

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.

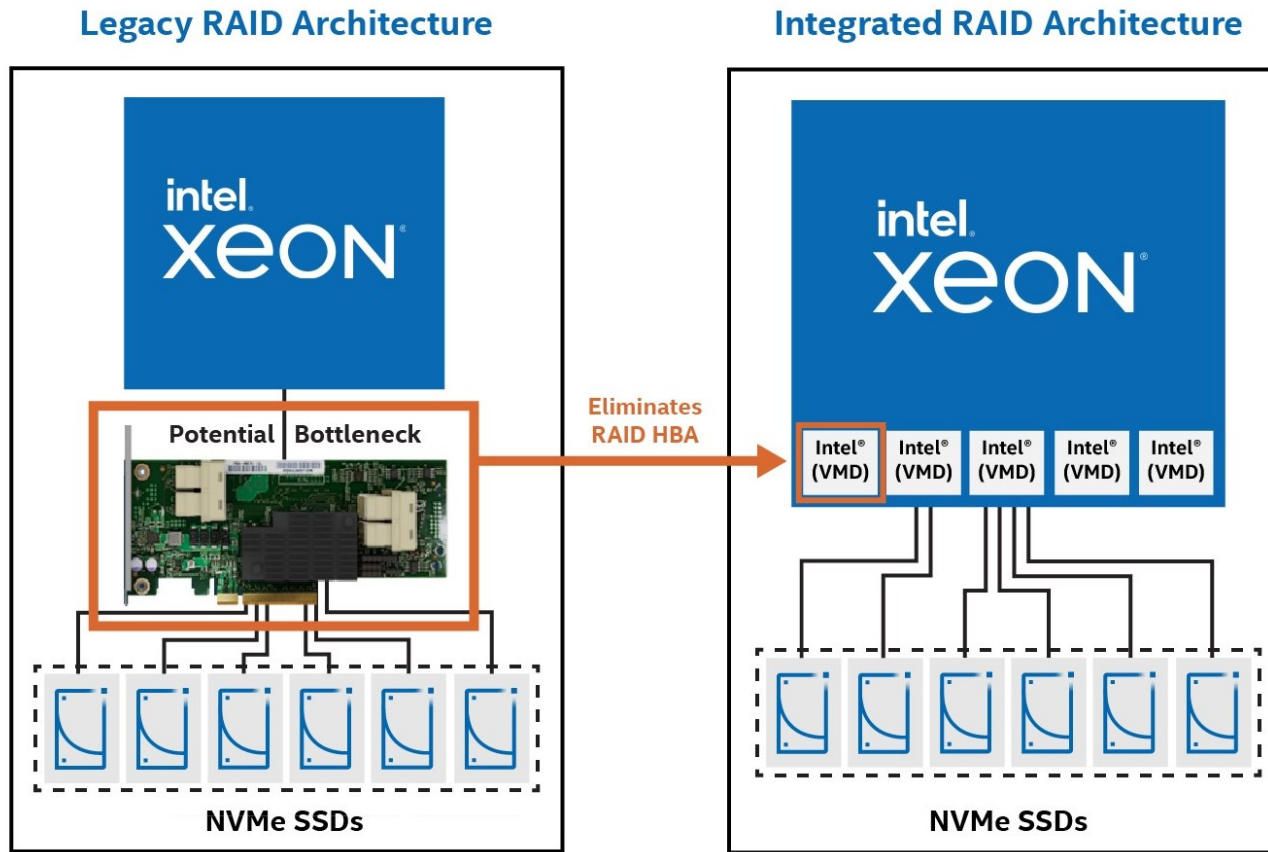
Intel technologies may require enabled hardware, software or service activation

© 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.

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

Intel® Virtual RAID on CPU

5th Gen Intel® Xeon® Processors with **Integrated RAID**



Performance
Up to **165%** increase (IOPS)
Latency reduced by **62%**



Cost
Up to **60%** lower



Power Efficiency
Up to **155%** better

See backup B for configuration details. Results may vary

Intel® VROC Product Family

Features		Pass-thru (Included with PCH)	Standard SKU	Premium SKU	RAID1 Only SKU **
Intel® VMD	Hot-plug	✓	✓	✓	✓
	Fault Isolation	✓	✓	✓	✓
	LED Management	✓	✓	✓	✓
NVMe 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

4 th Gen Intel® Xeon® Processors	5 th Gen Intel® Xeon® Processors	Birch Stream / Kaseyville (In Planning)
<ul style="list-style-type: none"> ❑ 5 VMDs ❑ Out-Of-Band management support ❑ Firmware Management Tools for VMWare ESXi ❑ SED Local Key Management ❑ Pre-boot DMA protection in PreOS environment ❑ NVMe secure erase via OOB 	<ul style="list-style-type: none"> ❑ 5 VMDs ❑ RAID1 Only SKU ❑ Local key management support (TPM) for SED drives ❑ Intel® On Demand licensing support ❑ VROC OOB integration with AMI BMC ❑ Data RAID10 volume with more than 4 drives on Linux ❑ NVMe Secure Erase in PreOS environment ❑ Enhanced cryptographic algorithm for VROC SED 	<ul style="list-style-type: none"> ❑ 6 VMDs ❑ Data RAID 0 and 5 support for VMWare ESXi ❑ Montage IOH support (SATA) ❑ Performance enhancements ❑ Debian OS Support ❑ NPEM support for CPU-attached SSD ❑ OOB for NVMe SSDs behind VMD via MCTP
<p>OS: Windows, VMWare ESXi, Linux</p>		

Intel® VROC is Integrated RAID

RAID Features	HW RAID	SW RAID	VROC	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
LED Management / Hot-plug	✓	x	✓	Hot-plug and LED Management through Intel® VMD
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
Reliability	x	✓	✓	Removes HBA single point-of-failure, less HW. Supports auto rebuild on spare devices
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



↑60% cost saving (vs HBA)



Improves NVMe performance ↑165%



↑155% power efficiency and Improves reliability

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, HCI (vSAN)
- Node and system data

DC / Edge


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

Telco/Embedded

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

DB Application

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



VROC & VMD Usage Scenarios

Application Server with Network Storage

Local storage for boot, network storage for data

- **Data:** Network Storage System
- **Boot:** VROC/VMD

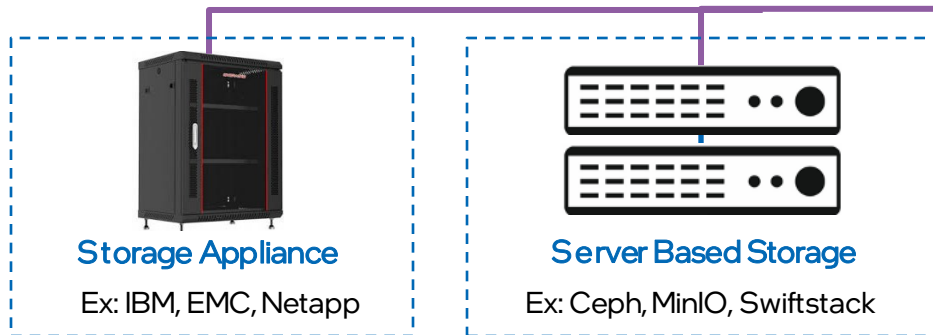


Network Storage System

Wide variety of options providing a wide variety of features.

Examples: SDS, NAS/SAN, FC, NVMe-oF, iSCSI, etc.


- **Data:** VMD with System Storage SW
- **Boot:** VROC/VMD



Hyperconverged (HCI)

Special case of SBS/SDS with the application and storage on the same system. Ex: Nutanix, vSAN, Hyperflex


- **Data:** VMD with HCI SW
- **Boot:** VROC/VMD



Direct Attached Storage

Local storage for both boot and data

- **Data:** VROC/VMD
- **Boot:** VROC/VMD



Intel® VROC Value Proposition!

Problem

- RAID cards create a performance 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) Increased reliability

Benefits

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

Competitive Advantage

- Differentiator from AMD
- DCAI Affinity
- Differentiator from Dell


Boost Customer Workloads

- 165% better performance
- 60% cost savings
- 155% 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



Integrated RAID Architecture

\$758
BOM Savings

↑ Performance
↓ Latency



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 10/09/2023

See backup B for configuration details. Results may vary

Intel® VROC Support Matrix

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

How to Buy – OEM Readiness

Dell Supported Platforms	Intel® VROC SKU Description	Dell PN
Precision Workstation: 5860 Tower, 7960 Tower, 7960 Rack Servers:	Intel® VROC Premium SKU	780-BCVC
	Intel® VROC Standard SKU	780-BCIP
HCI Appliance: XC660-10, XC760-24, XC760-24N, XC660-12N, XC760xa-6N, XC750-24, XC750-16N, XC650-10, XC650-10N	VMD Only	VMD Only

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
ThinkStation P5, P7, PX	Intel® VROC Premium SKU	SC50A13986
	Intel® VROC Standard SKU	SC50A13985
ThinkAgile HX650 V3, HX630 V3, HX3330/1, HX7530/1	VMD Only	VMD Only

HPE Supported Platforms	Intel® VROC SKU Description	HPE PN
ProLiant Gen10plus: DL360, DL380, DL110 ProLiant Gen11: DL20, ML30, ML110, ML350, DL110, DL320, DL360, DL380, DL560 Edgeline: e920, e920d, e920t, e930t Synergy: 480 Altera: 4110, 4120, 4140	Intel® VROC Premium SKU	R7J57A (FIO ¹) R7J59AAE (E-RTU ²)
	Intel® VROC RAID1 Only SKU	S3Q19A (FIO ¹) S3Q39AAE (E-RTU ²)
HPC: XD220v (Hardware Key)	Intel® VROC Standard SKU	P51777-B21 (FIO ¹)
ProLiant: DL20 Gen10 Plus, ML30 Gen10 Plus Superdome: Flex 280	SATA RAID Mode	No License Req
ProLiant Gen10plus DX360 ProLiant Gen11 DX360, DX4120, DX380a, DX560	VMD Only	VMD Only

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

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

How to Buy – OEM Readiness

Cisco Supported Platforms	Intel® VROC SKU Description	Cisco VROCPN
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

Mitac Supported Platforms	Intel® VROC SKU Description	Mitac PN
M50FCPIUR204 M50FCPIUR212 M50FCP2UR208 M50FCP2UR312	Intel® VROC Premium SKU	5262C4721734 (FIO)
M50CYPPIUR204 M50CYPPIUR212 M50CYP2UR208 M50CYP2UR312	Intel® VROC Standard SKU	5262C5140058 (FIO)

Jabil Supported Platforms	Intel VROC SKU Description	Jabil PN
2U Eagle Stream Server – J322-S	Intel VROC Standard SKU	Standard PN

Nutanix Supported Platforms	Intel® VROC SKU Description	Nutanix PN
NX-1175S G9 NX-3035 G9, NX-8035 G8* NX-3060 G9, NX-3060 G8* NX-3155 G9, NX-3155G G8* NX-8150 G9, NX-8150 G8* NX-8155 G9, NX-8155 G8* NX-8170 G9, NX-8170 G8* NX-3170 G8*	VMD Only	VMD Only

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)

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

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)

Compal Supported Platforms	Intel® VROC SKU Description	Compal PN
General Purpose: SR110-2, SR111-2, SR120-2, SR122-2, SR210-2, SR211-2, SR220-2, SR222-2 Multi-Node: SD220-8 GPU Server: SG220-2	Intel® VROC Premium SKU	DC210000W0Q

¹FIO - Factory Integrated Order (CTO)

²FUO - Field Upgrade Order (BTO)

OEM Enabled License Models



Software Key

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



Hardware Key

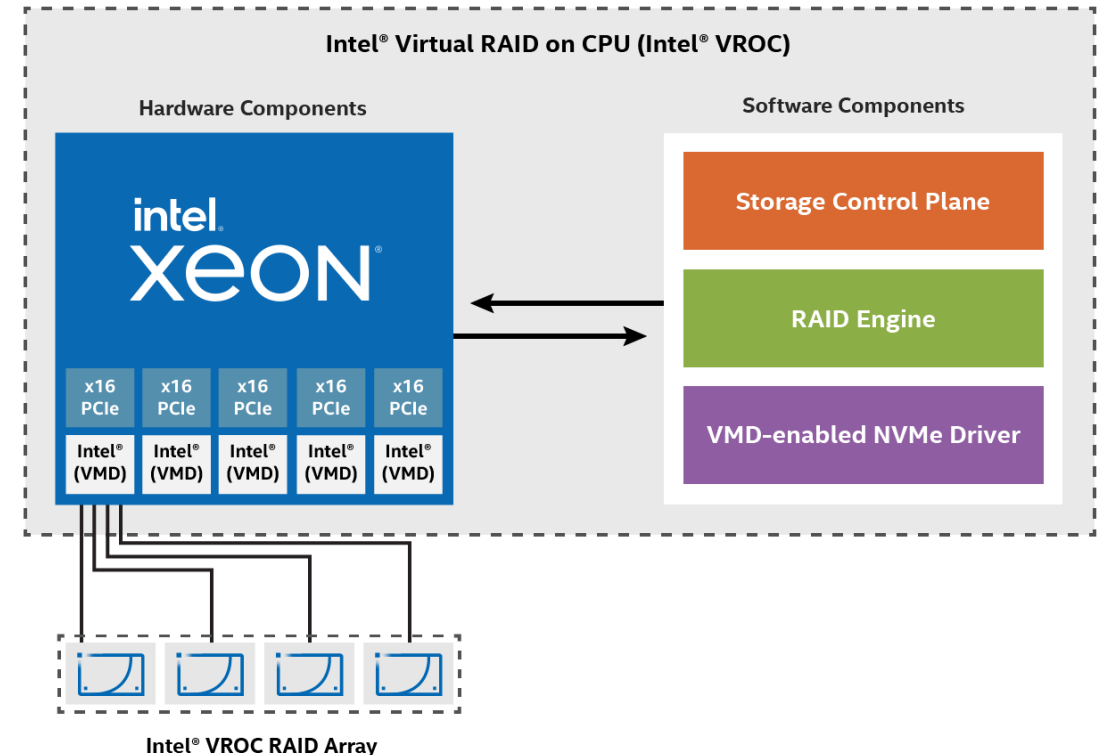
- OEM design into motherboard
- OEM purchases key from Intel
- Field upgradable

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

Key Features

Intel® VROC Features

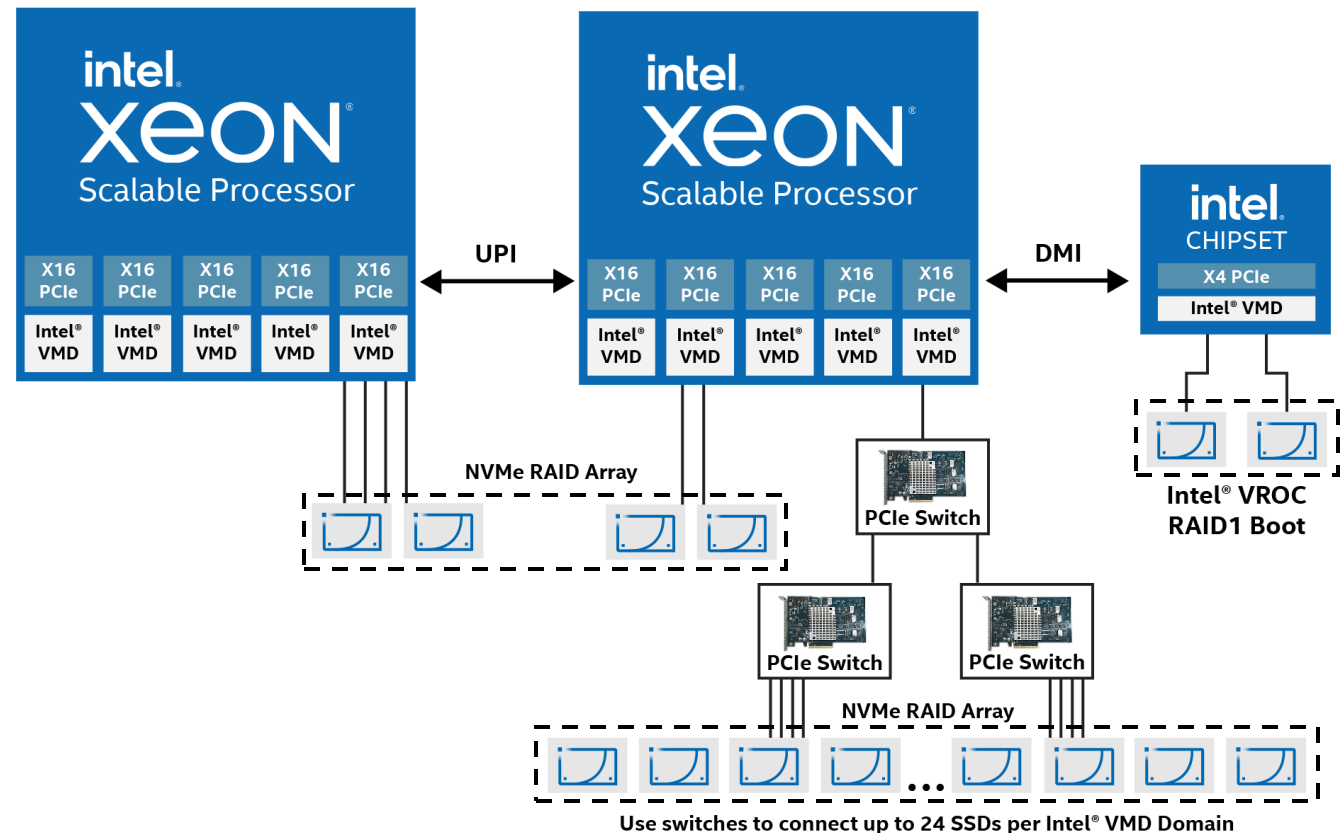
- Volume Management Device (VMD)
- Boot and Data RAID 0/1/5/10
- Data RAID spanning
- NVMe and SATA SSD Drives
- Hot-plug/ LED Management/ Fault Isolation
- SED Key Management
- Out-of-band management (OOB)
- VMD Direct Assign for hypervisor use case
- OS and UEFI Pre-OS supported:
 - Windows, Linux,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



Windows and Linux: Supported Configurations

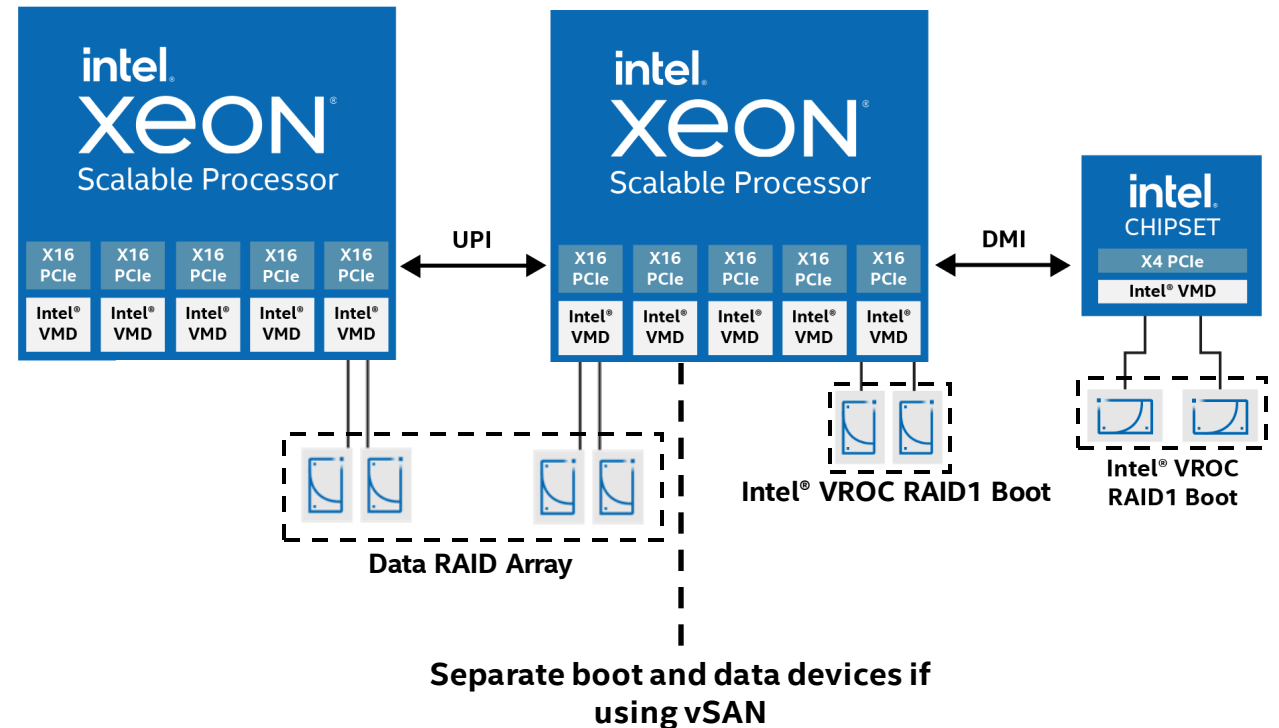
Intel VROC Configuration Guidelines

- 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

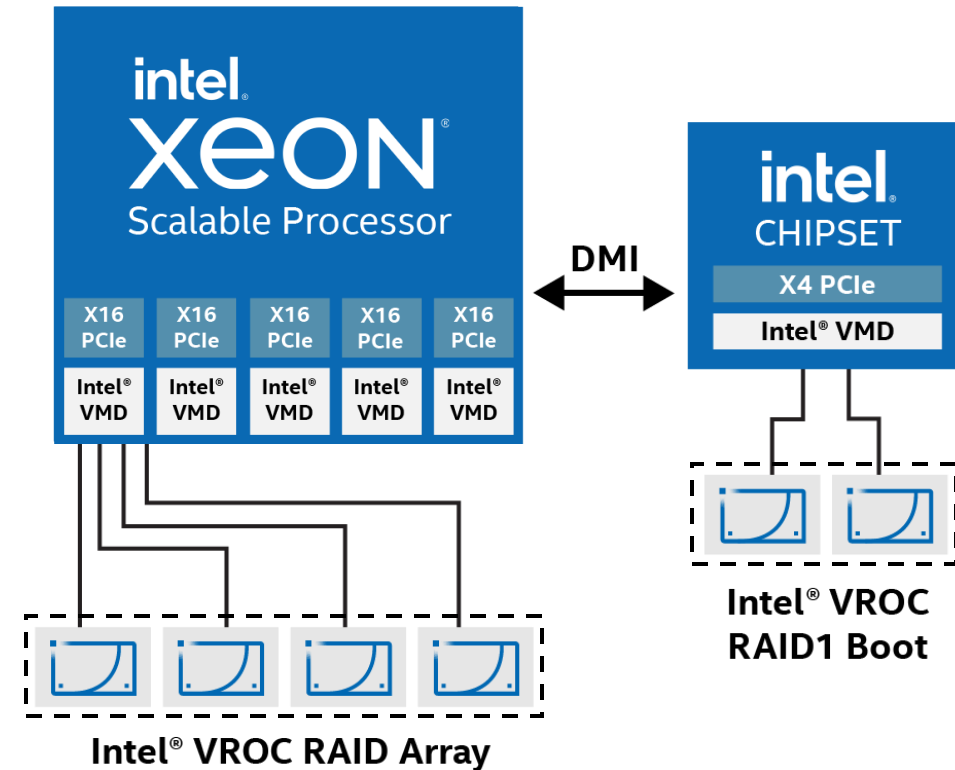
- **Intel VMD-enabled NVMe RAID1 Driver**
 - In-boxed in VMware ESXi
- VMware vSAN certified
 - If vSAN is being used: Boot volumes must be on separate VMD Domain from data devices
 - If vSAN is NOT being used: Boot volumes and data devices can share a VMD domain
- RAID Level Support: Boot and Data RAID1
 - 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



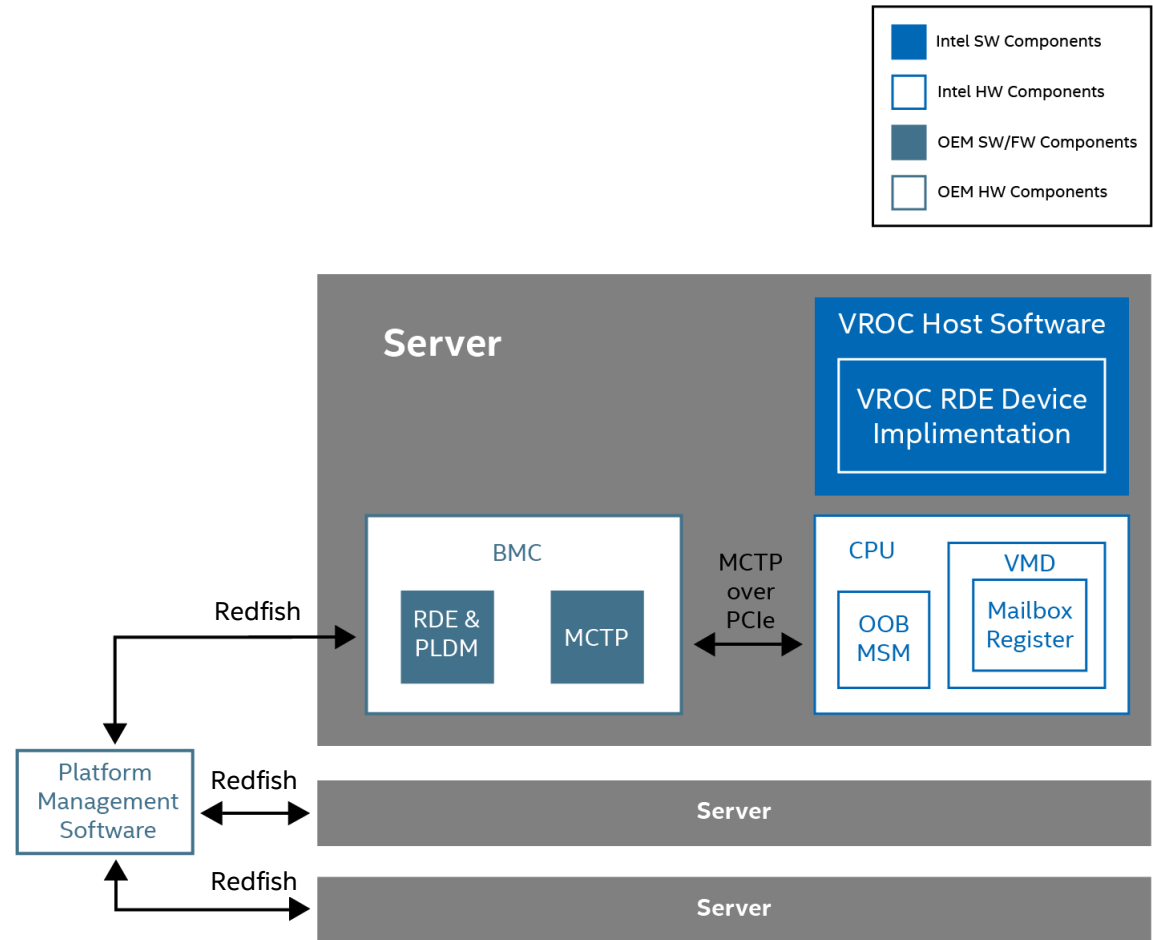
Note: New VROC license option for boot only RAID1 will be introduced in VROC 8.5

* VMWare limited to RAID1 boot

Out-of-Band Management

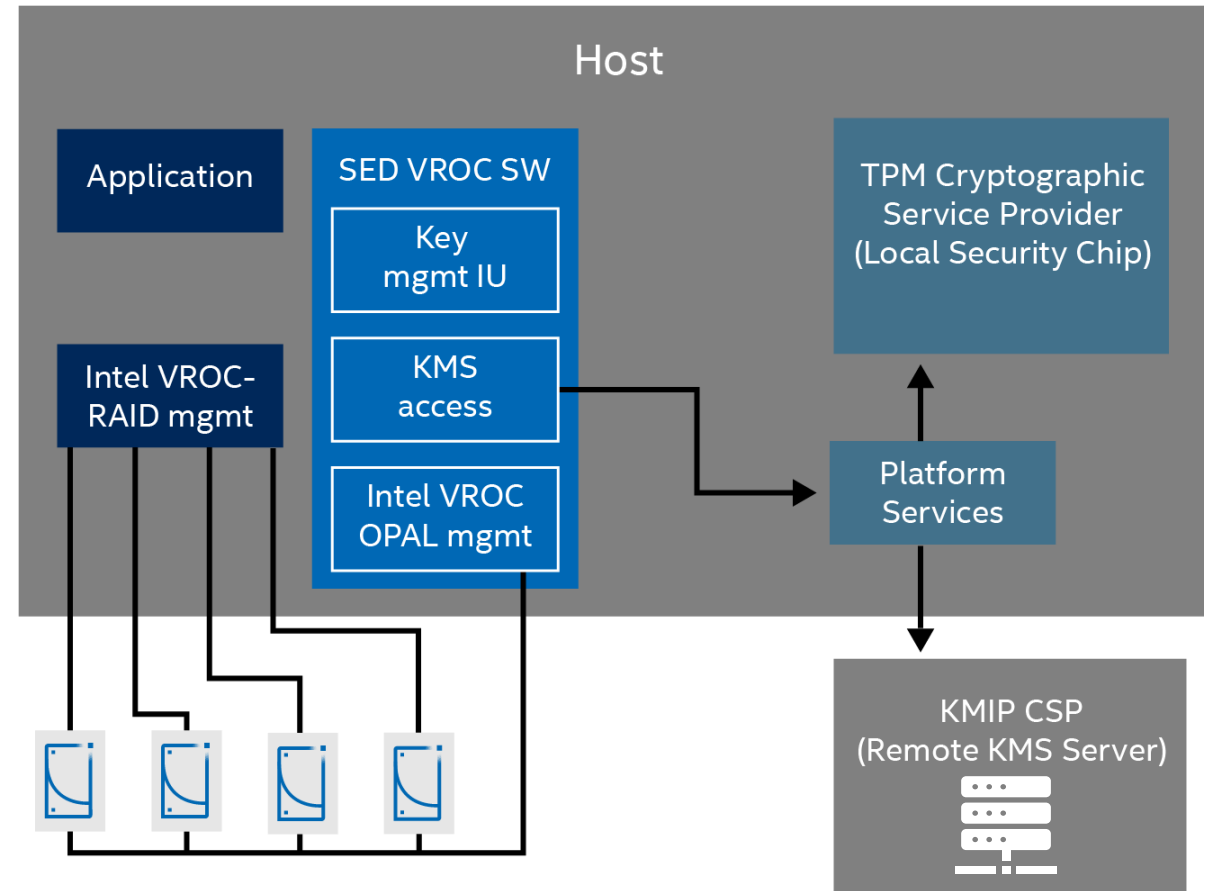
VROC 8.0 OOB Architecture

- **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 Management¹

- 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



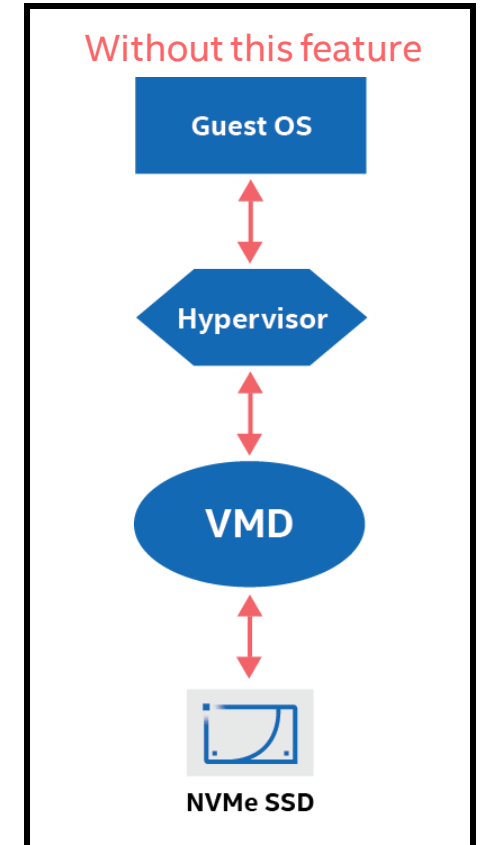
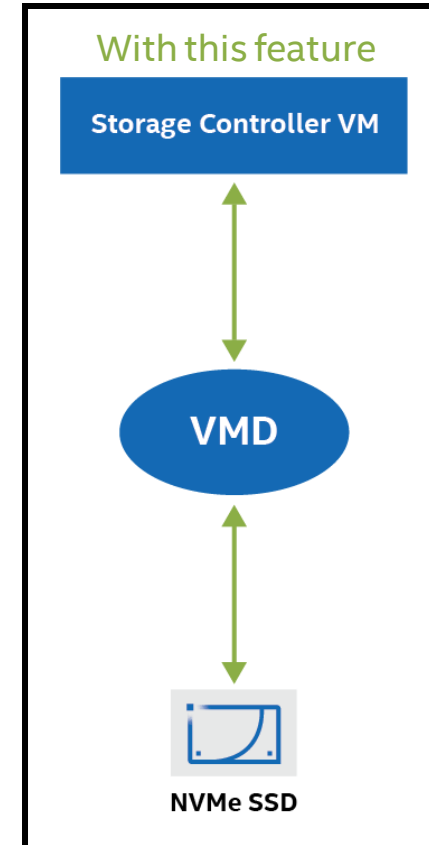
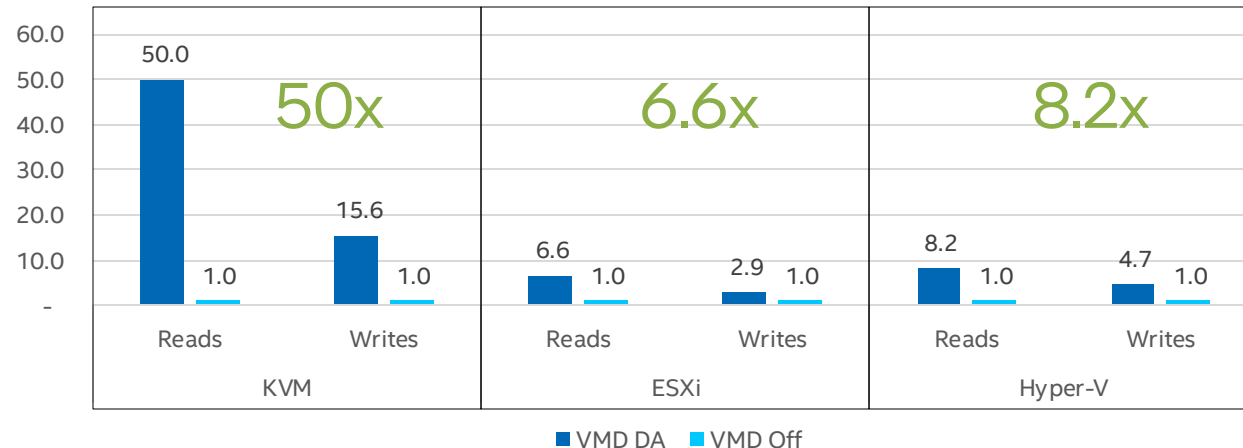
¹-Requires additional driver integration at platform level. May not be available on all Intel VROC enabled platforms. Specific functionality depends on platform provider integration into their preferred KMS

Virtualization Support

VMD Direct Assign

Intel® VMD Direct Assign Access Path

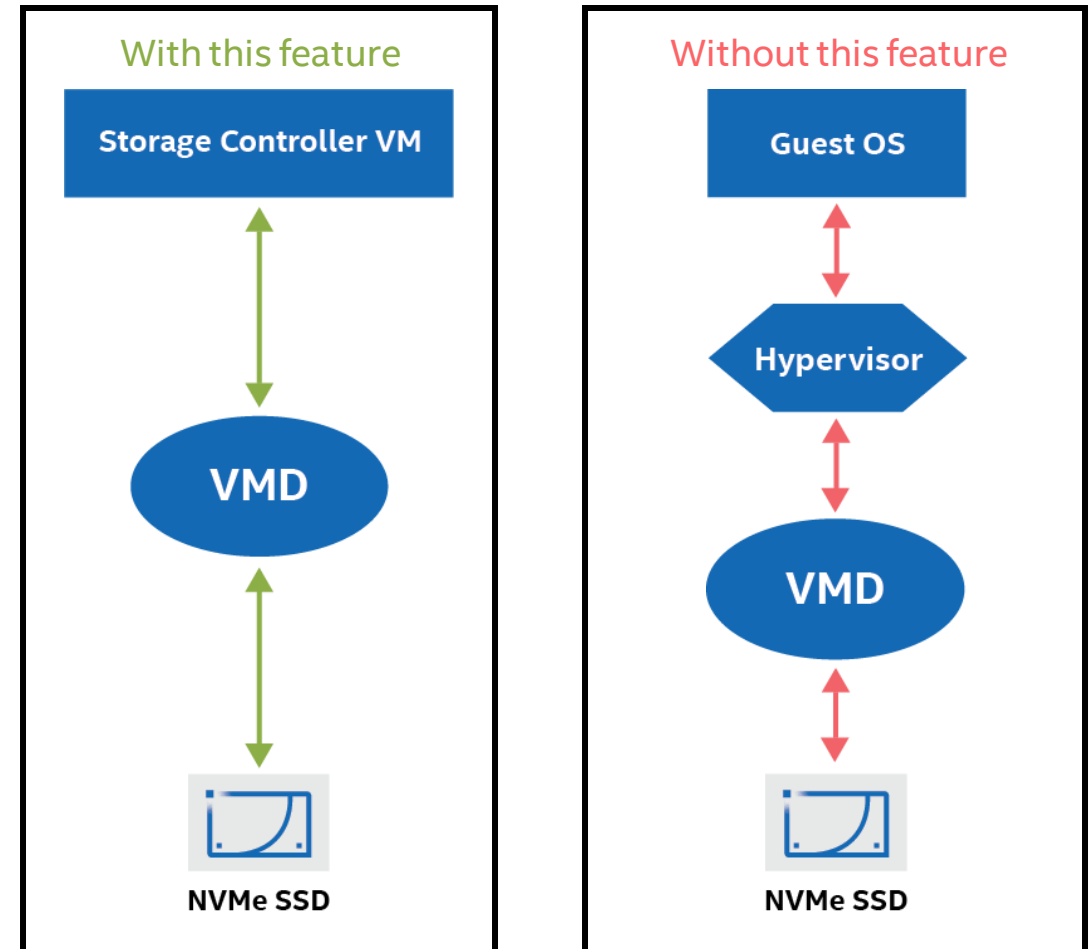
- Up to 50x performance benefit with KVM
- VMD attached devices appear in Guest OS as physical NVMe with hot-add, hot-remove and LED management functionality
- The VMD Domain and all NVMe attached must be Direct Assigned
 - Use switches to assign up to 24 devices to a VM



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

System Reliability, Data Integrity & Data Recovery

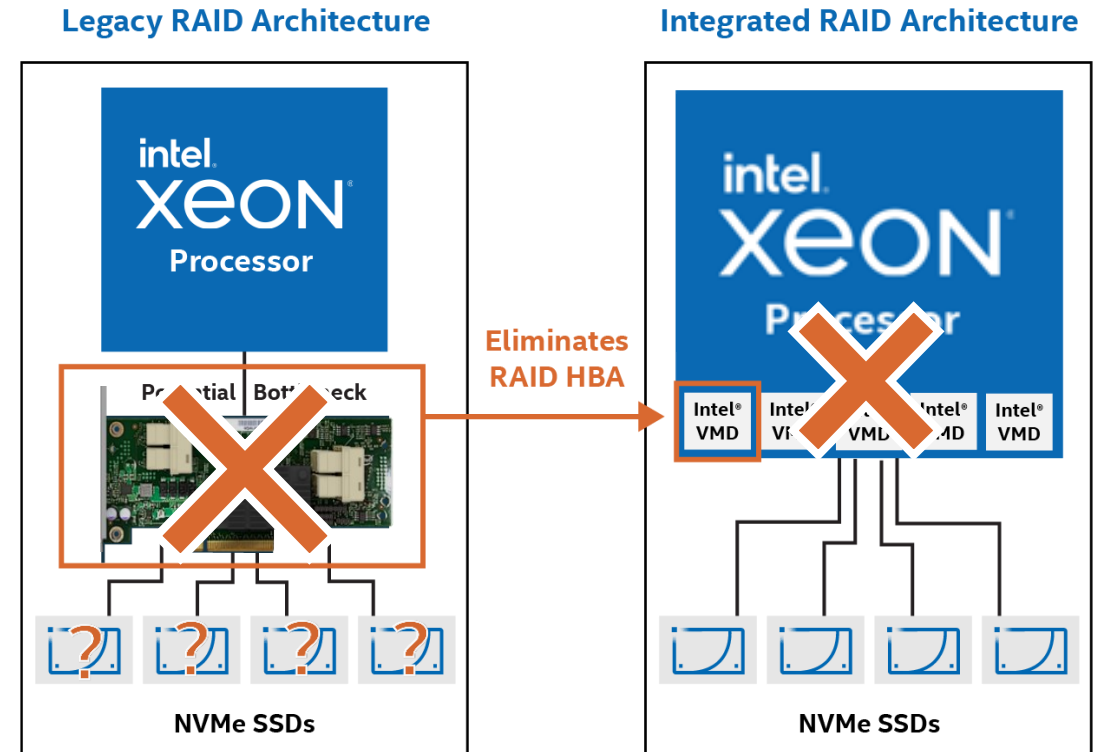
Intel® VROC Improves System Reliability

Eliminating HBA improves system MTBF

- Intel® VROC removes a single point of failure from the system
 - Both systems rely on the same CPU, memory, and storage
 - Both Intel VROC and HBA have firmware and software

Roaming data security with Intel® VROC

- When system fails, you can just replace the CPU or connect the array to another VROC enabled system



Data Integrity with Intel® VROC

NVMe Power Loss Imminent (PLI)

- Power Loss Imminent (PLI) consists of capacitors present on Enterprise Devices to ensure WRITES are fully persistent

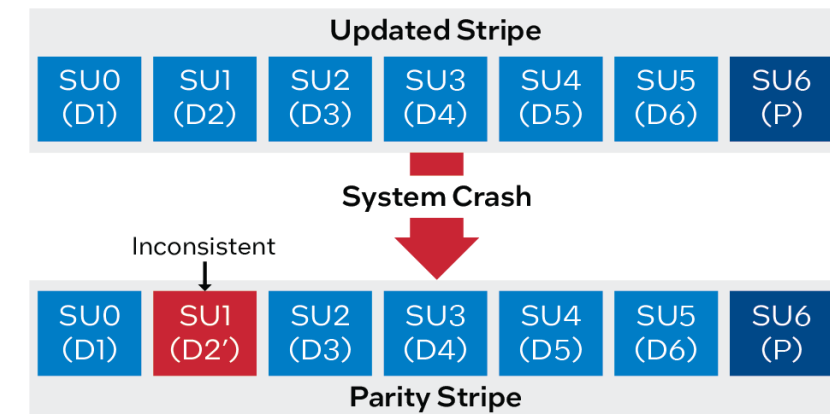

Preventing data loss with PPL

- For parity-based RAID, power loss could lead to scenarios like RAID Write Hole (RWH)
- Instead of relying on a battery backup unit (BBU), Intel VROC addresses RWH by using patented¹ journaling that creates a Partial Parity Log (PPL)

Root Port Isolation with Intel® VMD

- Intel® VMD captures any Non-Fatal Errors at the root port before they can make it up to the system

1. Patent # 9921914




Fault Recovery – Rebuild

2DRI – 70/30 R/W Workload

- Intel VROC delivers **85%** better performance and utilizes only **0.09%** higher CPU utilization
- Intel VROC offers **similar rebuild times** when compared to the HBA

	VROC	HBA
Time to rebuild	3 hours & 44 minutes	3 hours & 51 minutes

- While rebuilding, Intel VROC delivers **4.3x** better performance with only **0.65%** higher CPU utilization



But that is not the complete story...


See backup D for configuration details. Results may vary

Fault Recovery – Rebuild

2DRI – 70/30 R/W Workload

- Just like discreet RAID cards, Intel VROC can be configured to deliver even faster rebuild times
- Intel VROC delivers **2x** better performance while rebuilding with only **0.32%** higher CPU utilization
- Intel VROC rebuild times are **3.7x** better than the HBA, delivering excellent reliability

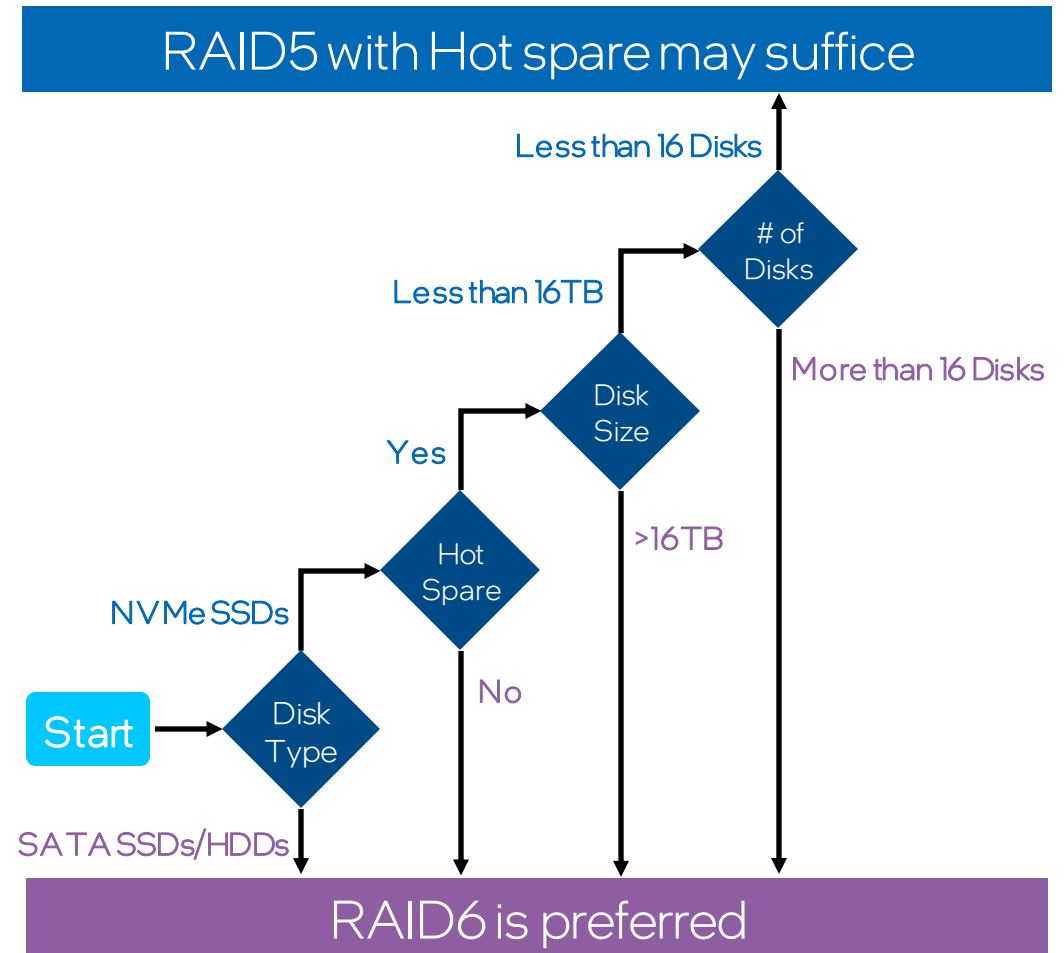
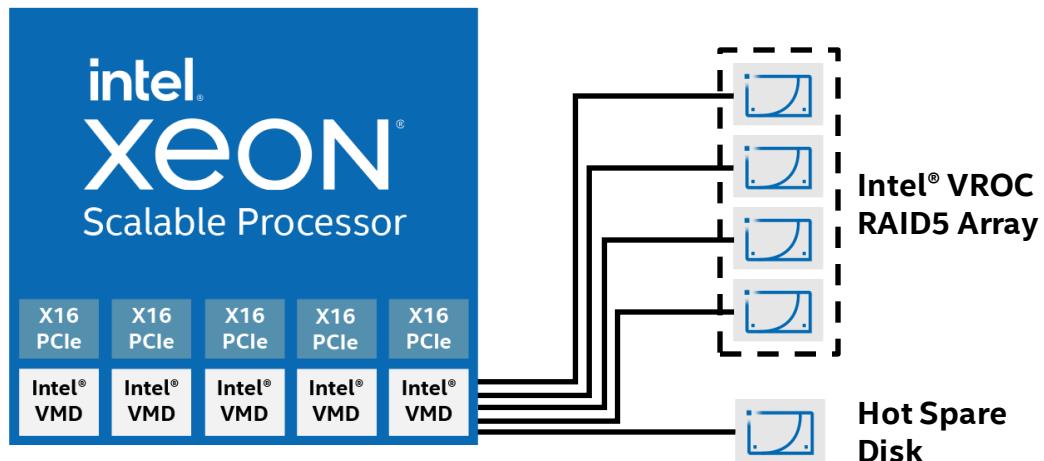
	HBA	VROC (Tuned)
Time to rebuild	3 hours & 51 minutes	1 hours & 2 minutes



See backup D for configuration details. Results may vary

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

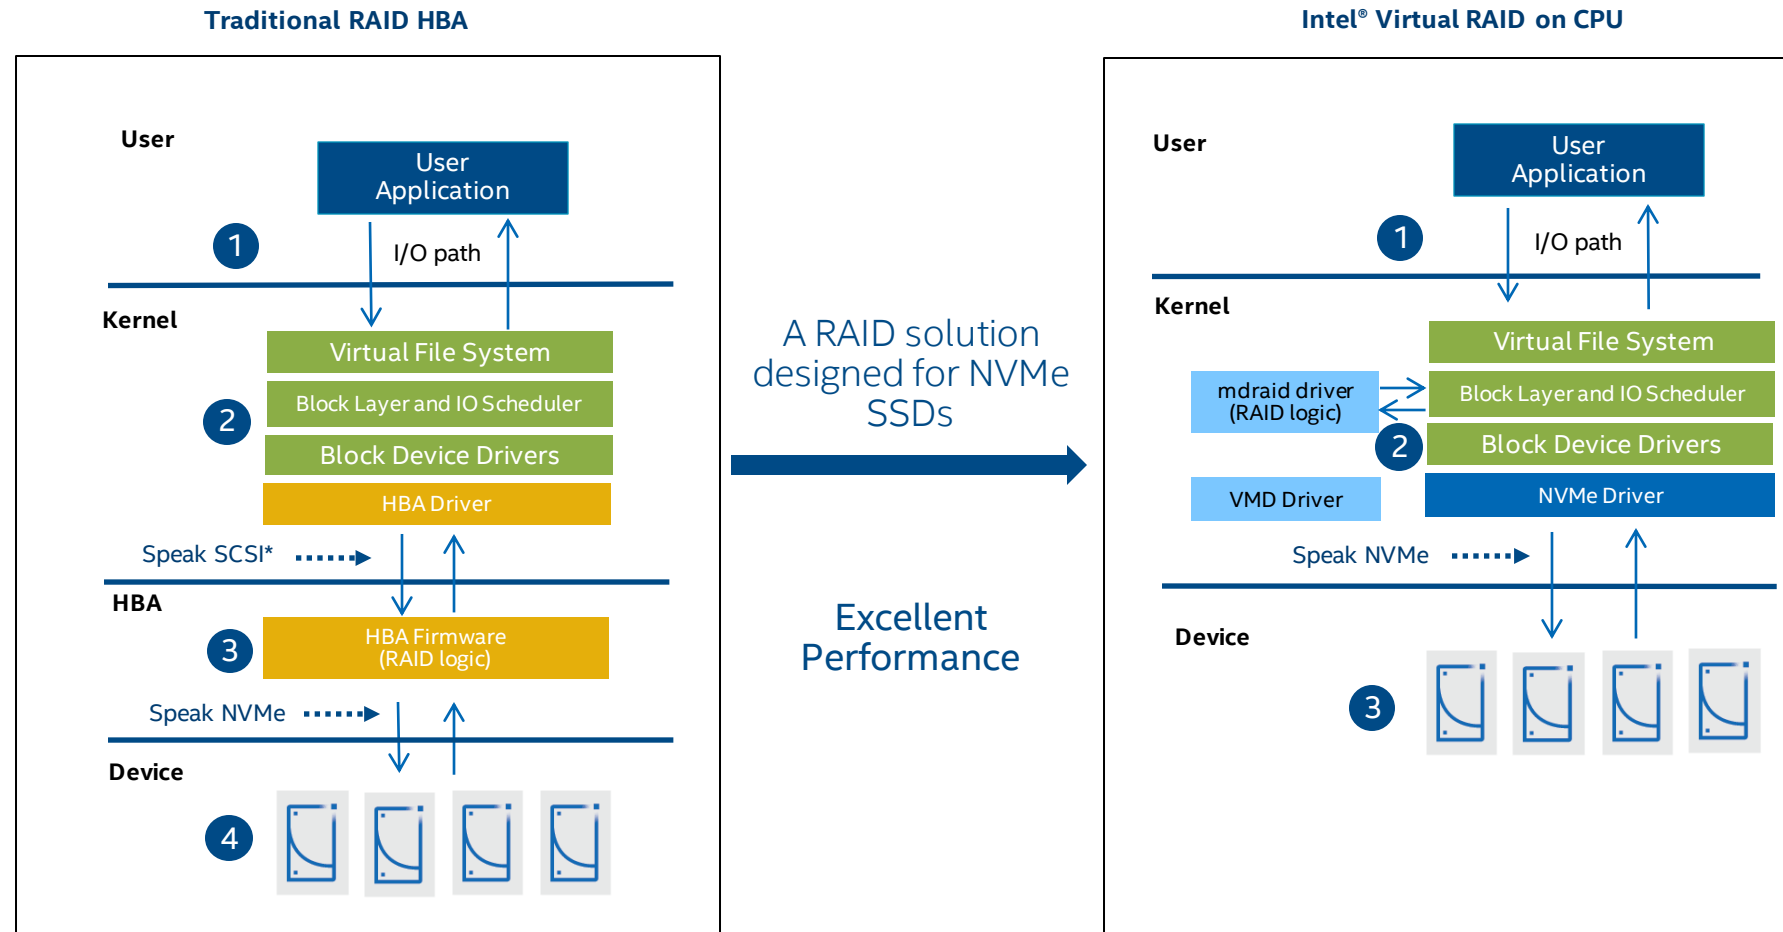


Power



Application
Performance

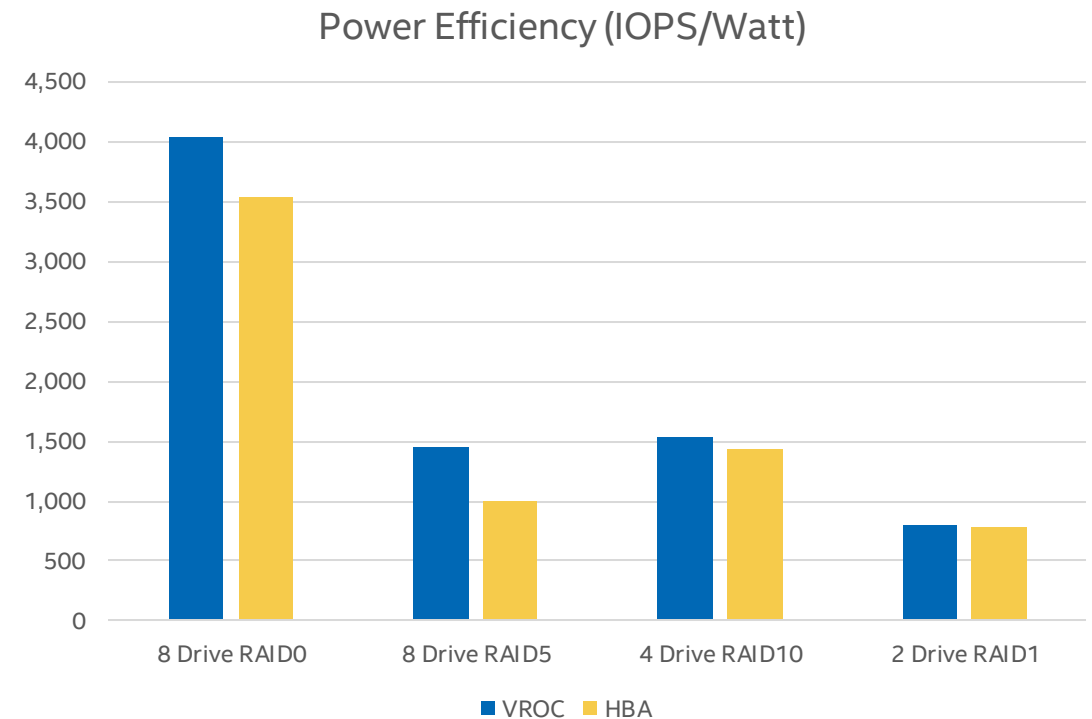

Intel® VROC Software Stack vs HBA



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

Intel® VROC Performance

Intel VROC achieves up to 3.2 million IOPS with RAID0 on mixed workloads



Result is based on Solidigm P5620 with 4KB random 70/30 R/W workload using 16 Threads & 32 IODepth

See backup C for configuration details. Results may vary

CPU Utilization, Throughput & Latency Comparison

RAID0 8-drive on Solidigm P5620 with 4KB random 70/30 R/W workload using 16 Threads



~0.17% ↓

CPU Utilization



~15% ↑

Throughput



~13% ↓

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
32	4.00%	4.30%	-0.30%	13,481	12,080	1.12x	155	173	0.90x
64	4.10%	4.50%	-0.40%	13,881	12,478	1.11x	302	336	0.90x
128	4.70%	4.50%	0.20%	15,254	12,451	1.23x	550	673	0.82x

See backup C for configuration details. Results may vary

CPU Utilization, Throughput & Latency Comparison

RAID5 8-drive on Solidigm P5620 with 4KB random 70/30 R/W workload using 16 Threads



~6.87% ↑

CPU Utilization



~69% ↑

Throughput



~36% ↓

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
32	8.10%	1.50%	6.60%	4,797	3,044	1.58x	535	689	0.78x
64	8.30%	1.50%	6.80%	4,779	2,785	1.72x	849	1,506	0.56x
128	8.60%	1.40%	7.20%	4,928	2,787	1.77x	1,780	3,010	0.59x

See backup C for configuration details. Results may vary

Better Overall Performance

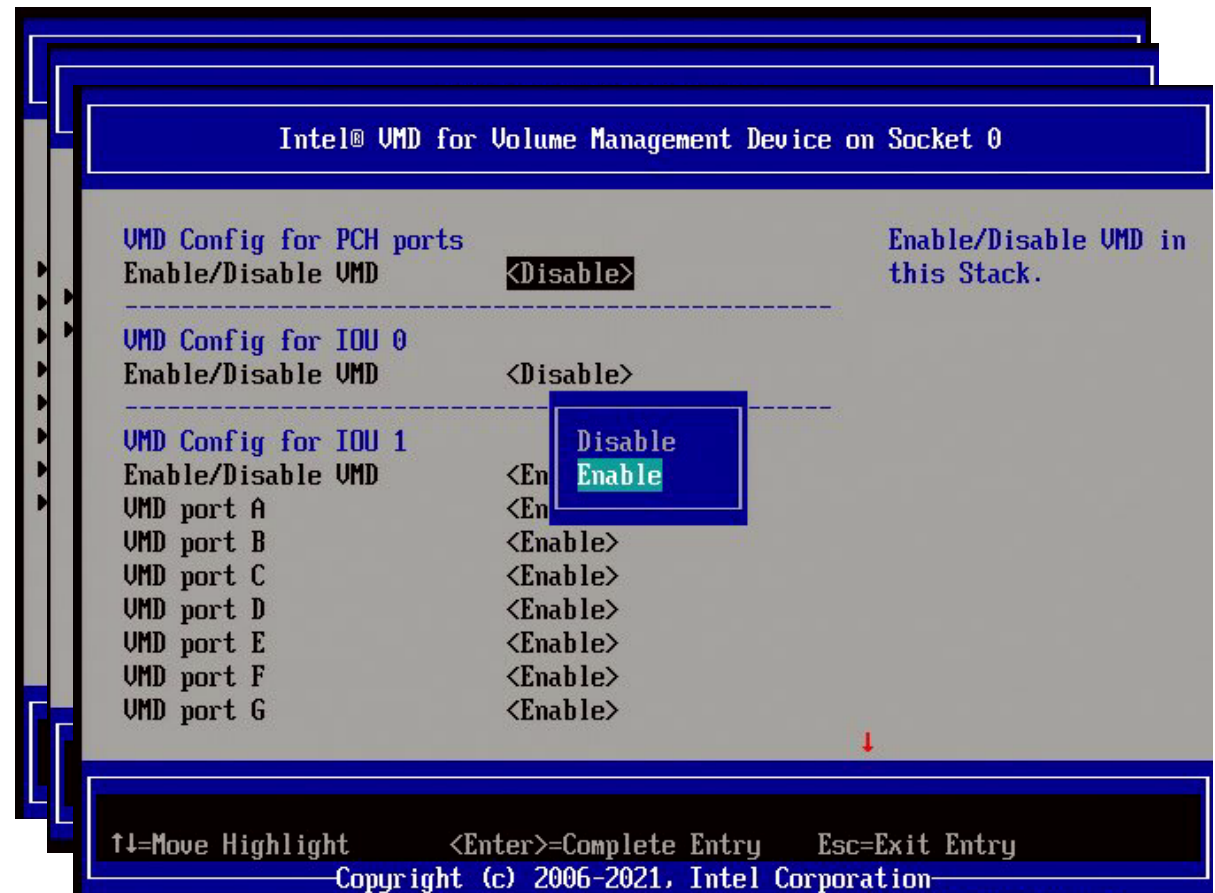
1. Intel VROC performance for all RAID levels is equal or better than RAID HBA **↑ IOPS, ↓ Latency**
2. Intel VROC implementation efficiently utilizes CPU resources where HBA is limited by bandwidth bottleneck **Improved Application Performance at Peak Workload**
3. Intel VROC is designed to **scale** 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

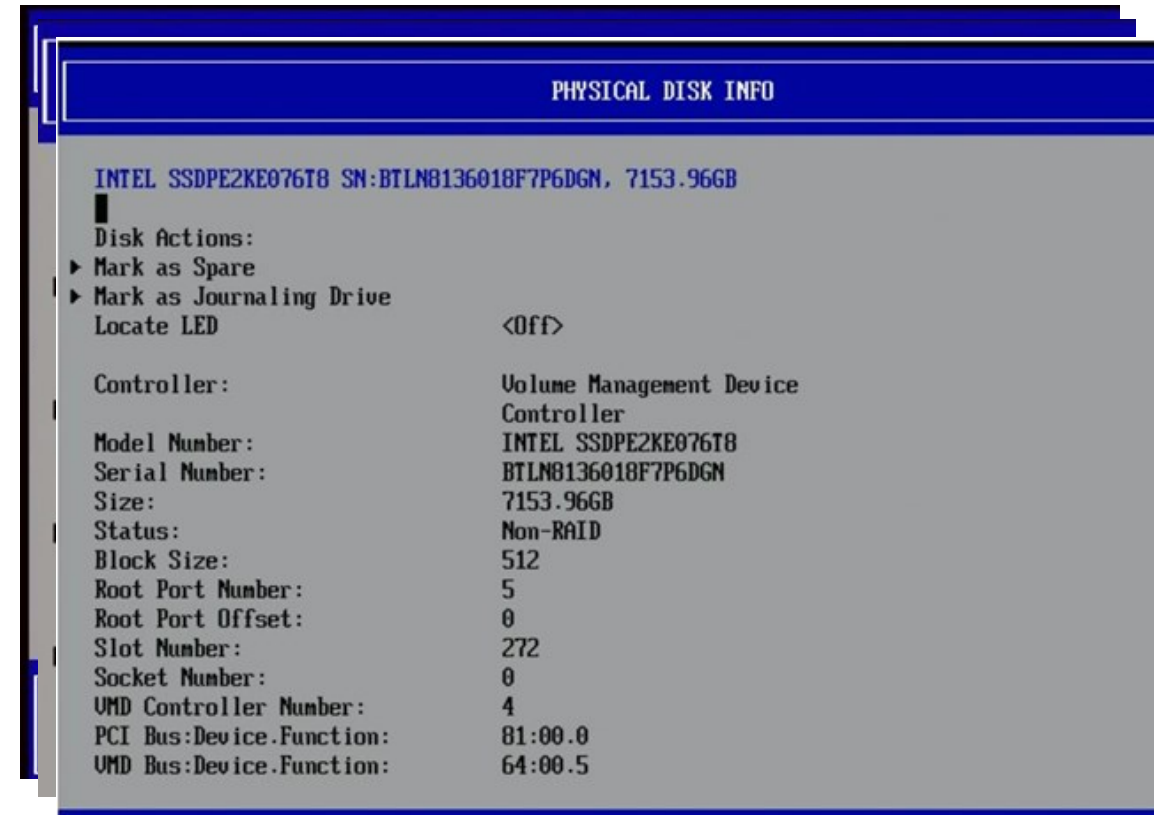
1. Enable VMD
2. Choose the Processor
3. Choose to enable each x4 PCIe lane for VMD Config



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

Configuring Intel VROC RAID Volumes from the CRB BIOS Setup Environment

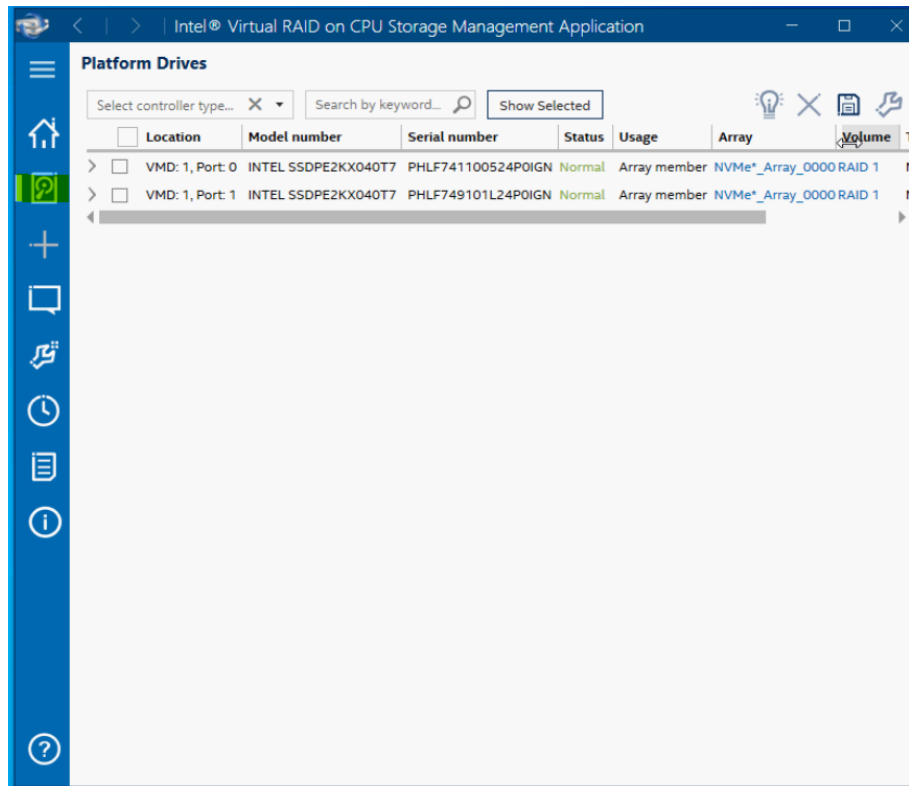
1. Enter Intel® Virtual RAID on CPU Menu
2. Create RAID and View Existing RAID Volumes and NVMe Devices
3. View each device and perform actions



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/ism0 /dev/nvme[0-1]n1 -n 2 -e ism  
# mdadm -C /dev/md/md0 /dev/md/ism0 -n 2 -1 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 -1 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 -1 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

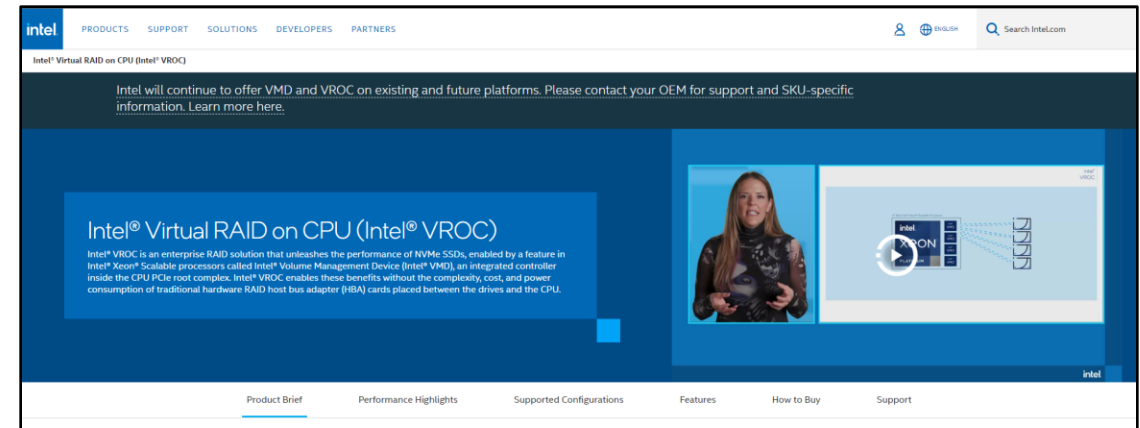
- Intel.com/VROC for:

- Product Brief
- FAQ
- Supported OS/HW Information
- Related Links


- Intel® VROC Documentation:

- User Guides
- Tech Briefs
- Performance Documentation

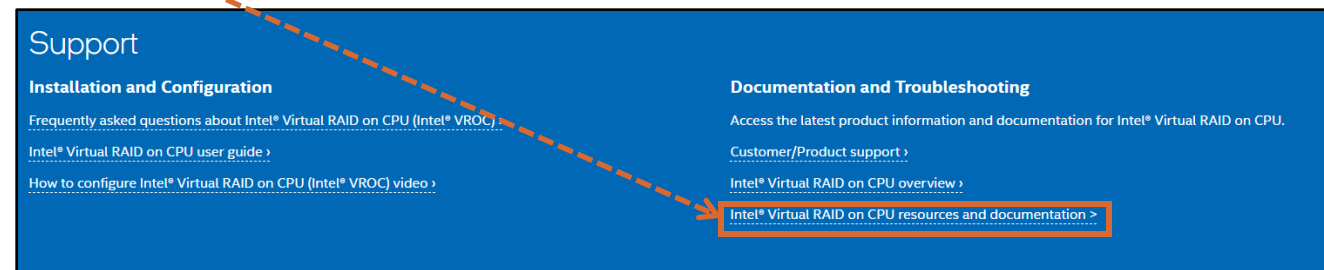
- Support available via IPS



Intel® Virtual RAID on CPU Website



Download support



Get product support

Backup

Intel® VROC Supported Processors/Chipsets

Intel® Xeon® Processor	Chipset	Platform Codename	VMD NVMe RAID	SATA RAID
Intel® Xeon® Scalable Processors	C620 Series	Purley	✓	✓
Intel® Xeon® W Processor	C400 Series	Basin Falls	✓	✓
Intel® Xeon® E Processor	C240 Series	Mehlow		✓
Intel® Xeon® D Processor	Integrated in CPU	Bakerville / Idaville	✓	✓
2 nd Generation Intel® Xeon® Scalable Processors	C620 Series	Purley R	✓	✓
3 rd Generation Intel® Xeon® Scalable Processors 4S/8S (-H)	C620 Series	Cedar Island	✓	✓
3 rd Generation Intel® Xeon® Scalable Processors 1S/2S	C620 Series	Whitley	✓	✓
4 th Generation Intel® Xeon® Scalable Processors 2S/4S	C741 Series	Eagle Stream	✓	✓
4 th Generation Intel® Xeon® Scalable Processors 1S/2S	W790 Series	Fishhawk Falls	✓	✓

Backup A – Configuration Details

Performance results are based on testing by Intel as of November 27, 2023 and may not reflect all publicly available updates. Results may vary.

- **Platform:** Platform: ArcherCity CRB; 2x INTEL(R) XEON(R) PLATINUM 8592+ (QDF: Q3W7) (64 cores each) (EagleStreamRefresh - Emerald Rapids); 128GB RAM (8 x 16GB Hynix HMCG78AGBRA191N 5800 MT/s DDR5 Synchronous Registered (Buffered) DIMMs); BIOS Version: EGSDREL1.SYS.0105.D74.2308261931 (Microcode revision: 0xa1000161)
- **BIOS Settings:** SpeedStep(Enabled), Turbo(Enabled), ProcessorC6(Enabled), C1E (Disabled), PackageC-State(C0/C1 State), C1 Auto Demotion (Disabled), CPU_PowerAndPerformancePolicy(Performance), HardwareP-States(NativeMode), WorkloadConfiguration(I/O Sensitive), Hyperthreading (Enabled)
- **Storage:** 8x 3.2TB Solidigm P5620 PCIe Gen4 U.2 SSDs (Model: SSDPF2KE032TIES, Firmware: 9CV10410) connected to backplane which is connected to CPU0 3A-D, 4A-D MCIO PCIe ports (NUMA Node 1, CPU 0)
- **Virtual Disk: KVM:** Storage pool with "disk" type created on all disks; Disks attached to the virtual machine as block devices (on "virtio" bus); **ESXi:** VMFS partitions with "Thick Provisioned - Eager zeroed" Virtual Disks; Virtual controllers attached to VM 4x"VMware Paravirtual"; Disks attached to the virtual machine with physical compatibility. (2 virtual disks for each of the virtual controllers); **Hyper-V:** Fixed size virtual disk file on NTFS partition
- **Direct Assign:** VMD enabled on any applicable ports containing NVMe devices (with Intel(R) VROC Premium license); Intel(R) VROC PreOS Version: 8.5.0.1082; mdadm version = mdadm-4.2-8.el9.x86_64; vmd driver version = inbox; 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 9.2 GA, Kernel: 5.14.0-284.11.1.el9_2.x86_64
- **RAID Configurations:** 6-Disk RAID0 with 4KB random 70/30 R/W workload using 16 Threads and 32 IODepth with Intel VROC and Intel RS3P4TF160F
- **FIO version:** Fio version: 3.35

Backup B – Configuration Details

Performance results are based on testing by Intel as of October 25, 2023 and may not reflect all publicly available updates. Results may vary.

- **Platform:** Platform: ArcherCity CRB; 2x INTEL(R) XEON(R) PLATINUM 8592+ (QDF: Q3W7) (64 cores each) (EagleStreamRefresh - Emerald Rapids); 128GB RAM (8 x 16GB Hynix HMCG78AGBRA191N 5800 MT/s DDR5 Synchronous Registered (Buffered) DIMMs); BIOS Version: EGSDREL1.SYS.0105.D74.2308261931 (Microcode revision: 0xa1000161)
- **BIOS Settings:** SpeedStep(Enabled), Turbo(Enabled), ProcessorC6(Enabled), C1E (Disabled), PackageC-State(C0/C1 State), C1 Auto Demotion (Disabled), CPU_PowerAndPerformancePolicy(Performance), HardwareP-States(NativeMode), WorkloadConfiguration(I/O Sensitive), Hyperthreading (Enabled)
- **Storage:** VROC: 6 x 400GB Intel Optane P5810X PCIe Gen4 U.2 SSDs (Model: SSDPF21Q400GA, Firmware: L0310600) connected to backplane which is connected to a midplane switch via special Broadcom SlimSAS cables which is connected to CPU0 3A-D, 4A-DMC10 PCIe ports (NUMA Node 1, CPU0); Trimode with Midplane switch: 6 x 400GB Intel Optane P5810X PCIe Gen4 U.2 SSDs (Model: SSDPF21Q400GA, Firmware: L0310600) connected to backplane which is connected to a midplane switch via special Broadcom SlimSAS cables which is connected to an IntelRS3P4TF160F (x8) (equivalent to Broadcom 9560-16i) card on Riser1, PCIe slot 1 (NUMA Node 0, CPU0); Model: CYP SWITCHMP; Firmware package version: 4.6; Firmware version: 1.06.00.00; Firmware mode: HWRX12; OS is on NVMe 118GB Intel Optane P1600X M.2 SSDs (Model: SSDPEK1F118GA, Firmware: U50A0550) connected to M.2 sockets on the motherboard
- **RAID Controller:** IntelRS3P4TF160F (x8) (equivalent to Broadcom MegaRAID 9560-16i); Firmware Package Build = 52.25.0-4902; Firmware Version = 5.250.02-3847; megaraid_sas driver version = 07.726.02.00; CLI Version = 007.2612.0000.0000 June 13, 2023; Added "scsi_mod.use_blk_mq=y" to grub boot option for maximum throughput on the Broadcom card; When creating RAID volumes "pdcache=off, Write-Through, No Read Ahead, Direct I/O" **OR** Intel(R) VROC PreOS Version: 8.0.0.4006; mdadm version = mdadm-4.2-8.el9.x86_64; vmd driver version = inbox; Added "pci=pcie_bus_perf" to grub boot option which sets MaxPayload to the maximum for each of the NVMe devices; VMD enabled on any applicable ports containing NVMe devices (with Intel(R) VROC Premium license)
- **OS:** Red Hat Enterprise Linux 9.2 GA, Kernel: 5.14.0-284.11.1.el9_2.x86_64
- **RAID Configurations:** 6-Disk RAID0 with 4KB random 70/30 R/W workload using 16 Threads and 32 IODepth with Intel VROC and IntelRS3P4TF160F
- **FIO version:** Fio version: 3.35

Intel® VROC Cost

	RAID HBA ¹	Intel® VROC (Premium SKU)
RAID Cost	\$1257	\$499

60% lower cost

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 10/09/2023

Power efficiency

	VROC	HBA
IOPS	6,439,585	2,425,548
Power (W)	828	796
IOPS/Watt	7,779	3,048

155% better Power Efficiency

Backup C – Configuration Details

Performance results are based on testing by Intel as of November 27, 2023 and may not reflect all publicly available updates. Results may vary.

- **Platform:** Platform: ArcherCity CRB; 2x INTEL(R) XEON(R) PLATINUM 8592+ (QDF: Q3W7) (64 cores each) (EagleStreamRefresh - Emerald Rapids); 128GB RAM (8 x 16GB Hynix HMCG78AGBRA191N 5800 MT/s DDR5 Synchronous Registered (Buffered) DIMMs); BIOS Version: EGSDREL1.SYS.0105.D74.2308261931 (Microcode revision: 0xa1000161)
- **BIOS Settings:** SpeedStep(Enabled), Turbo(Enabled), ProcessorC6(Enabled), C1E (Disabled), PackageC-State(C0/C1 State), C1 Auto Demotion (Disabled), CPU_PowerAndPerformancePolicy(Performance), HardwareP-States(NativeMode), WorkloadConfiguration(I/O Sensitive), Hyperthreading (Enabled)
- **Storage:** VROC: 8 x 3.2TB Solidigm P5620 PCIe Gen4 U.2 SSDs (Model: SSDPF2KE032T1, Firmware: 9CV10410) connected to backplane which is connected to CPU0 3A-D, 4A-D MCIO PCIe ports (NUMA Node 1, CPU 0); Trimode with Midplane switch: 8 x 3.2TB Solidigm P5620 PCIe Gen4 U.2 SSDs (Model: SSDPF2KE032T1, Firmware: 9CV10410) connected to backplane which is connected to a midplane switch via special Broadcom SlimSAS cables which is connected to an Intel RS3P4TF160F (x8) (equivalent to Broadcom 9560-16i) card on Riser 1, PCIe slot 1 (NUMA Node 0, CPU0); OS is on NVMe 118GB Intel Optane P1600X M.2 SSDs (Model: SSDPEK1F118GA, Firmware: U50A0550) connected to M.2 sockets on the motherboard
- **RAID Controller:** Intel RS3P4TF160F (x8) (equivalent to Broadcom MegaRAID 9560-16i); Firmware Package Build = 52.25.0-4902; Firmware Version = 5.250.02-3847; megaraid_sas driver version = 07.726.02.00; CLI Version = 007.2612.0000.0000 June 13, 2023; Added "scsi_mod.use_blk_mq=y" to grub boot option for maximum throughput on the Broadcom card; When creating RAID volumes "pdcache=off, Write-Through, No Read Ahead, Direct I/O" OR Intel(R) VROC PreOS Version: 8.0.0.4006; mdadm version = mdadm-4.2-8.el9.x86_64; vmd driver version = inbox; Added "pci=pcie_bus_perf" to grub boot option which sets MaxPayload to the maximum for each of the NVMe devices; VMD enabled on any applicable ports containing NVMe devices (with Intel(R) VROC Premium license)
- **OS:** Red Hat Enterprise Linux 9.2 GA, Kernel: 5.14.0-284.11.1.el9_2.x86_64
- **RAID Configurations:** 6-Disk RAID0 with 4KB random 70/30 R/W workload using 16 Threads and 32 IO Depth with Intel VROC and Intel RS3P4TF160F
- **FIO version:** Fio version: 3.35

Backup D – Rebuild Performance

Performance results are based on testing by Intel as of May 2, 2021 and may not reflect all publicly available security updates.

- **System configuration:** Intel® Beta Coyote Pass M50CYP2SBSTD (chassis M50CYP2UR208BPP), 2x Intel® Xeon® Platinum 8358 processor (2.60 GHz, 32 cores each), DRAM 128 GB. BIOS version (released October 19, 2021):SE5C620.86B.01.01.0004.2110190142 (Microcode revision: 0xd000311)
- **OS:** Red Hat Enterprise Linux Server 8.1 (Ootpa), Kernel: 4.18.0-147.el8.x86_64, Intel(R) VROC PreOS Version: 7.6.0.1012, mdadm version: mdadm - v4.1 - 2018-10-01
- **Storage:** Both configurations used 4 x 3.84TB Intel P5510 PCIe Gen4 U.2 SSDs (Model:SSDPF2KX038TZ, Firmware: JCV10100) connected to backplane which is connected via SlimSAS cables to the motherboard on the 2nd CPU (NUMA Node 1)
- **BIOS setting:** SpeedStep (Enabled), Turbo(Enabled), ProcessorC6(Enabled), PackageC-State(C0/C1 State), CPU_PowerAndPerformancePolicy (Performance), HardwareP-States(NativeMode), WorkloadConfiguration(I/O Sensitive)
- **RAID configuration:** 2-disk RAID1 with Intel VROC and Broadcom MegaRAID 9560-16i.
- **Workload generator:** FIO 3.25

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:

- $p(t) = 1 - e^{-t/MTTDL}$

Model ran in R to calculate probability of failure

intel®