Intel® VROC

Intel® Virtual RAID on CPU

Product 30-3-30

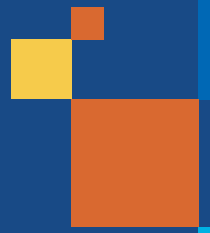

intel®

Notices and Disclaimers

Performance varies by use, configuration and other factors. Learn more at www.Intel.com/PerformanceIndex.

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details. No product or component can be absolutely secure.

Your costs and results may vary.

Intel technologies may require enabled hardware, software or service activation

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

Intel® Virtual RAID on CPU

4th Gen Intel® Xeon® Processors with Integrated RAID

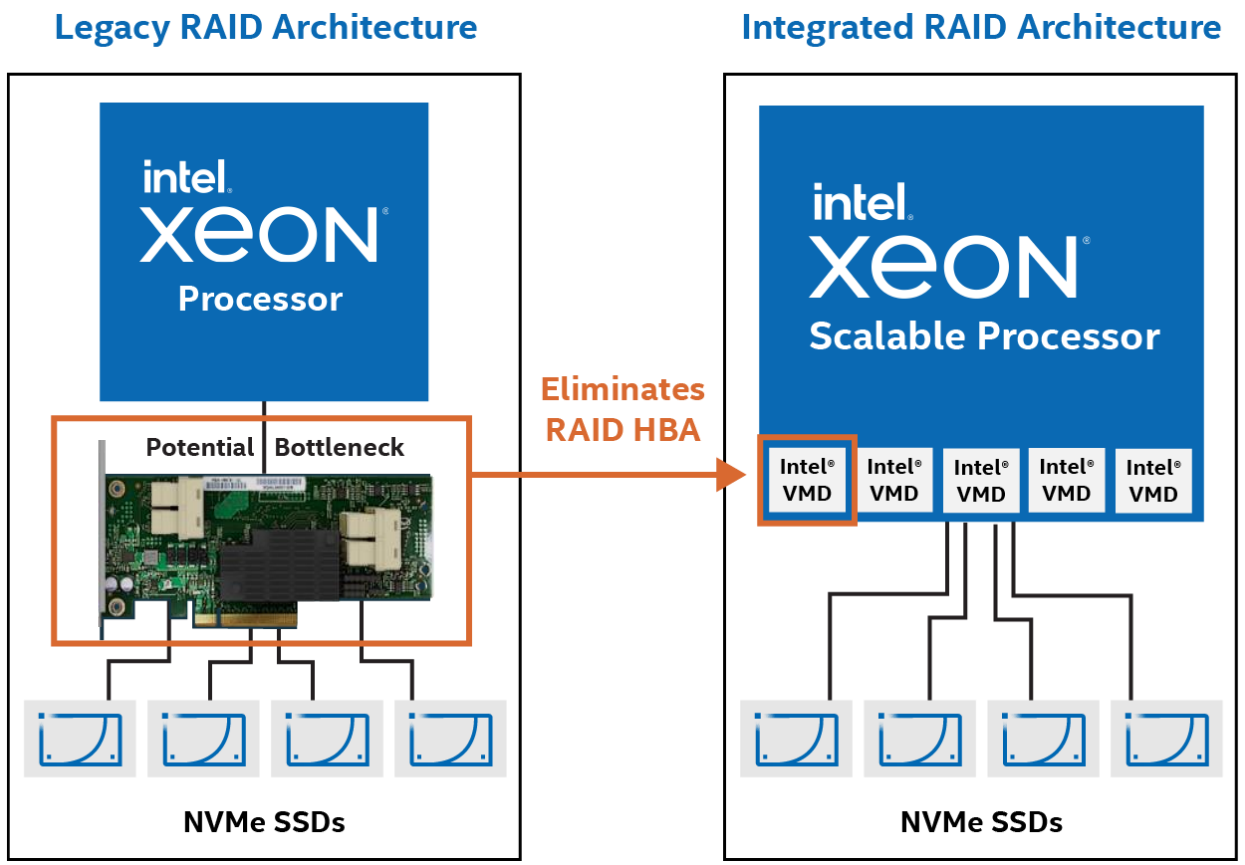

Performance
Up to **66%** increase (IOPS)
Latency reduced by **40%**

Cost
Up to **70%** lower

Power Efficiency
Up to **62%** better

See backup B for configuration details. Results may vary

Intel® VROC Product Family

Features		Pass-thru (Included with PCH)	Standard SKU	Premium SKU	Boot SKU **
Intel® VMD	Hot-plug/ Fault Isolation	✓	✓	✓	✓
	LED Management	✓	✓	✓	✓
	3rd Party SSD Support (non-RAID)	✓	✓	✓	✓
NVMe RAID	3rd party SSD support (RAID)	-	✓	✓	✓
	Bootable RAID	-	✓	✓	RAID1 only
	RAID 0/1/10	-	✓	✓	RAID1 only
	RAID 5	-	-	✓	-
SATA RAID	SATA RAID on PCH • Bootable RAID • RAID 0/1/5/10	✓	✓	✓	✓

** available with 5th Gen Intel® Xeon® Processors

Intel® VROC Product Roadmap

Ice Lake	Sapphire/Emerald Rapids
❑ Gen4 PCIe	❑ Out-Of-Band management support
❑ Boot & Data RAID 0/1/5/10 (Windows and Linux)	❑ Firmware Management Tools for VMWare ESXi
❑ Boot & Data RAID1 support for VMWare ESXi	❑ SED Local Key Management
❑ Hot-plug / LED Management / Fault Isolation	❑ Data RAID5 support for VMWare ESXi **(EMR) (2024)**
❑ Self-Encrypting Drive (SED) Key Management	❑ Boot Only SKU **(EMR)**
❑ Intel® VMD Direct-Assign for Virtualization	❑ Secure Erase feature for NVMe **(EMR)**
	❑ Intel OnDemand licensing **(EMR)**
OS: Windows, VMWare ESXi, Linux	

Intel® VROC is Integrated RAID

RAID Features	HW RAID	VROC	SW RAID	Intel® VROC Comment
Error Handling Isolation	✓	✓	x	Intel® VMD isolates SSD error/event handling from OS to reduce system crash or reboot due to error
Reliability	✓	✓	x	Removes HBA single point-of-failure, less HW Supports auto rebuild on spare devices
Boot support	✓	✓	x	Redundant system volume = less down-time/crashes
Complete Management Tools	✓	✓	x	UEFI, GUI, CLI, remote web, deployment tools. Compatible with BMC
SED Key Management	✓	✓	x	Xeon-based Platform integrated Key Management solution
Lower power requirement	x	✓	✓	No additional HW ensures lower power consumption
Supply Chain	x	✓	✓	No impact from supply chain constraints
Easily upgraded	x	✓	✓	Software update vs new HW purchase
Less hardware required	x	✓	✓	No need for HBA, BBU. Save power and PCIe* lanes

Intel® VROC Value

Eliminates supply chain concerns

↑70% cost saving (vs HBA)

Improves NVMe performance ↑66%

Improves reliability ↑62% power efficiency

Call-2-Action: Move to NVMe & replace RAID HBA

See backup B for configuration details. Results may vary

Intel® VROC Use Cases

Boot

- Protects system from failure
- OS, vSAN (hypervisor), HCI
- Node and system data

DC Edge

- Fast access to storage (NVMe)
- Performance (demanding app/ workload)
- Realtime edge processing

Telco/Embedded

- Serviceability (repair, config, boot support)
- System integrity, faulty recovery, resiliency
- Space limitation

DB Application

- Performance
- High-Capacity RAID arrays
- Integrity/Redundancy

Intel® VROC Value Proposition!

Problem

- RAID cards create a bottleneck which negatively impacts demanding applications & workloads
- Doesn't deliver NVMe SSD full performance potential (8GB/s)

Solution → Intel® VROC

1) Eliminates RAID Card
2) Drives down system cost
3) Improves performance
4) Maintains high availability

Benefits

- High Reliability, Availability and Serviceability (RAS)
- Doesn't use PCIe slots
- Reduce sys complexity/BOM

Competitive Advantage

- Differentiator from Dell
- Differentiator from AMD

Boost Customer Workloads

- 66% better performance
- 70% cost savings
- 62% better power efficiency
- Quicker access to storage data means fast decisions
- Cost effective solution that improve productivity

See backup B for configuration details. Results may vary

Intel® VROC: Integrated RAID

Legacy RAID Architecture

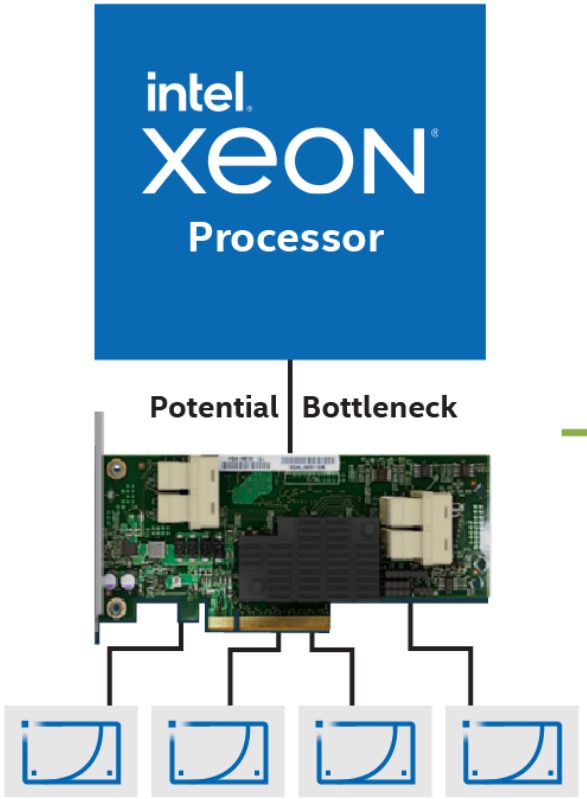

NVMe HBA Cost: $1,797

$1,298
BOM Savings

↑ Performance
↓ Latency

Integrated RAID Architecture

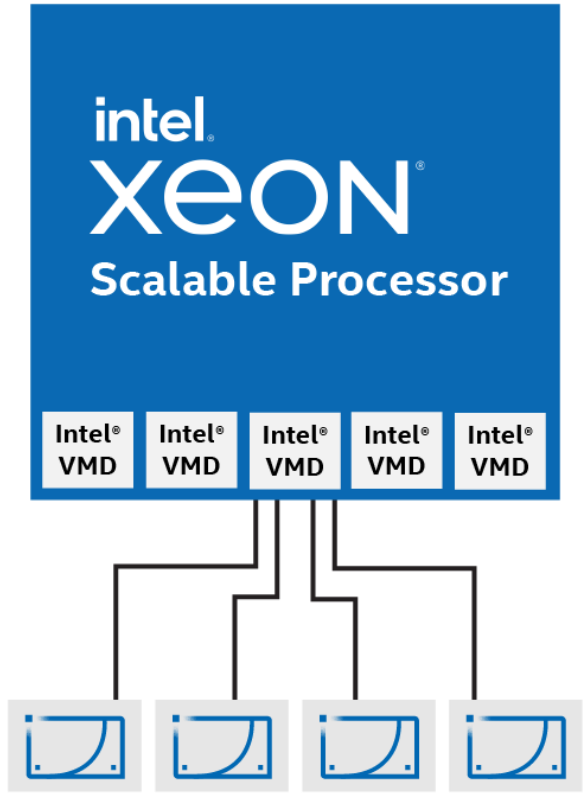

VROC RCP: $499

Broadcom MegaRAID 9560-16i - storage controller (RAID), https://www.cdw.com/product/broadcom-megaraid-9560-16i-storage-controller-raid-sata-6gb-s-sas-1/6392393
Pricing captured on 01/24/2023

See backup B for configuration details. Results may vary

How to Buy – OEM Readiness

HPE Supported Platforms	Intel® VROC SKU Description	HPE PN
ProLiant Gen10 plus: DL360, DL380, DL110 **ProLiant Gen11**: DL20, ML30, ML110, ML350, DL110, DL320, DL360, DL380, DL560 **Edgeline:** E920, E920d, E92t **Synergy:** 400, 660 **Apollo:** 2000, 6000	Intel® VROC Premium SKU	R7J57A (**FIO**[1]) R7J59AAE (**E-RTU**[2])
	Intel® VROC Standard SKU	S0E37A (**FIO**[1]) S0E38AAE (**E-RTU**[2])
ProLiant: DL20 Gen10 Plus, ML30 Gen10 Plus **Superdome:** Flex 280	SATA RAID Mode	No License Req

Supermicro Supported Platforms	Intel® VROC SKU Description	Supermicro PN
X11, B11 X12, B12 X13, B13	Intel® VROC Standard SKU	AOC-VROCSTNMOD
	Intel® VROC Premium SKU	AOC-VROCPREMOD

[1]FIO - Factory Integrated Order
[2]E-RTU - Electronic Field Upgrade

Lenovo Supported Platforms	Intel® VROC SKU Description	Lenovo PN
ThinkSystem SR630 V2, SR650 V2, SR670 V2, ST650 V2, HR630X V2, HR650X V2, SR850 V2, SR860 V2, ST50 V2, ST250 V2, SR250 V2 **ThinkEdge** SE350 V2, SE450 V2	Intel® VROC Premium SKU	4L47A39164
ThinkSystem SR630 V3, SR650 V3, ST650 V3, SR850 V3, SR860 V3, SR950 V3, HR631X V3, HR651X V3, HR860X V3, HS350X V3	Intel® VROC Premium SKU	4L47A39164
	Intel® VROC Standard SKU	4L47A83669

Intel® VROC SKU Description	DSG VROC PN
Intel® VROC (VMD NVMe RAID) Standard	951605
Intel® VROC (VMD NVMe RAID) Premium	951606

Cisco Supported Platforms	Intel® VROC SKU Description	Cisco VROC PN
UCS M5: UCSC-C240-M5xx; UCSC-C220-M5xx; UCSB-B200-M5xx; UCSB-B480-M5xx **UCS M6:** UCSC-C240-M6xx; UCSC-C220-M6xx; UCSX-210C-M6xx	Intel® SSD Only	No License Req

How to Buy – OEM Readiness

Inspur Supported Platforms	Intel® VROC SKU Description	Inspur PN
General-Purpose: NF5180M5, NF5280M5, NF5270M5, NF8480M5, NF8260M5, NF5466M5, NF5266M5, NF5468M5, TS860M5, NF5180M6, NF5280M6, NF5260M6, NF5260FM6, NF5270M6, NF8480M6, NF8260M6, NF5466M6, NF5266M6, NF5468M6, NF5488M6 **Blade:** NX5460M5 **Multi-Node:** i24M5, i48M5, i24M6, i48M6 **Rack Scale**: ORS6000S, ORS3000S, SN5161M5	Intel® VROC Premium SKU	V08902E000000000
	Intel® VROC Standard SKU	V08902F000000000

H3C Supported Platforms	Intel® VROC SKU Description	H3C PN
General-Purpose: R2700G3, R2900G3, R4700G3, R4900G3, R6700G3, R6900G3, R8900G3, R4300G3, R4400G3, R5300G3, R4700G5, R4900G5, R6900G5, R4300G5, R5500G5, R5300G5 **Blade:** B5700G3, B5800G3, B7800G3, B5700G5 **Edge**: E3200G3 **UniStor**: X10828G5 **Rack Scale**: S4703G5, S2703G5	Intel® VROC Premium SKU	0231A6R8 (FIO[1]) 0231A6TA (FUO[2])
	Intel® VROC Standard SKU	0231A6R6 (FIO[1]) 0231A6TB (FUO[2])

xFusion Supported Platforms	Intel® VROC SKU Description	xFusion PN
FusionServer: 1288H V5, 2288H V5, 1288H V6, 2288H V6, 2488H V6	Intel® VROC Premium SKU	43020237 (FIO[1])
	Intel® VROC Standard SKU	43020236 (FIO[1])

[1]FIO - Factory Integrated Order (CTO)
[2]FUO - Field Upgrade Order (BTO)

Intel® VROC Support Matrix

CPU	MS	Linux	ESXi (NVMe)	SATA	NVMe	PCH
3rd Generation Intel® Xeon® CPU	√	√	RAID1	√	√	√
4th Generation Intel® Xeon® CPU	√	√	RAID1/5	√	√	√
5th Generation Intel® Xeon® CPU	√	√	RAID 1/5	√	√	√
Granite Rapids	√	√	RAID 1/5/10		√	

OEM Enabled License Models

Software/Keyless

- Intel® On Demand**
- BIOS enables VMD/VROC
- OEM enables platform
- Field upgradable

Hardware Key

- OEM design into motherboard
- OEM purchases key from Intel
- Field upgradeable

** available with 5th Gen Intel® Xeon® Processors

Key Features

Intel® VROC Features

- Volume Management Device (VMD)
- Boot and Data RAID 0/1/5/10
- Data RAID spanning
- NVMe and SATA SSD Drives
- Hot-plug/ LED Management/ Fault Isolation
- SED Key Management
- Out-of-band management (OOB)
- VMD Direct Assign for hypervisor use case
- OS and UEFI Pre-OS supported:
  - Windows, Linux, and VMWare ESXi
- Supported and Validated by Intel

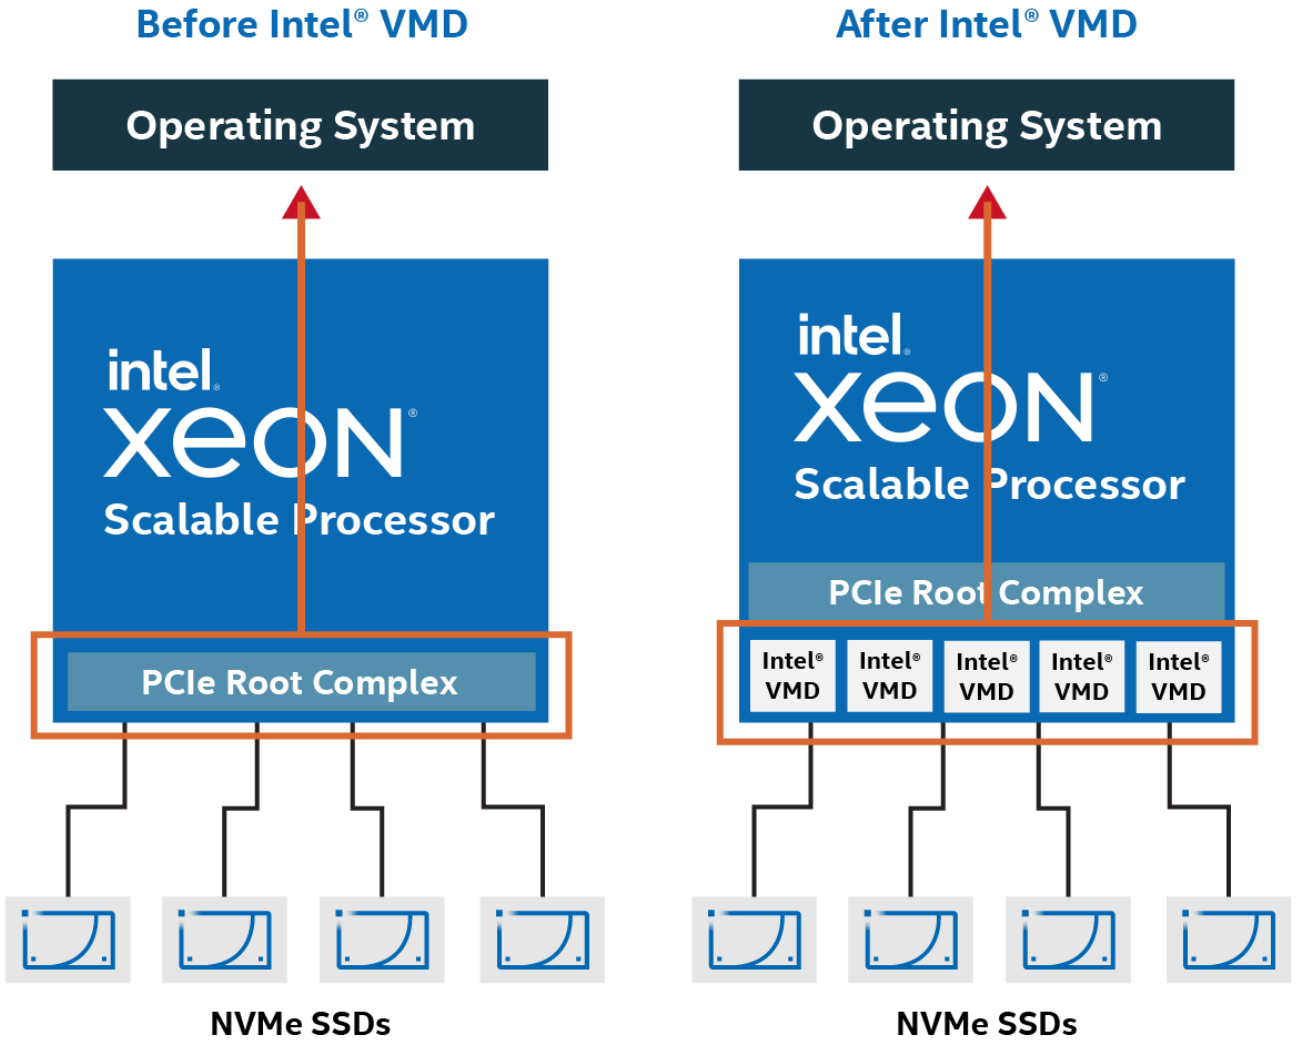

Windows and Linux: Supported Configurations

Intel VROC Configuration Guidelines

- RAID Level Support: 0, 1, 5, 10
- RAID Spanning
  - Data Volumes can span VMD Domains and CPUs
  - Boot Volumes must be within a single VMD Domain (e.g., PCH VMD)
- Max. 96 NVMe SSDs/platform
  - Up to 24 SSDs per RAID0/5 Array
  - Up to 4 SSDs per RAID10 Array
  - Up to 2 SSDs per RAID1 Array

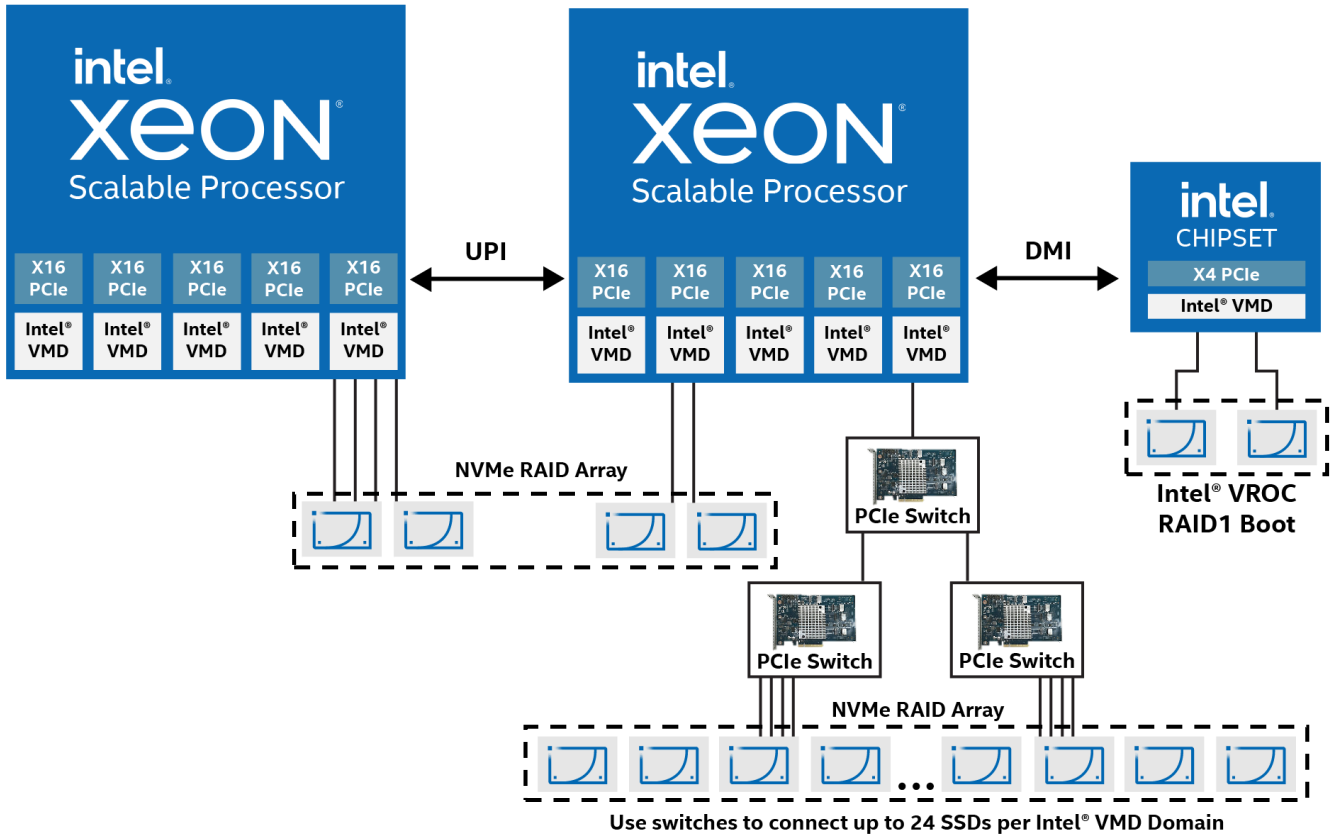

VMware ESXi – Supported Configurations

- **Intel VMD-enabled NVMe RAID1 Driver**
  - In-boxed in VMware ESXi
- VMware vSAN certified
  - If vSAN is being used: Boot volumes must be on separate VMD Domain from data devices
  - If vSAN is NOT being used: Boot volumes and data devices can share a VMD domain
- RAID Level Support: Boot and Data RAID1
  - Boot Volumes must be within a single VMD Domain (e.g., PCH VMD)
- New NVMe Firmware Management tools in ESXi
  - Allow end users to update NVMe SSD firmware inside ESXi without reboot
- Data RAID5 support coming in 2024

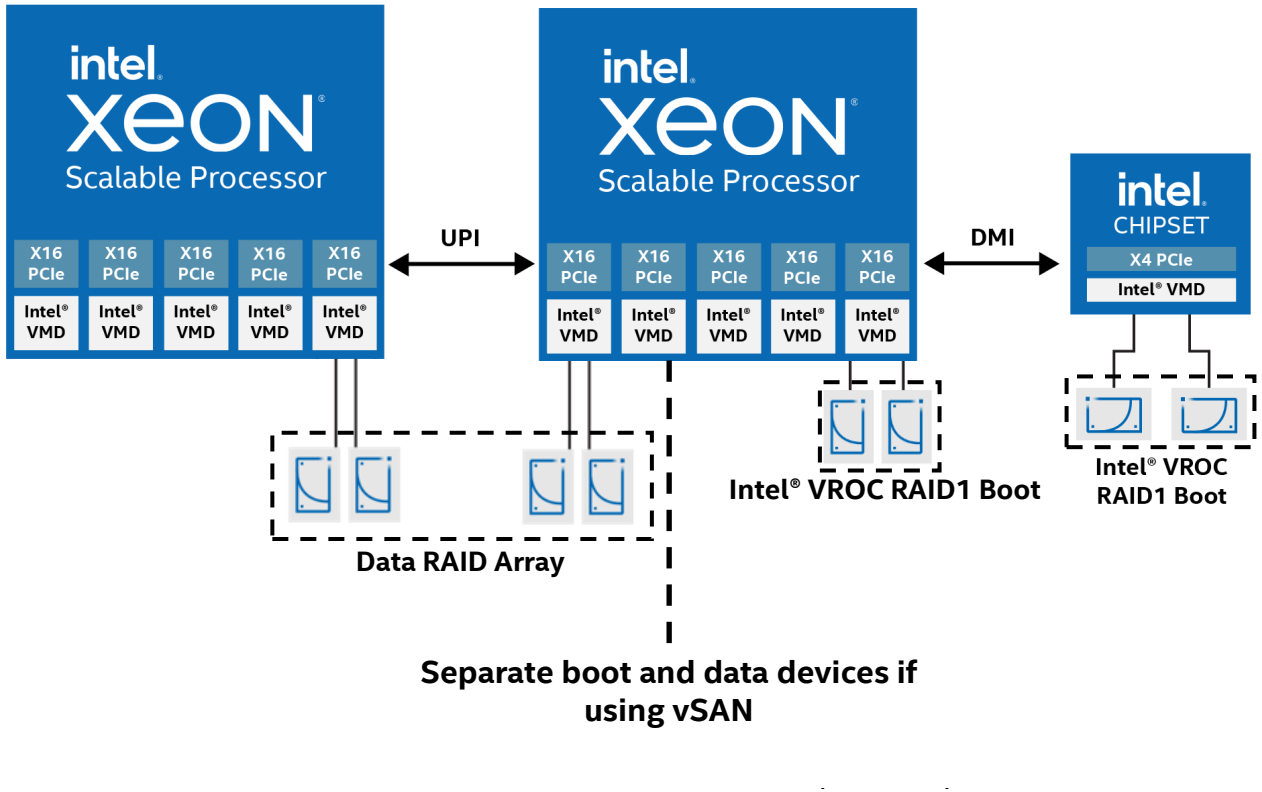

Boot RAID

Intel® VROC for Boot Volumes

- Redundant OS and High Availability
- UEFI/BIOS RAID configuration
- CPU or PCH Intel VMD domains for boot attach points
- Broad OS Support:
  - Linux, Windows, VMware ESXi*
- Form Factor Flexibility
  - M.2, U.2, E1.S, U.3
- RAID0/10 configurations for larger boot volumes

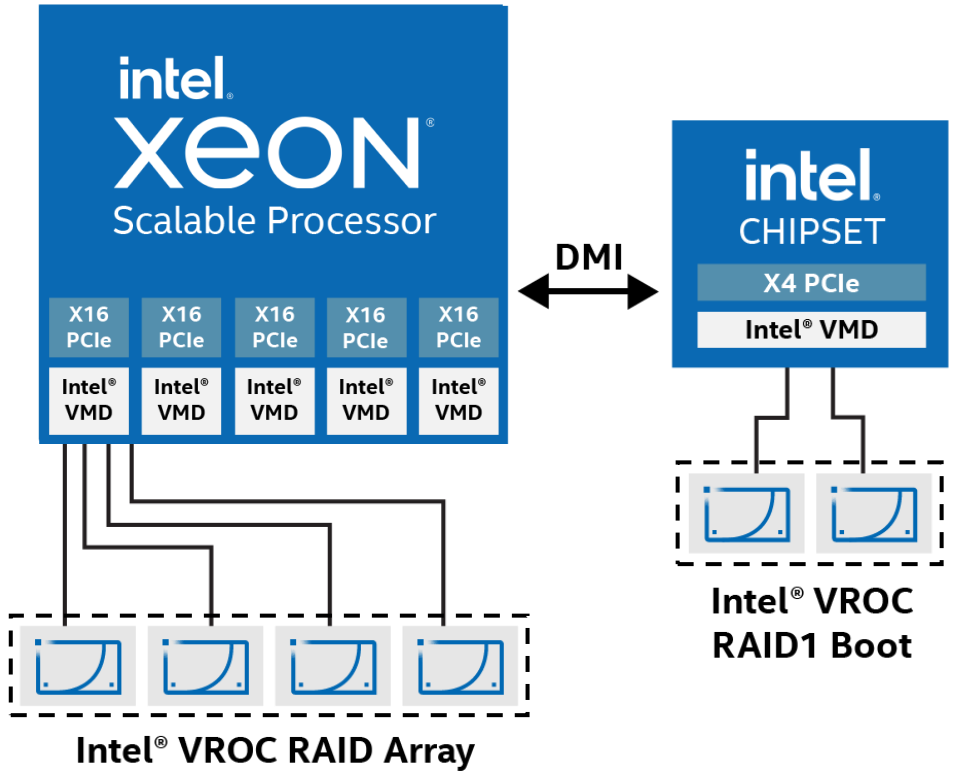

Note: New VROC license option for boot only RAID1 will be introduced in VROC 8.5
* VMWare limited to RAID1 boot

Out-of-Band Management

VROC 8.0 OOB Architecture

- **BMC**: PLDM/RDE over MCTP aware software, no VROC component necessary
- **VMD**: exposes a set of registers that constitute a mailbox
- **OOB-MSM**: responsible for MCTP message transfer through the VMD mailbox register
- **Host software**: OOB agent running on the host OS presenting an RDE device
  - **Linux**: user mode daemon accessing the VMD mailbox through sysfs executing commands via mdadm(8)
  - **Windows**: VROC VMD driver accessing the VMD mailbox passing commands to the VROC OOB agent
  - **UEFI**: additional VROC OOB driver responsible for handling OOB commands using VROC RAID and VROC VMD drivers to execute them

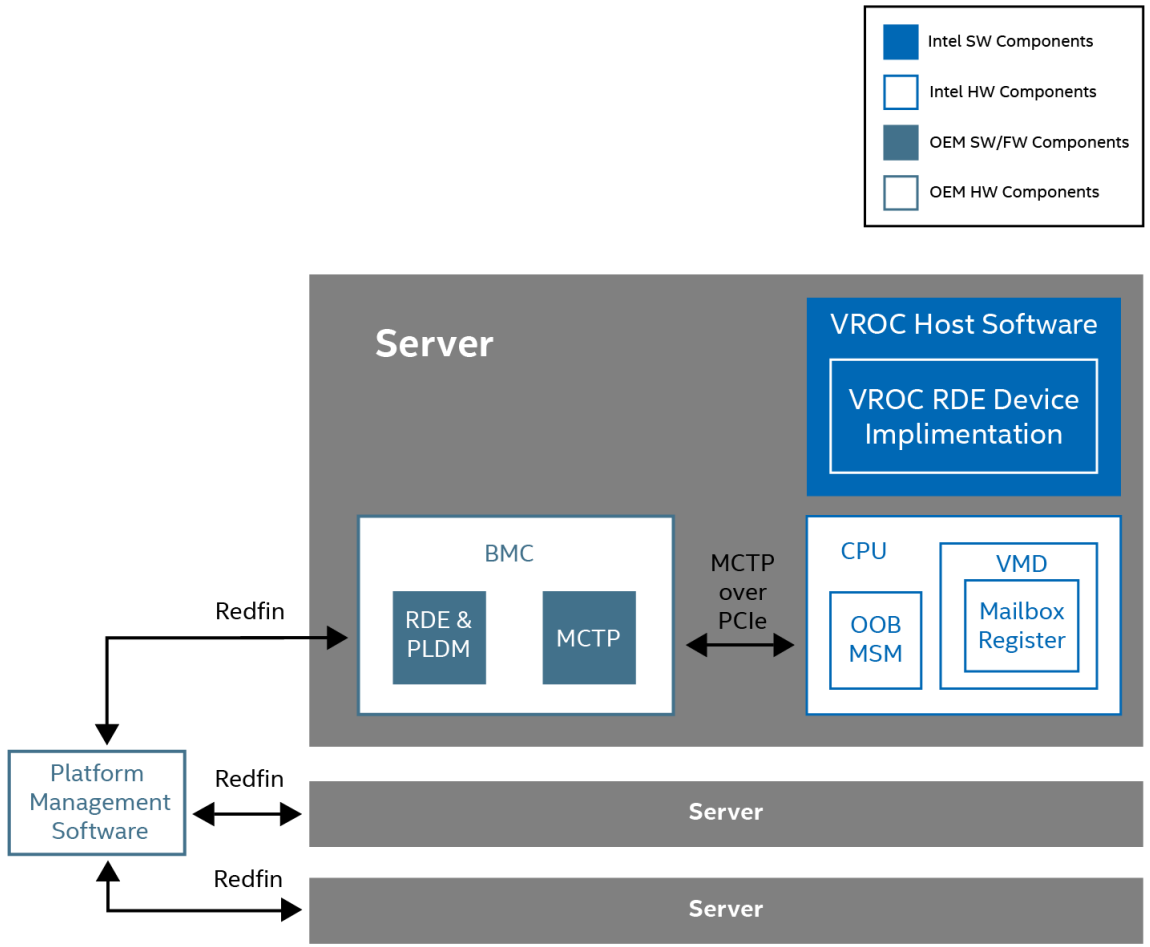

Intel® VROC SED Key Management[1]

- UEFI utility with HII interface for SED setup
- Automatic drive provisioning and unlock on system boot
- Boot from a secured RAID volume or secured single drive into any Intel VROC supported OS
- Secure sensitive data volumes with SED
- Support for multiple key managers:
  - OASIS based KMIP – industry standard KMS
  - Local Security Chip/TPM
- Intel VROC SED itself does no encryption, just key management

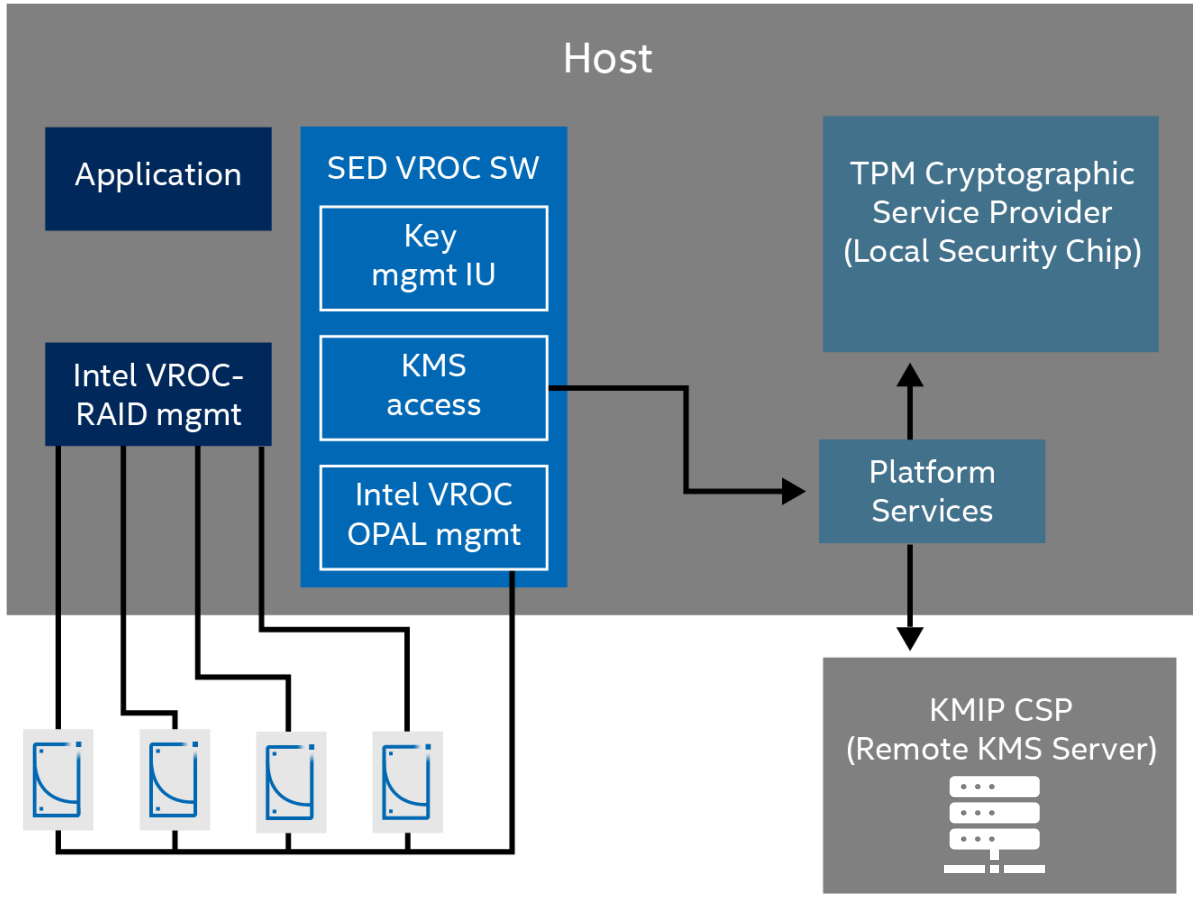

1-Requires additional driver integration at platform level. May not be available on all Intel VROC enabled platforms.
Specific functionality depends on platform provider integration into their preferred KMS

Virtualization Support

VMD Direct Assign

Intel® VMD Direct Assign Access Path

- Up to 20x performance benefit depending on Hypervisor
- VMD attached devices appear in Guest OS as physical NVMe with hot-add, hot-remove and LED management functionality
- The VMD Domain and all NVMe attached must be Direct Assigned
  - 4 Direct Attached Devices
  - Use switches to assign up to 24 devices to a VM

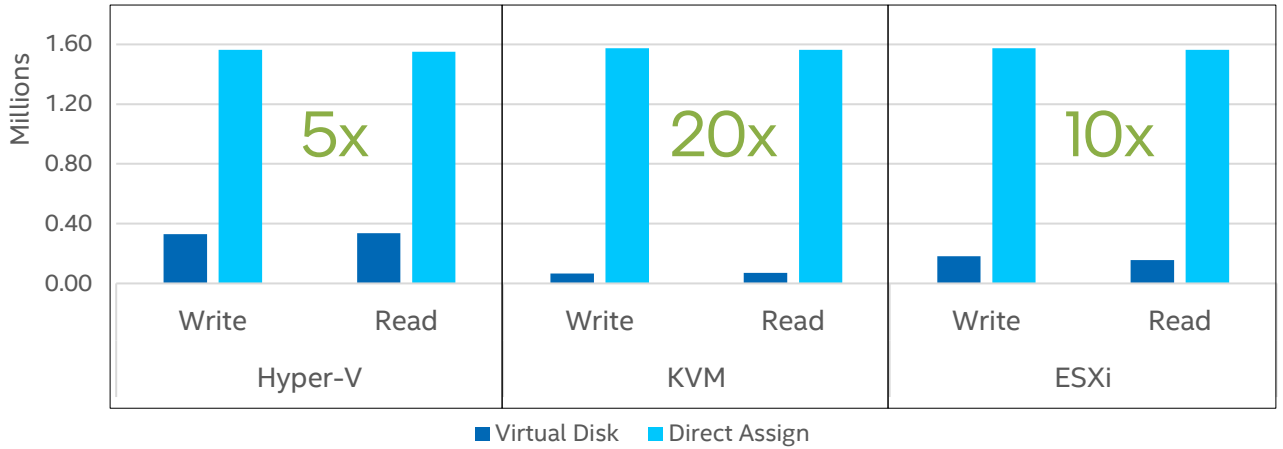

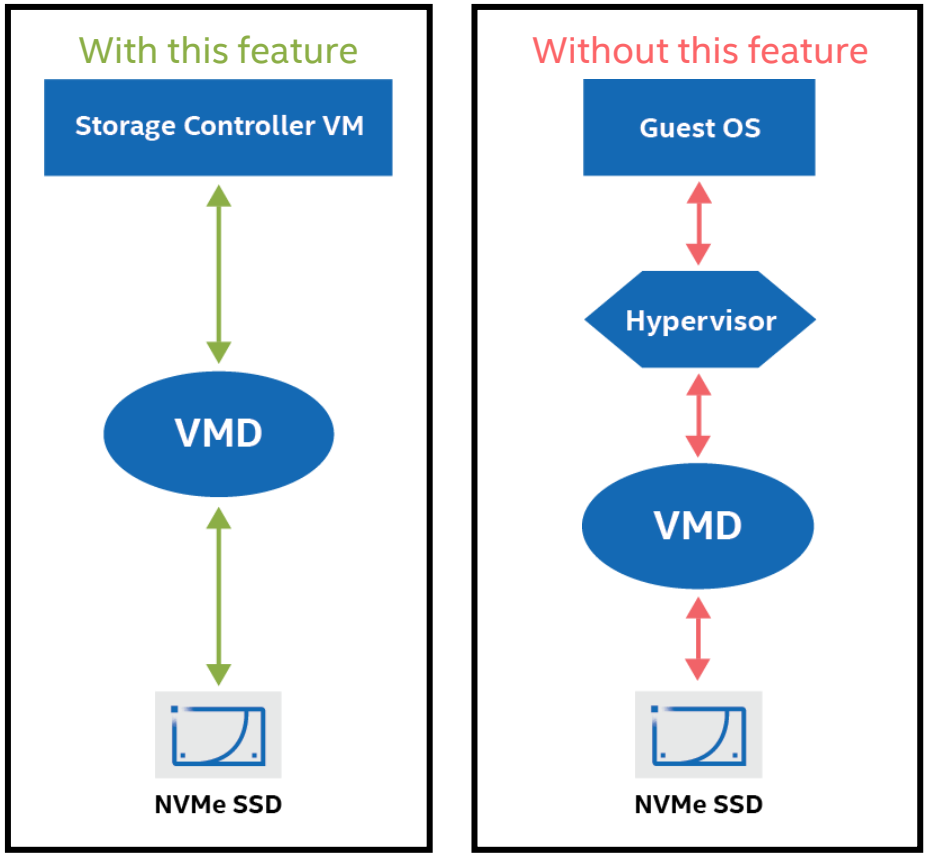

See backup A for configuration details. Results may vary

Direct Assign Ecosystem Support

- Guest OS Requires Direct Assign End-point
  - Linux In-box (RHEL 8.2)
  - Linux Out-of-box (RHEL 7.X, Ubuntu)
- Broader Hypervisor Support
  - VMWare ESXi
  - Linux KVM
  - Custom AHVs
- Intel VMD on the Platform
  - VMD1.0 on Gen 1 or 2, Intel Xeon Scalable processors requires custom patch
  - VMD2.0 on Gen 3 or 4, Intel Xeon Scalable processors supports Direct Assign Natively with Linux 8.2 and BIOS settings for VMD Direct Assign on VMD Domain

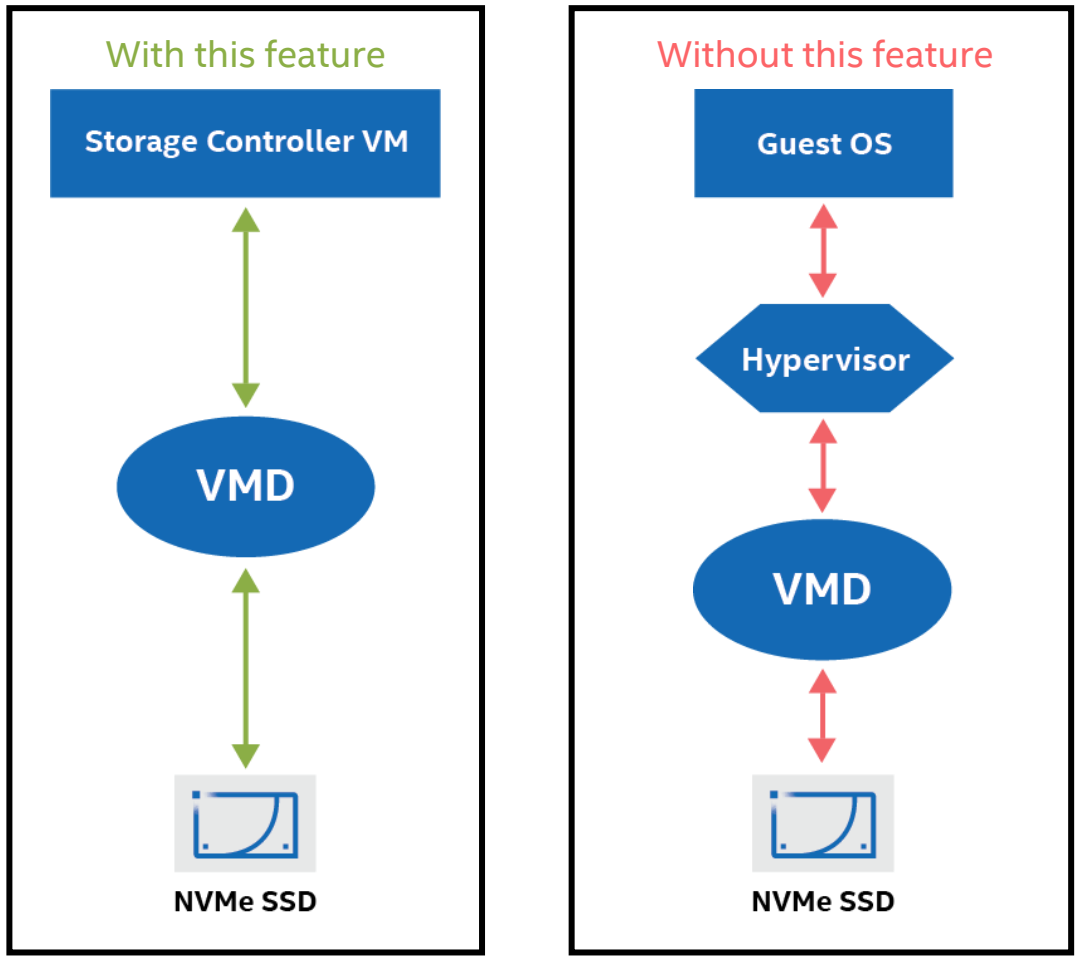

*Other names and brands may be claimed as the property of others.

Data Reliability

RAID5 with Hot Spare

- NVMe SSDs have 2 to 3 orders of magnitude higher reliability than enterprise HDDs
- Single NVMe drive Probability of Failure is less than 0.3%* over a period of 10 years; see criteria on the right…
- Automatic rebuilds enhance reliability by eliminating the need for human intervention
- RAID5 has significantly lower performance penalty than RAID6

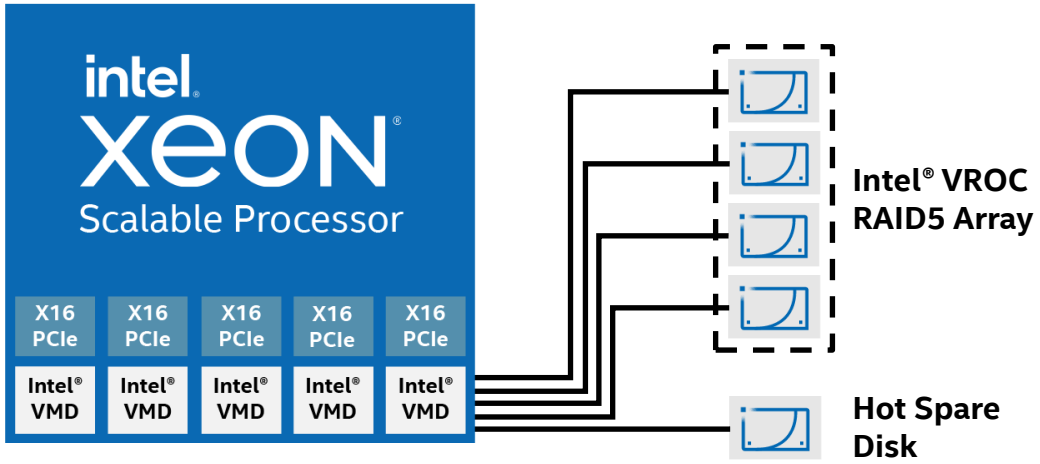

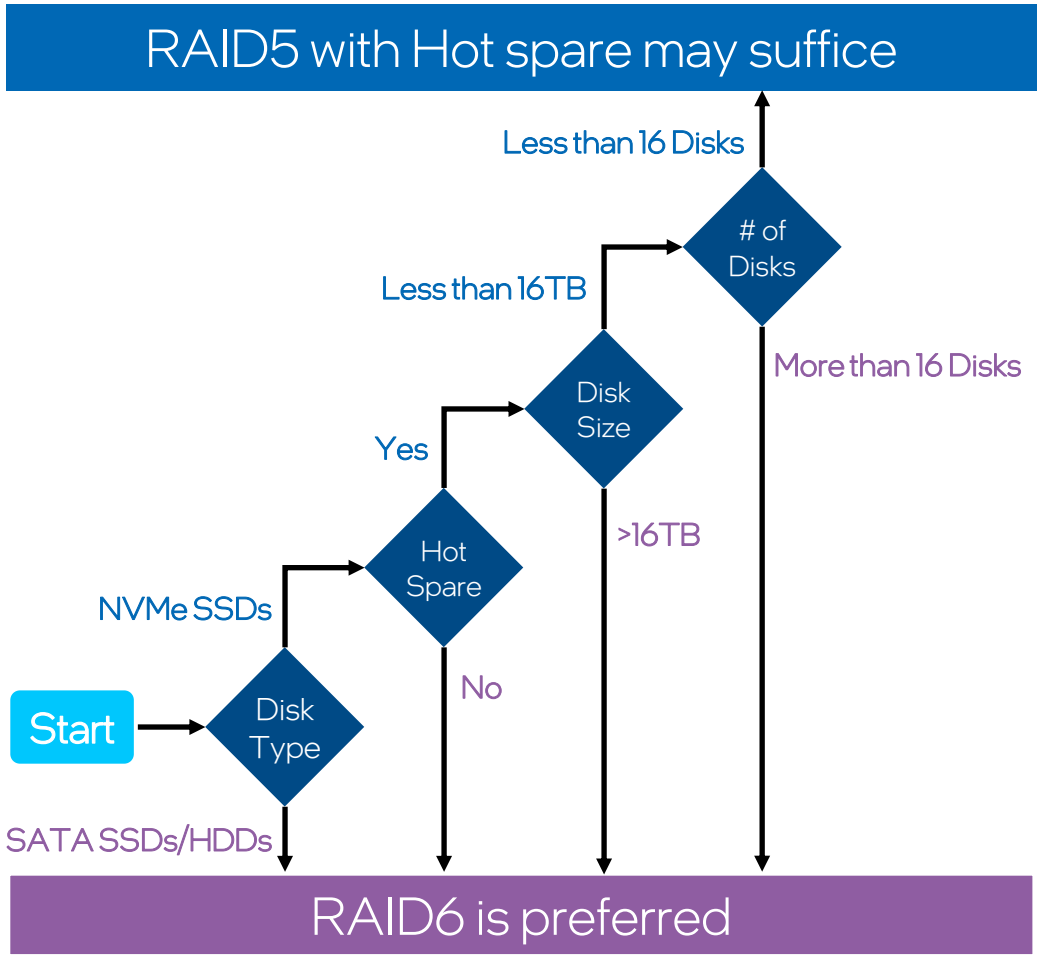

*See backup for model details. Results may vary

Performance

Storage Performance

Key Performance Indicators (KPIs)

IOPS and Bandwidth

Latency

CPU Utilization

Power

Application Performance

Intel® VROC Software Stack vs HBA

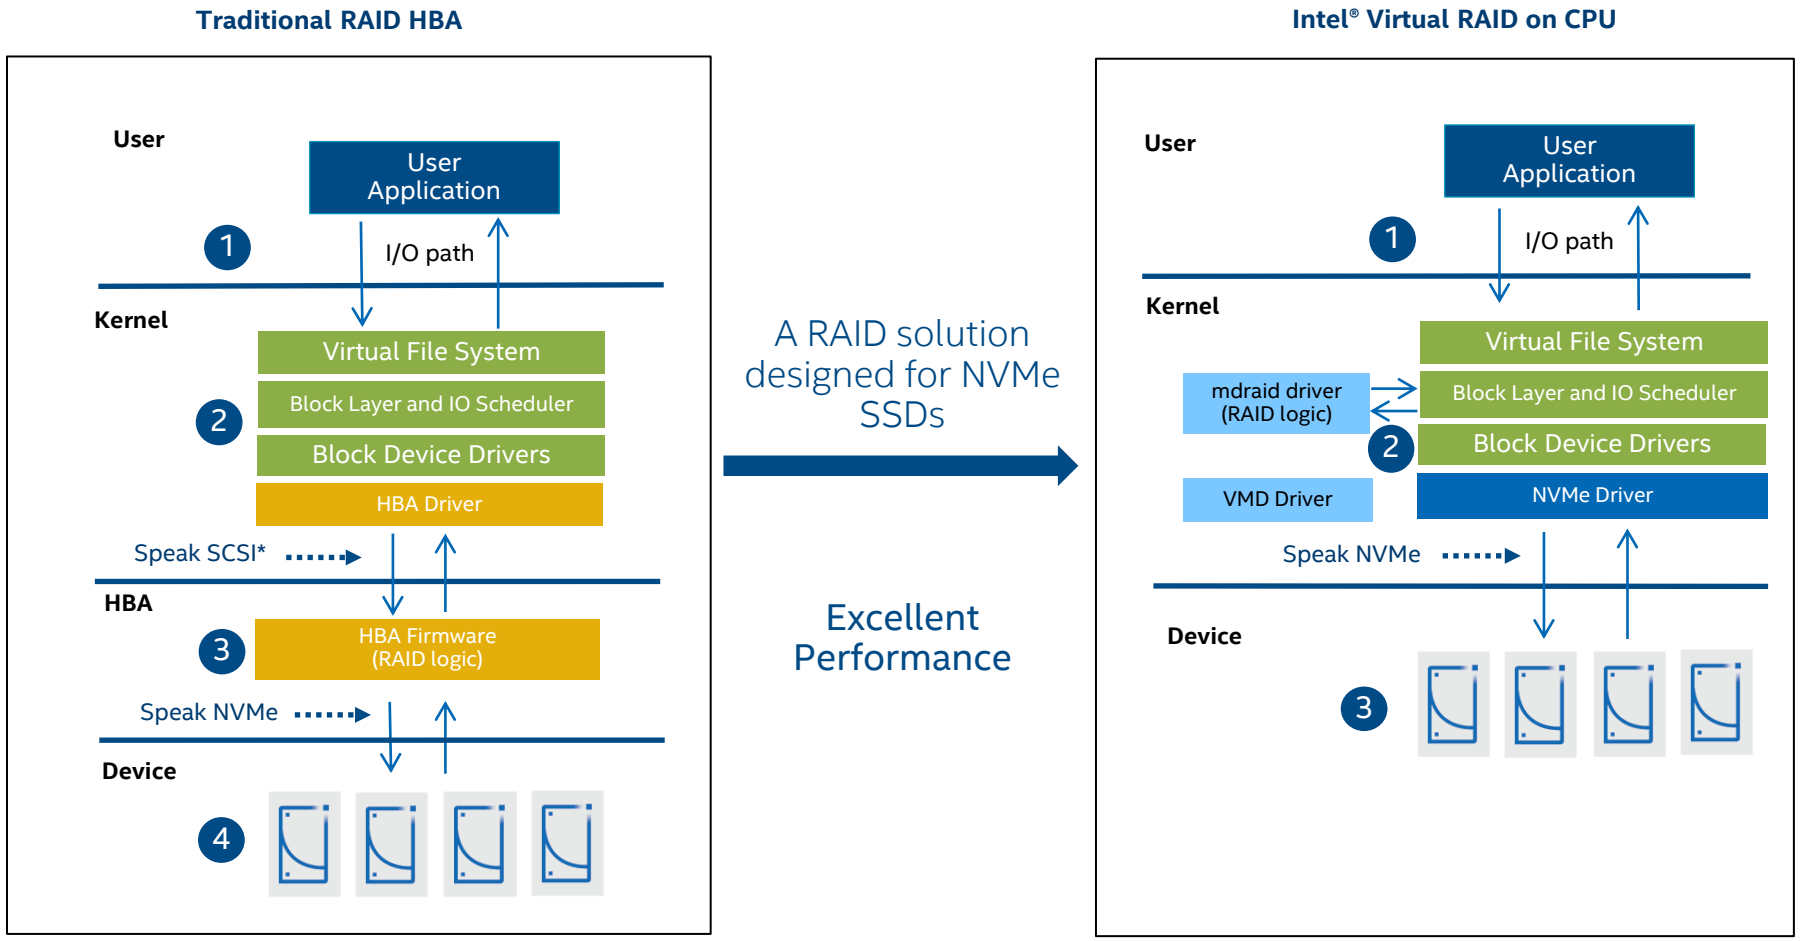

* Refers to Tri-mode HBA cards supporting SATA, SAS and NVMe interfaces

Intel® VROC Performance

Intel VROC achieves up to 5.7 million IOPS with RAID0 on mixed workloads

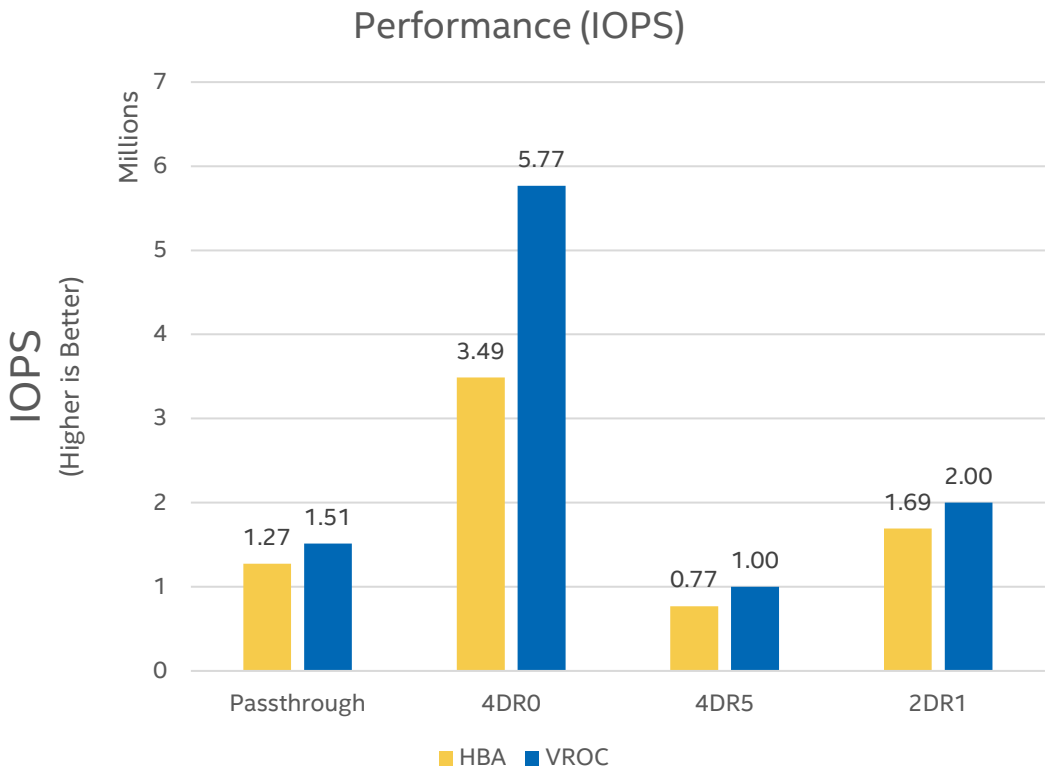

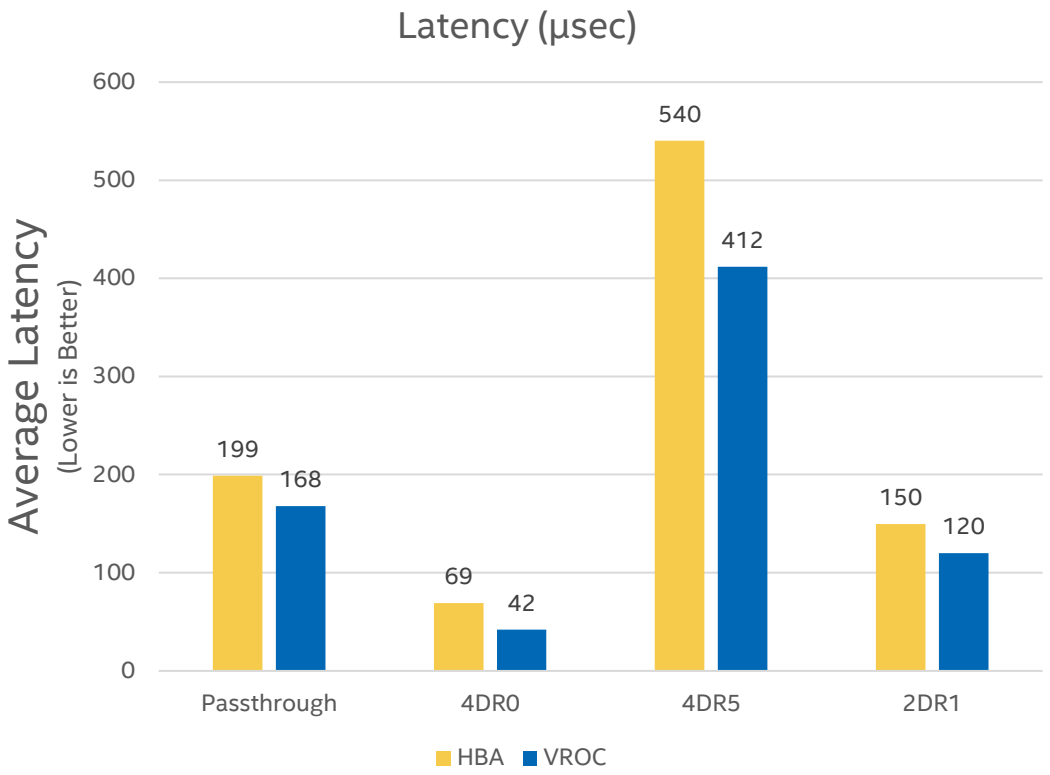

Result is based on Intel P5810X with 4KB random 70/30 R/W workload using 16 Threads &16 IODepth

See backup B for configuration details. Results may vary

CPU Utilization, Throughput & Latency Comparison

RAID0 4-drive on Intel P5810X with 4KB random 70/30 R/W workload using 16 Threads

~0.13% ↑

CPU Utilization

~65% ↑

Throughput

~39% ↓

Latency

IO Depth	CPU Utilization (100 – Idle%)			Throughput (MB/s)			Average Latency (µsec)		
	VROC	HBA	VROC-HBA	VROC	HBA	VROC/HBA	VROC	HBA	VROC/HBA
64	9.64%	9.51%	0.13%	23,126	13,971	1.65x	175	289	0.605x
128	9.64%	9.51%	0.13%	23,107	13,965	1.65x	352	582	0.604x
256	9.63%	9.51%	0.12%	23,049	13,965	1.65x	708	1,169	0.605x

See backup B for configuration details. Results may vary

Better Overall Performance

1. Intel VROC performance for all RAID levels is equal or better than RAID HBA **↑ IOPS, ↓ Latency**

2. Intel VROC implementation efficiently utilizes CPU resources where HBA is limited by bandwidth bottleneck **Improved Application Performance at Peak Workload**

3. Intel VROC is designed to **scale** with workload demand and unlocks the performance potential of NVMe SSDs **↑ IOPS/CPU Core**

See backup for configuration details. Results may vary

Enabling and Configuration

To Enable VMD Within the Intel CRB BIOS

1. Enable VMD

2. Choose the Processor

3. Choose to enable each x4 PCIe lane for VMD Config

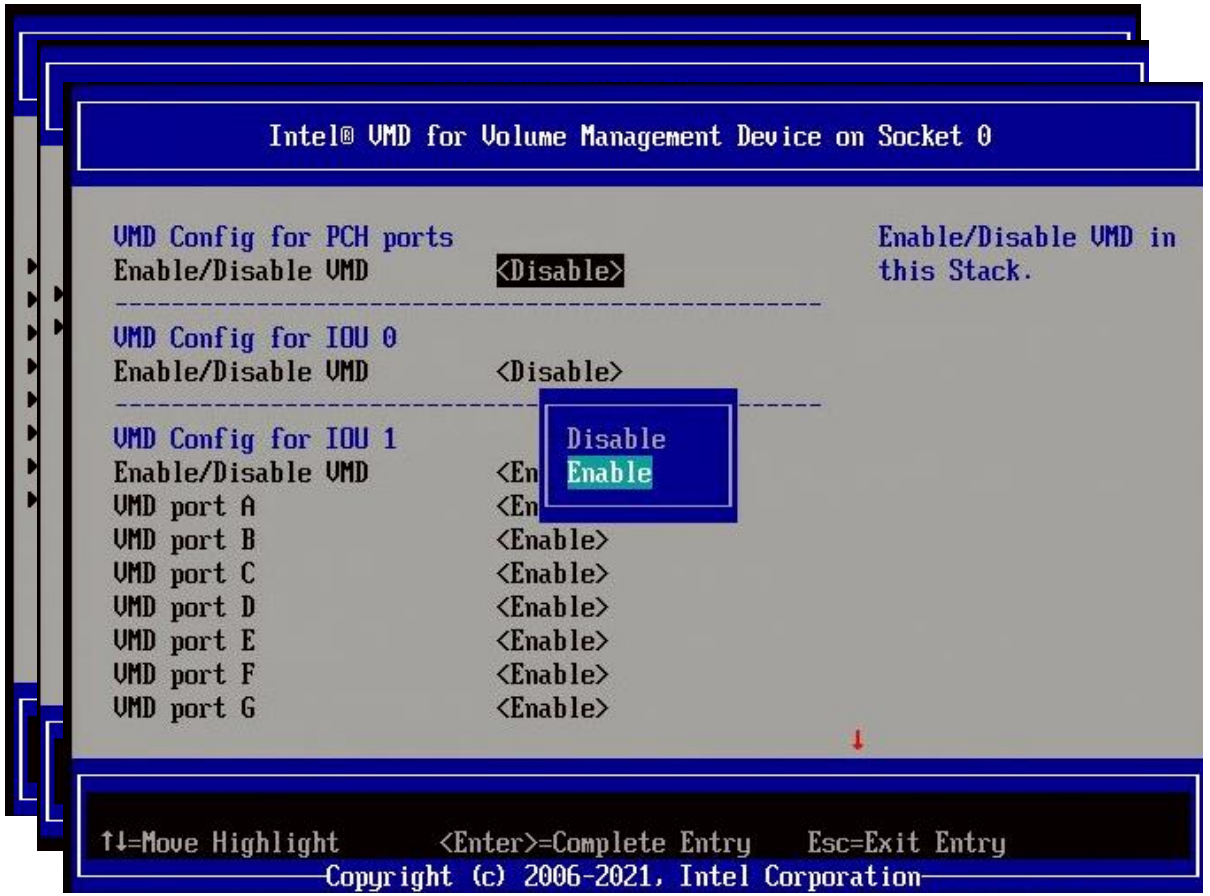

Screenshots captured on Intel® BIOS. OEM BIOS UI may look different

Configuring Intel VROC RAID Volumes from the CRB BIOS Setup Environment

1. Enter Intel® Virtual RAID on CPU Menu
2. Create RAID and View Existing RAID Volumes and NVMe Devices
3. View each device and perform actions

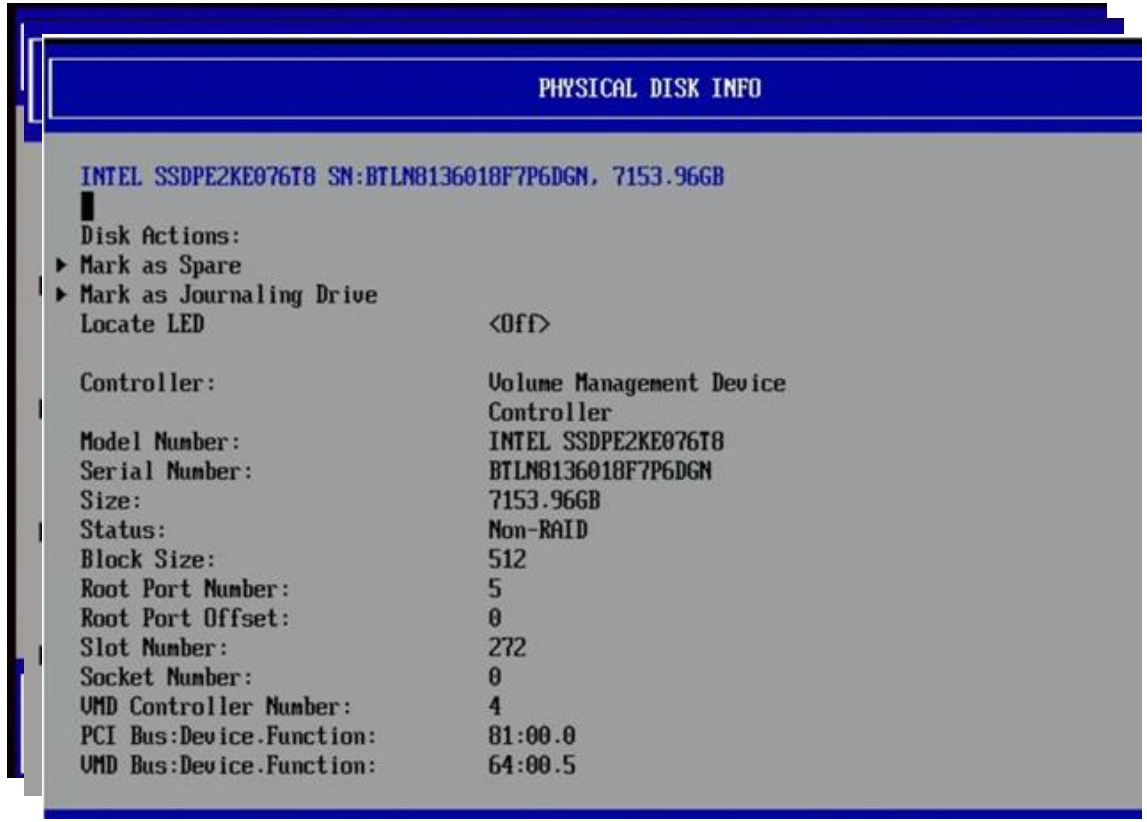

Screenshots captured on Intel® BIOS. OEM BIOS UI may look different

Configuring Intel VROC RAID Volumes in the OS Environment

Microsoft Windows Server VROC GUI

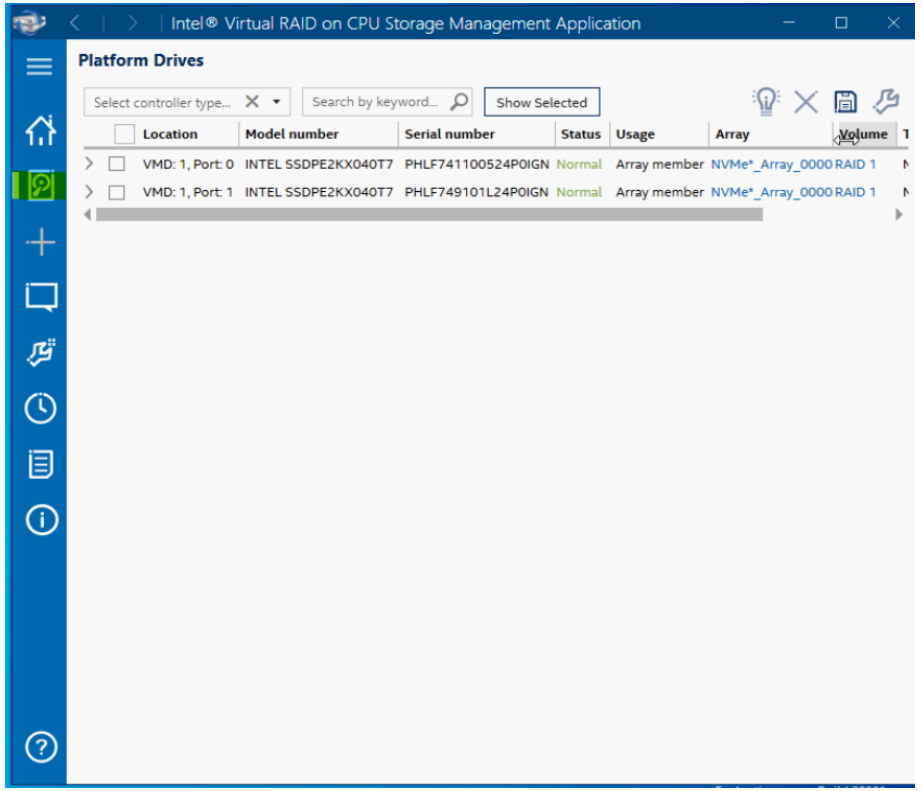

Linux Shell Intel VROC Command Line

```
# mdadm -C /dev/md/imsm0 /dev/nvme[0-1]n1 -n 2 -e imsm
# mdadm -C /dev/md/md0 /dev/md/imsm0 -n 2 -l 1
```

ESXi VMD RAID Command Line Tool

```
[root@localhost:/opt/intel/bin] ./intel-vmdr-user disklist
Disk Name: "INTEL SSDPF21Q40"   Controller: vmhba2      TargetId: 0     Serial#: "PHAL0274001E400JGN  "         State: SPARE
Disk Name: "INTEL SSDPF2KX01"   Controller: vmhba2      TargetId: 1     Serial#: "PHAX107400261P9BGN  "         State: PASSTHROUGH
Disk Name: "INTEL SSDPF21Q03"   Controller: vmhba2      TargetId: 2     Serial#: "PHAL1046000P3P2GGN  "         State: MEMBER
Disk Name: "INTEL SSDPF2KE12"   Controller: vmhba2      TargetId: 3     Serial#: "PHAX1346001L15PFGN  "         State: PASSTHROUGH
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin] ./intel-vmdr-user getled vmhba2-1 -d 1
get_led Called
Targeting Disk:1
LED state:off
[root@localhost:/opt/intel/bin] ./intel-vmdr-user setled vmhba2-1 -d 1 -l identify
Targeting Disk:1 with LED:identify
Request to Set LED on disk has completed.
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin] ./intel-vmdr-user getled vmhba2-1 -d 1
get_led Called
Targeting Disk:1
LED state:identify
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin] ./intel-vmdr-user setled vmhba2-1 -d 1 -l off
Targeting Disk:1 with LED:off
Request to Set LED on disk has completed.
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin]
[root@localhost:/opt/intel/bin] ./intel-vmdr-user getled vmhba2-1 -d 1
get_led Called
Targeting Disk:1
LED state:off
[root@localhost:/opt/intel/bin]
```

Intel® VROC Resources and Support

Intel® VROC Resources and Support Docs

- **Intel.com/VROC** for:
  - Product Brief
  - FAQ
  - Supported OS/HW Information
  - Related Links
- **Intel® VROC Support Page for:**
  - User Guides
  - Tech Briefs
  - Performance Documentation
- **Technical Docs on RDC**
- **Support available via IPS**

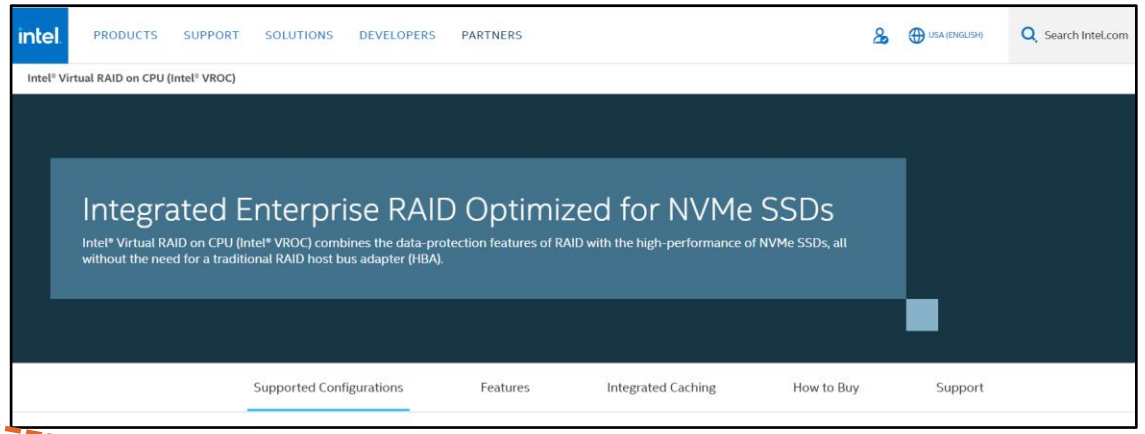

Intel® Virtual RAID on CPU Website

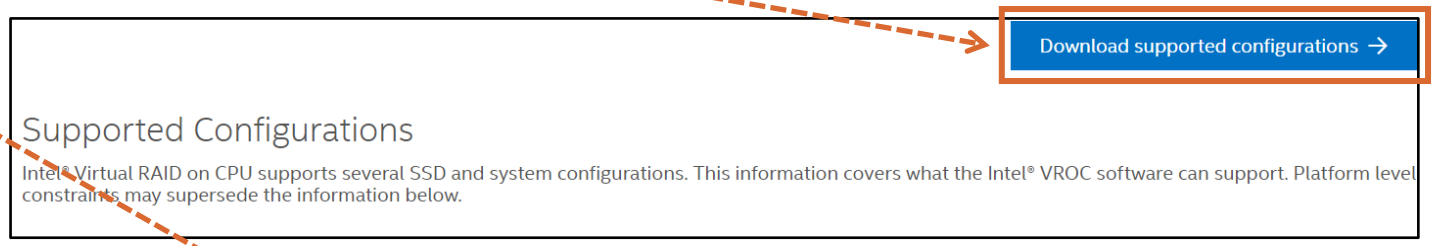

Download support

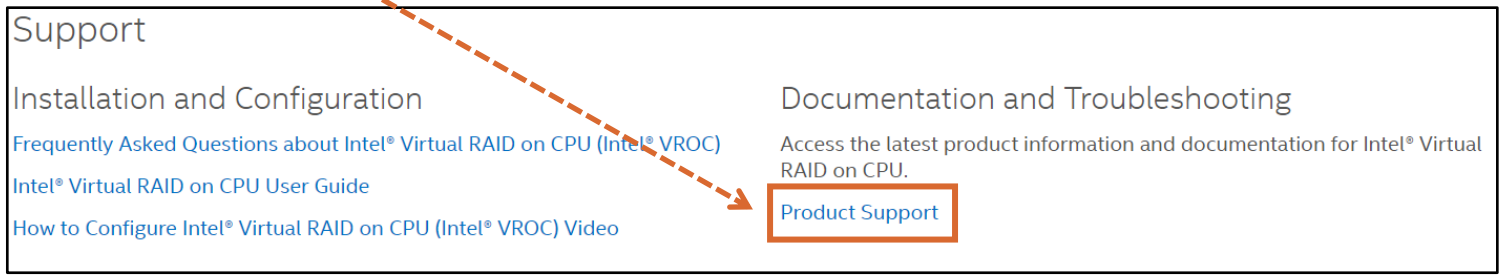

Get product support

Backup

CPU Utilization, Throughput & Latency Comparison

RAID5 4-drive on Intel P5810X with 4KB random 70/30 R/W workload using 16 Threads

~4.34% ↑

CPU Utilization

~48% ↑

Throughput

~33% ↓

Latency

IO Depth	CPU Utilization (100 – Idle%)			Throughput (MB/s)			Average Latency (µsec)		
	VROC	HBA	VROC-HBA	VROC	HBA	VROC/HBA	VROC	HBA	VROC/HBA
64	6.34%	1.48%	4.86%	3,943	2,637	1.48x	1,714	2,575	0.66x
128	6.46%	1.49%	4.97%	3,913	2,636	1.48x	3,471	5,166	0.67x
256	6.27%	3.09%	3.18%	3,870	2,665	1.48x	7,037	10,196	0.67x

See backup B for configuration details. Results may vary

Backup A – Configuration Details

Performance results are based on testing by Intel as of October 28, 2022 and may not reflect all publicly available updates. Results may vary.

- **Platform**: ArcherCity_012, 2x SVR SPR CPU XCCSP D0 QY0E CPU, SVR EBG PCH B0 QXRH PCH, BKC: BKC#44_AR, BIOS EGSDCRB.SYS.OR.64.2021.47.2.02.0319.0_SPR_EBG_SPS_Production__RSS_BIOS_1.0.0.644, PreOS 8.0.0.1196, RAM: 8 GB
- **OS:** Linux RHEL 8.4_2054.01.0.2101, VMD_LINUX_iavmd_1.0.0.1590 Driver
- **Additional kernel parameters:** pci=pcie_bus_perf, initcall_blacklist=vmd_drv_init,
- **Storage:** Alderstream drive connected to HSBP1_Slot1-4, PHAL029300AE1P6MGN size=1,6 TB model=INTEL SSDPF21Q016TB model name=P5800X FW=L0310300
- **BIOS settings:** EDKII Menu > Miscellaneous Configuration >> Fan PWM Offset [100], > Console Redirection Configuration >> Console Redirection <Enable>, > Socket Configuration >> Processor Configuration >>> Hyper-Threading [ALL] <Enable>, >> IIO Configuration >>> Intel? VT for Directed I/O (VT-d) >>>> Intel? VT for Directed I/O <Enable>, >> IIO Configuration >>> IIO Global Performance Tuning >>>> Performance Tuning Mode <Performance Enable Mode>, >> Advanced Power Management Configuration >>> CPU P State Control >>>> SpeedStep (Pstates) <Enable>, >>>> Turbo Mode <Enable>, >>>> Energy Efficient Turbo <Disable>, >>> Hardware PM State Control >>>> Hardware P-States <Native Mode>, >>> Package C State Control >>>> Package C State <Auto>, >>> CPU - Advanced PM Tuning >>>> Energy Perf BIAS >>>>> Workload Configuration <I/O sensitive>

Backup B – Configuration Details

Performance results are based on testing by Intel as of March 21, 2023 and may not reflect all publicly available updates. Results may vary.

- **Platform**: Intel Beta Fox Creek Pass M50FCP2SBSTD (chassis M50FCP2UR208BPP), 2 x Intel® Xeon® Platinum 8468H @ 2.1GHz (XCC CPUs, QDF: Q242, Stepping: E5) (48 cores each) (EagleStream - Sapphire Rapids), 256GB RAM (16 x 16GB Micron MTC10F1084S1RC48BAW 4800 MT/s DDR5 Synchronous Registered (Buffered) DIMMs), BIOS Version: SE5C7411.86B.8805.D02.2209220021 (Microcode revision: 0x2b000081), BIOS Release Date: 09/22/2022, BMC version: 1.27-0-gfedbbf-3cc10000, ME version: 06.00.03.0248, FRU version: 0.02, CPLD version: 2.0
- **BIOS Settings**: SpeedStep(Enabled), Turbo(Enabled), ProcessorC6(Enabled), PackageC-State(C0/C1 State), CPU_PowerAndPerformancePolicy(Performance), HardwareP-States(NativeMode), WorkloadConfiguration(I/O Sensitive), Hyperthreading enabled
- **Storage**: 4 x 400GB Intel Optane P5810X PCIe Gen4 U.2 SSDs (Model: SSDPF21Q400GA, Firmware: L0310351) connected to backplane which is connected via Broadcom SlimSAS to SlimSAS connections that connect to an Intel RS3P4TF160F RAID controller card on PCIe slot 1 on Riser card 2 on the 2nd CPU (NUMA Node 1). OS on 1 of the 2 x 118GB Intel Optane P1600X M.2 SSDs (Model: SSDPEK1A118GA, Firmware: U5110550) connected to M.2 sockets on the motherboard on the 1st CPU (NUMA Node 0), CPU affinitized on 2nd CPU (NUMA Node 1)
- **RAID Controller**: Intel RS3P4TF160F (x8) (equivalent to Broadcom MegaRAID 9560-16i) card with Broadcom firmware, Firmware Package Build = 52.22.0-4544, Firmware Version = 5.220.02-3691, Driver Version = 07.721.02.00, CLI Version = 007.1912.0000.0000 Nov 23, 2021, Added "scsi_mod.use_blk_mq=y" to grub boot option for maximum throughput on the Broadcom card, When creating RAID volumes "pdcache=on, Write-Back, No Read Ahead, Direct I/O". **OR** Intel® VROC PreOS Version: 8.0.0.1336, mdadm version: mdadm - v4.2-rc2 - 2021-08-02, Installed kmod-iavmd-1.0.0.1600-rhel_85.x86_64, "initcall_blacklist=vmd_drv_init" was added to grub boot option which disables inbox VMD and enables the kmod-iavmd driver, Added "pci=pcie_bus_perf" to grub boot option which sets MaxPayload to the maximum for each of the NVMe devices
- **OS**: Red Hat Enterprise Linux Server 8.5, Kernel: 4.18.0-348.el8.x86_64
- **RAID Configurations**: 4-Disk RAID0 with Intel VROC and Intel RS3P4TF160F
- **FIO version**: 3.30 (fio config files will not include the "iodepth_batch_complete_min" parameter for all testing)

RAID5 with Hot Spare – Reliability Model

Intel modeled RAID5 and RAID6 Mean Time to Data Loss (MTTDL) using Continuous Time Markov Chains (CTMC).

MTTF Formulae for RAID

$MTTF_{RAID5} = MTTF^2 / (MTTR * N * (N-1))$
$MTTF_{RAID6} = MTTF^3 / (MTTR^2 * N * (N-1) * (N-2))$
Where MTTF is the Mean Time to Failure of the individual disk in a RAID array of N disks including parity disks.

MTTF Formulae for Disk Failure(s) Followed by an Uber Event

$MTTF_{UBER} = MTTF_{RAIDX}/P_{FAILURE\_DURING\_REBUILD}$
$P_{FAILURE\_DURING\_REBUILD} = 1 – (1 – UBER)^{BITS\_PER\_DISK*REMAINING\_DISKS}$

MTTDL

RAID MTTDL is the harmonic sum of the above two drivers

- $MTTDL_{RAID} = (MTTF_{RAIDX}^{-1} + MTTF_{UBER}^{-1})^{-1}$

Reliability Pass/Fail Criteria

We assume arrivals of failures to be a Poisson process with constant failure rates.
So, probability of failure over time t is denoted by p(t) and calculated as follows:

- p(t) = 1−e(−t/MTTDL)p(t) = 1−e−t/MTTDL

Model ran in R to calculate probability of failure

Intel® VROC Cost

	RAID HBA[1]	Intel® VROC (Premium SKU)
RAID Cost	$1797	$499

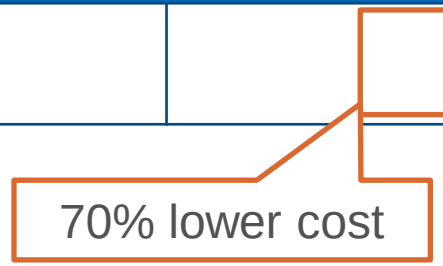

1- Broadcom MegaRAID 9560-16i - storage controller (RAID), https://www.cdw.com/product/broadcom-megaraid-9560-16i-storage-controller-raid-sata-6gb-s-sas-1/6392393
Pricing captured on 01/24/2023

Power efficiency

	HBA	VROC
IOPS	3,487,507	5,793,256
Power (W)	825	847
IOPS/Watt	4,228	6,836

62% better Power Efficiency


intel