriate Intel BIOS support personnel.

7.1.1.1 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.



7.1.1.2 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.

7.1.2 ATAPI on SATA/sSATA/tSATA Controller

The Intel® VROC (SATA RAID) product will provide support for ATAPI devices connected to the SATA/sSATA/tSATA controller.

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

Note: If using an ATAPI device connected to either the sSATA or tSATA controllers to install a Windows operating system on a VROC-managed drive or RAID volume on another controller (ex: SATA controller, VMD controller), the installation may be blocked with a pop-up error message in the Windows pre-install window stating that a media driver your computer needs is missing. This is a limitation with Windows. Please see below for 2 workaround scenarios:

1. If the SATA controller is set to "AHCI" mode in the BIOS, connect the ATAPI device to the that controller (not the sSATA or tSATA controllers) and proceed with the Windows installation process as usual.
2. Set the sSATA or tSATA controller, whichever the ATAPI device is connected to, in the BIOS to "RAID" mode. Then during the Windows pre-install phase, load the VROC driver for that controller. Then proceed with the OS installation process as usual.

7.1.3 Support for EFI_ATA_PASS_THRU_PROTOCOL

Intel® VROC (SATA RAID) provides limited support for providing information on the SATA drives managed by Intel® VROC through the **EFI_ATA_PASS_THRU_PROTOCOL**. This feature will be supported for drives in a RAID volume, pass-through drives (not in a RAID volume), drives marked as *Incompatible*, drives marked as *Spare*s and *Journaling* drives. The following is the list of supported ATA pass-through commands for RAID members:

- CHECK_POWER_MODE DATA_SET_MANAGEMENT
- DATA_SET_MANAGEMENT_XL
- DOWNLOAD_MICROCODE
- EXECUTE_DEVICE_DIAGNOSTIC
- GET_PHYSICAL_ELEMENT_STATUS
- IDENTIFY_DEVICE
- IDLE
- IDLE_IMMEDIATE
- MUTATE_EXT
- NCQ_NON_DATA

- NOP
- READ_LOG_EXT
- REQUEST_SENSE_DATA_EXT
- SECURITY_DISABLE_PASSWORD
- SECURITY_ERASE_PREPARE
- SECURITY_ERASE_UNIT
- SECURITY_FREEZE_LOCK
- SECURITY_SET_PASSWORD
- SECURITY_UNLOCK
- SET_DATE_AND_TIME_EXT
- SET_FEATURES
- SLEEP
- SMART
- STANDBY
- STANDBY_IMMEDIATE
- WRITE_LOG_EXT
- DATA_SET_MANAGMENT
- FLUSH_CACHE_EXT

For any not listed command, the `EFI_ATA_PASS_THRU_PROTOCOL.PassThru` function will return `EFI_UNSUPPORTED` status.

7.1.4 **EFI_ATA_PASS_THRU_PROTOCOL Supported Functions**

Intel® VROC (SATA RAID) will support and handle the following `EFI_ATA_PASS_THRU_PROTOCOL` functions on all drives including RAID member drives:

- `BuildDevicePath()`
- `GetDevice()`
- `GetNextDevice()`
- `GetNextPort()`
- `PassThru()`

Intel® VROC (SATA RAID) will report as `EFI_UNSUPPORTED` for the following `EFI_ATA_PASS_THRU_PROTOCOL` functions on all drives including RAID member drives:

- `ResetDevice()`
- `ResetPort()`



7.2 Intel® VROC (SATA RAID) Miscellaneous Features and Functionality

7.2.1 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 or tSATA controllers.

7.2.2 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/tSATA controllers in RAID mode. The following commands are supported by Intel® VROC:

- 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

7.2.3 NCQ (SATA/sSATA/tSATA)

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).

7.2.4 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 the SATA controller on the platform.

7.2.5 Intel® VROC (SATA RAID) Maximum Number of sSATA/tSATA drives in a RAID Volume

The Intel® VROC (SATA RAID) product will support to create sSATA/tSATA RAID volumes from maximum number of drives equal to the number of ports on the sSATA/tSATA controller on the platform.

7.2.6 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/tSATA controller).

7.2.7 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:

- Support for `EFI_EXT_SCSI_PASS_THRU_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.

7.2.8 Intel® VROC (SATA RAID) UEFI Support of for `EFI_ATA_PASSTHRU`

Intel® VROC 7.6 VROC (SATA RAID) introduces support for `EFI_ATA_PASS_THRU` protocol support by the Intel® VROC (SATA RAID) UEFI driver. The following specific options are supported in Intel® VROC 7.6:

- `PassThru;`
 - `IDENTIFY`
 - `ATA_READ_LOG_EXT`
- `GetNextPort;`
- `GetNextDevice;`
- `BuildDevicePath;`
- `GetDevice;`

For the following commands: `ResetPort` and `ResetDevice;` an `EFI_STATUS` equal to `EFI_UNSUPPORTED` will be returned.

For the `PassThru` functionality, Intel® VROC (SATA RAID) provides for the following support specific ATA commands:



- For pass-thru drives (drives not part of an Intel® VROC (SATA RAID) volume) all commands are supported.
- For Intel® VROC RAID (SATA RAID) member drives (RAID members, spares and/or journaling drives) Intel® VROC (SATA RAID) will support:
 - ATA_COMMAND_CHECK_POWER_MODE
 - ATA_COMMAND_DATA_SET_MANAGEMENT
 - ATA_COMMAND_DATA_SET_MANAGEMENT_XL
 - ATA_COMMAND_DOWNLOAD_MICROCODE
 - ATA_COMMAND_EXECUTE_DEVICE_DIAGNOSTIC
 - ATA_COMMAND_GET_PHYSICAL_ELEMENT_STATUS
 - ATA_COMMAND_IDENTIFY_DEVICE
 - ATA_COMMAND_IDLE
 - ATA_COMMAND_IDLE_IMMEDIATE
 - ATA_COMMAND_MUTATE_EXT
 - ATA_COMMAND_NCQ_NON_DATA
 - ATA_COMMAND_NOP
 - ATA_COMMAND_READ_LOG_EXT
 - ATA_COMMAND_REQUEST_SENSE_DATA_EXT
 - ATA_COMMAND_SECURITY_DISABLE_PASSWORD
 - ATA_COMMAND_SECURITY_ERASE_PREPARE
 - ATA_COMMAND_SECURITY_ERASE_UNIT
 - ATA_COMMAND_SECURITY_FREEZE_LOCK
 - ATA_COMMAND_SECURITY_SET_PASSWORD
 - ATA_COMMAND_SECURITY_UNLOCK
 - ATA_COMMAND_SET_DATE_AND_TIME_EXT
 - ATA_COMMAND_SET_FEATURES
 - ATA_COMMAND_SLEEP
 - ATA_COMMAND_SMART
 - ATA_COMMAND_STANDBY
 - ATA_COMMAND_STANDBY_IMMEDIATE
 - ATA_COMMAND_WRITE_LOG_EXT

7.2.9 RAID Volume Read Cache

Intel® VROC (SATA RAID) will support the ability to enable/disable RAID Volume Read Cache through the Intel® VROC GUI. RAID Volume Read Cache will be enabled by default.

7.2.10 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, sSATA and tSATA controller for those drives that support this feature.

7.2.11 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.



Appendix A Related Documentation

A.1 Relevant Specifications

Specification	Location
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 I Specification	http://www.serialata.org
SATA II Specification	http://www.serialata.org
SATA III Specification	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
NVM Express Revision 1.2 Specification	https://nvmexpress.org/wp-content/uploads/NVM_Express_1_2_Gold_20141209.pdf



Appendix B Storage System Events

The following table lists storage system events detected by the monitor service (*IAStorDataMgrSvc*):

Event Type	Event Level	String	Event Displayed		
			NAI ¹ (Notification Area Icon)	Event Log	E-Mail Notify ²
Disk Triggered Events					
Failed	Error	Disk on port {n}: Failed. Open the application for details.	Yes	Yes	Yes
S.M.A.R.T.	Warning	Disk on port {n}: At risk. Open the application for details.	Yes	Yes	Yes
Unlocked	Info	Disk on port {n}: Unlocked.	Yes	Yes	Yes
Added	Info	Disk on port {n}: Detected.	Yes	Yes	Yes
Removed	Info	Disk on port {n}: Removed.	Yes	Yes	Yes
Volume Triggered Events					
Failed	Error	Volume {0}: Failed. Open the application for details.	Yes	Yes	Yes
Degraded	Warning	Volume {0}: Degraded. Open the application for details.	Yes	Yes	Yes
Detected	Info	A new volume was found.	Yes	Yes	Yes
RebuildComplete	Info	Volume {0}: Rebuilding complete.	Yes	Yes	Yes
VerifyStop	Info	Volume {0}: Verification complete.	Yes	Yes	Yes
VerifyAndRepairStop	Info	Volume {0}: Verification and repair complete.	Yes	Yes	Yes
MigrationComplete	Info	Volume {0}: Data migration complete.	Yes	Yes	Yes
InitializeComplete	Info	Volume {0}: Initialization complete.	Yes	Yes	Yes
Unlocked	Info	Volume {0}: Unlocked.	Yes	Yes	Yes
NotPresent	Info	Volume {0}: No longer present on system.	Yes	Yes	Yes
RebuildStarted	Info	Volume {0}: Rebuilding in progress.	Yes	No	No
VerifyStarted	Info	Volume {0}: Verification in progress.	Yes	No	No

Event Type	Event Level	String	Event Displayed		
			NAI ¹ (Notification Area Icon)	Event Log	E-Mail Notify ²
VerifyAndRepairStarted	Info	Volume {0}: Verification and repair in progress.	Yes	No	No
MigrationStarted	Info	Volume {0}: Data migration in progress.	Yes	No	No
InitializeStarted	Info	Volume {0}: Initialization in progress.	Yes	No	No
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

1. NAI true only if the user selected to receive notification under *Preferences* in the UI.
2. Refer to the section [Email Alerting and Notification](#) for email feature support.

The following table outlines the events that will be logged and their corresponding event.

Event ID	Message Definition	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	Message Definition	Severity/Level	Message
4112	DISK_MANAGE_UNPLUG_LOG	Informational	Disk %2 has been managed 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.
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.

Event ID	Message Definition	Severity/Level	Message
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 (VROC 8.5 and older)	Warning	IO on %2 has failed. (Supported up to VROC 8.5 PV. Replaced with 4162, 4163 and 4164 below starting w/ VROC 8.6 PV)
4156	IO_TIMEOUT_LOG	Warning	IO on %2 has reached timeout.
4157	BAD_BLOCK_DETECTED	Warning	Bad block detected on drive %2 under LBA %3.
4158	BAD_BLOCK_FIXED	Informational	Bad block successfully fixed on drive %2 under LBA %3.
4159	UNRECOVERABLE_BAD_BLOCK	Error	Bad block under drive's LBA %2 which is member of a RAID volume %3 could not be recovered.

Event ID	Message Definition	Severity/Level	Message
4160	AHCI_PORT_INTERRUPT_STATUS	Warning	AHCI port %2 interrupt error status: PxIS = %3.
4162	ODD_EMPTY_DRIVE_IO_FAILED_LOG (VROC 8.6+)	Warning	I/O issued to empty Optical Disk Drive (media not present) model: <Model>, serial: <SerialNumber>, port: <PortNumber>.
4163	SATA_DRIVE_IO_FAILED_LOG (VROC 8.6+)	Warning	I/O on drive model: <Model>, serial: <SerialNumber>, port: <PortNumber> has failed.
4164	NVME_DRIVE_IO_FAILED_LOG (VROC 8.6+)	Warning	I/O on drive model:<Model>, serial: <SerialNumber> has failed.

§§

Appendix C External Hardware Capability

C.1 Non-Intel® NVMe Drives

For the latest list of supported NVMe SSDs can be found at the [Intel® Virtual RAID on CPU \(Intel® VROC\) Supported Configurations](#).

C.2 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® VROC (SATA RAID) Port Bitmap Implementation

D.1 UEFI Driver

The Intel® VROC 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).

D.2 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 SDD 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 E 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® VROC's online support site](#) which provides the user with self-help resources and contact support options.

E.1 Failed Volumes

RAID 0	
<p>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.</p>	
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 Pre-OS (UEFI) 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 Pre-OS (UEFI) 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. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 0 volume using the new disk. If the failed disk was part of the system volume, the user will also need to reinstall the operating system.
RAID 1	
<p>A RAID 1 volume is not reported as failed because two drives are mirror image.</p>	

RAID 5 A RAID 5 volume is reported as failed when two or more of its members have failed.	
Cause	Solution
Two or more array disks failed	In most cases, the volume cannot be recovered and any data on the volume is lost. However, before deleting the volume, 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. This procedure deletes the failed volume: <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 Pre-OS (UEFI) 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. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 5 volume using the new disks. If the failed disk was part of the system volume, 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	In most cases, the volume cannot be recovered and any data on the volume is lost. This procedure deletes the failed volume: <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 Pre-OS (UEFI) 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. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. The user will then need to reinstall the operating system on the new volume.

<p>Two or more array disks failed</p>	<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 Pre-OS (UEFI) 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. Select Delete RAID Volume from the main menu. 5. From the Delete Volume menu, select the failed RAID volume, using the up and down arrow keys. 6. Press the 'Delete' key to delete the volume, then 'Y' to confirm. 7. Create a new RAID 10 volume using the new disks. 8. The user will then need to reinstall the operating system on the new volume.
---------------------------------------	---

NOTE: Intel makes no guarantee of successful recovery from a failed state. This must be treated as a last chance effort and there is no guarantee that there won't be some data loss. Intel always recommends recovering a failed RAID volume by recreating the RAID volume from scratch and restore the data from the latest platform image backup.

E.2 Degraded Volumes

<p>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.</p>
<p>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.</p>
<p>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.</p>

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>

E.3 Other Volume States

Incompatible	
Cause	Solution
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.</p> <p>Here are the user's options:</p> <ol style="list-style-type: none"> 1. Reconnect the volume to the computer where the volume was originally created and continue using it. 2. Delete the volume, and then create a new volume with a RAID configuration that is supported by the current system. Follow the procedure described above to delete the volume. <p>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.</p>
Indicates that the Intel® VROC Upgrade Key is incorrect or missing.	<p>In this situation, volume data may not be accessible to the operating system.</p> <p>Here are the user's options:</p> <p>Install the proper Intel® VROC Upgrade Key</p>
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 Intel® VROC UEFI HII. Refer to the user's manual accessible from the Online Support area for details on using the Intel® VROC UEFI HII.
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.
--	---

E.4 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>

State	Cause	Solution
	<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>

E.5 Drive Information Structure Drive Status Values

Values	Status
0h	Unknown
1h	Offline – The drive is an array member of another RAID volume and as such, could not be configured
2h	Missing – The drive is not present or physically connected to the computer
3h	Failed – The drive has failed to properly complete read and write operations in a timely manner and it has exceeded its recoverable error threshold
4h	At risk (SMART event) – An impending error condition was detected on the drive, and it is now at risk of failing.
5h	Incompatible – Drive usage type is not supported in this configuration. It may occur when this drive has been used with previous versions of VROC or this drive does not meet requirements of SKU, which is already active. User can clear metadata from the drive and use it for another purpose
6h	Normal – The drive is present, functioning as expected, and unlocked
7h	Locked – The drive is locked
8h	Unsupported – This NVMe SSD vendor or model is not currently on supported compatibility list

E.6 Software Errors

Message	Cause	Solution
An unknown error occurred while running this application. If the problem persists, 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.

Message	Cause	Solution
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. 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. 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.

