

Intel® VROC for Better I/O Performance

Unlock the full performance potential of fast NVMe SSDs with Intel® Virtual RAID on CPU (Intel® VROC), which delivers improved I/O throughput without significant CPU overhead for mixed workloads

Table of Contents

Executive Summary	1
Business Challenge: HBAs Stand between Applications and Their Data... 2	
Achieve Better Performance without Compromising RAS	3
Conclusion.....	4
Learn More	4

Executive Summary

Conventional RAID configurations that use host bus adapters (HBAs) can protect data, but they inject an unnecessary bottleneck between modern, data-intensive applications and data stored on today’s fast NVMe storage devices.

Intel® Virtual RAID on CPU (Intel® VROC) is a unique, integrated (hardware and software) RAID solution that eliminates this HBA bottleneck. By directly connecting high-speed NVMe SSDs to powerful 3rd Generation Intel® Xeon® Scalable processors, Intel VROC provides extraordinary performance and enterprise-grade reliability, availability and serviceability (RAS). Using Intel VROC instead of HBAs can provide up to 161% more Input/Output Operations Per Second (IOPS)¹ and 25% higher throughput (for mixed workloads) without significantly increasing CPU utilization.²

Business Challenge: HBAs Stand between Applications and Their Data

One element of a successful enterprise is its ability to move and process data quickly. In today’s data-intensive business environment, high-performance access between application and storage is no longer optional—whether it’s a single SSD or a RAID. Better storage performance is measured by greater bandwidth, higher Input/Output Operations Per Second (IOPS) and lower latency; CPU utilization for targeted workloads is also a consideration.

The conventional approach to RAID using SATA and SAS drives and RAID host bus adapters (HBAs) fails to keep up with today’s data center needs. Legacy hardware RAID products traditionally isolated the storage sub-system from the host by managing the drives and RAID volumes connected to it. This design is sufficient for slower storage technologies, but with NVMe SSDs, a fundamentally new RAID architecture is required.

While SSDs and applications have evolved (such as the migration from SATA SSDs to NVMe SSDs), the HBA has not. RAID HBAs create an outdated bottleneck between applications and the data they need. Other problems associated with HBAs include:

- Being a single point of failure.
- Being difficult to obtain due to post-pandemic supply chain disruptions.

Intel® Virtual RAID on CPU (Intel® VROC) is an enterprise-grade integrated RAID solution that solves these dilemmas by connecting SSDs directly to 3rd Generation Intel® Xeon® Scalable processors, unlocking the full performance potential of NVMe SSDs to serve up data faster, without depending on expensive RAID HBA hardware.

Intel VROC Software Stack Overview

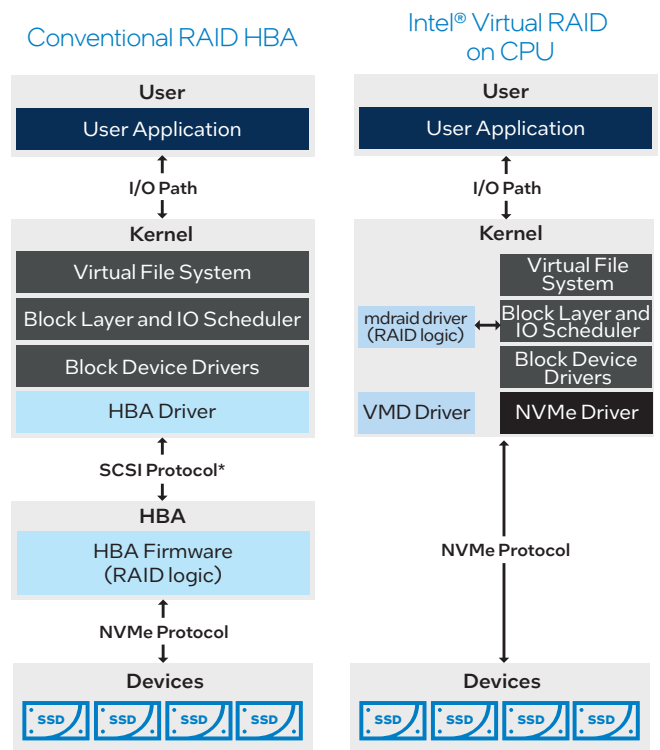
Figure 1 illustrates the I/O path in a simplified software stack, comparing a conventional HBA-based RAID set up and Intel VROC. The key difference is that the HBA offloads RAID logic (including parity calculation), whereas Intel VROC relies on the host CPU to execute RAID functions. It is a common misconception that an HBA does not require any host CPU cycles for RAID functions. In fact, the HBA device driver runs on the host CPU to communicate with HBA firmware across a fixed-width PCIe link. This architecture adds complexities (additional I/O protocol translations in the figure) and must be implemented in HBA software. As a result, it increases latency to service I/O requests. The fixed processing power on the HBA compute engine is not as scalable as multi-core Intel® Xeon® Scalable processors. Therefore, overall, the throughput and response time from an HBA is impacted. During normal workloads, both the HBA and Intel VROC deliver similar performance with comparable CPU usage; at peak workloads, Intel VROC provides better IOPS and lower latency with a small percentage of CPU overhead.¹

What Flavor of RAID Is Intel® VROC?

Hardware-only RAID: A hardware RAID solution has dedicated processor and memory to run the RAID application. It is implemented using a plug-in expansion card and offloads RAID processing from the host.

Software-only RAID: The RAID implementation is an application running on the host without any additional hardware. The pure software RAID solution is often integrated into the server OS and free of cost to users.

Integrated RAID: Intel VROC is a unique integrated hardware and software RAID solution that consists of a software layer built upon Intel® Volume Management Device (Intel® VMD), which is a unique root complex feature residing in Intel® Xeon® Scalable processors. This architecture enables a RAID implementation that maximizes PCIe bandwidth utilization (by directly attached NVMe SSDs), provides isolation of the CPU from the storage device and adds the flexibility inherent to software. This multi-character RAID design allows Intel VROC to differentiate itself from hardware-only RAID and software-only RAID.



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

Figure 1. Comparison between conventional HBA RAID and Intel® VROC stacks.

Achieve Better Performance without Compromising RAS

We performed a comparative analysis between Intel VROC and an equivalent HBA setup, using synthetic benchmark tools including Flexible I/O Tester (FIO) for all RAID levels and SysBench for online transaction processing (OLTP) workloads (e.g., MySQL database). The configurations of the tests are illustrated in Figure 2. All workloads in the tests are mixed read/writes.

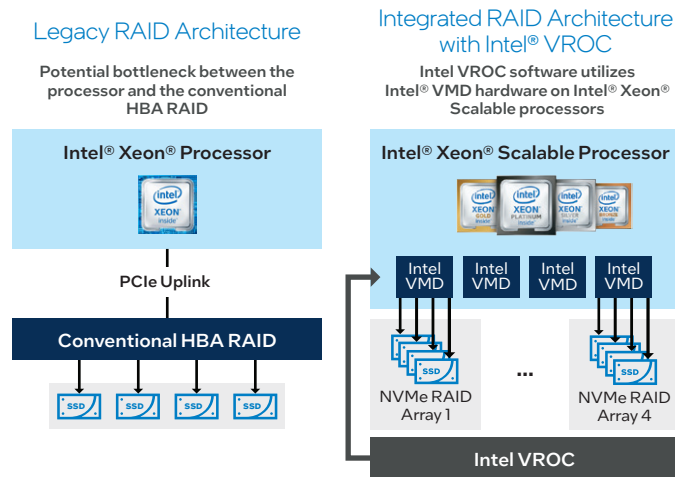


Figure 2. Test setup comparing HBA RAID to Intel® VROC performance.

Improved IOPS and Reduced Latency at All Supported RAID Levels

With a mixed workload (70:30 read/write ratio, random 4 KB block size), Intel VROC delivers excellent performance in terms of IOPS and latency. The test was performed using high-performance, low-latency Intel® Optane™ SSDs to reduce the effect of SSD limitations on the test results. As shown in Figure 3, Intel VROC delivers up to 5.6 million IOPS in RAID0. Overall, Intel VROC exceeds HBA in IOPS by 161%, 135%, 50% and 245% in RAID 0, RAID 1, RAID 5 and RAID 10, respectively. The latency results for Intel VROC are lower than those from HBA in all RAID levels as well.¹

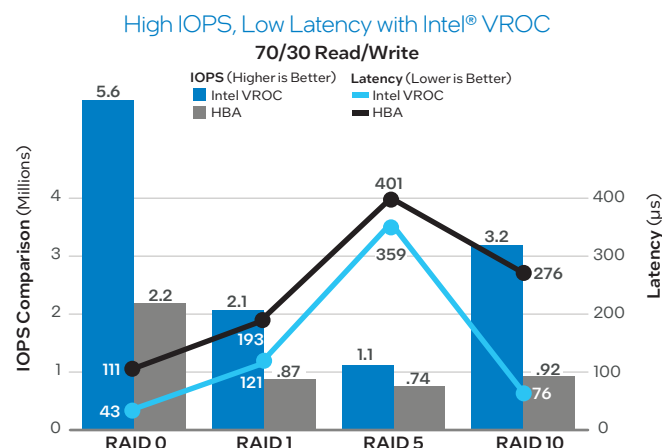


Figure 3. Compared to HBA-based RAID, Intel® VROC provides much higher IOPS and lower latency.¹

Up to 25% More Throughput with Similar CPU Utilization (RAID 5)

During conversations, customers have asked about the impact on CPU utilization compared to HBA. Our analysis shows that RAID 5 is the most CPU-intensive RAID operation among other RAID configurations; the results shared here are for RAID 5. The other RAID levels use the same or even less CPU resources compared to the HBA.

Intel VROC takes advantage of powerful Intel Xeon Scalable processors while HBAs rely on fixed RISC-based processors with limited onboard resources. With this difference in system architecture, the CPU utilization comparison is more meaningful when considered together with I/O throughput and application response time.

We used FIO to investigate the CPU utilization using different I/O depths instead of queue depths (number of pending I/O requests); Figure 4 and Table 1 provide the data for the RAID 5 test. Our results are summarized here:²

- Compared to the HBA solution, Intel VROC delivers over 25% more throughput with only a small amount of CPU overhead (less than 2%).
- Intel VROC also provides more than 25% lower latency than HBA does for mixed workloads using increasing I/O depths. Lower I/O latency means faster application response time.

In short, Intel VROC uses a small percentage of CPU overhead while providing a better overall application experience.

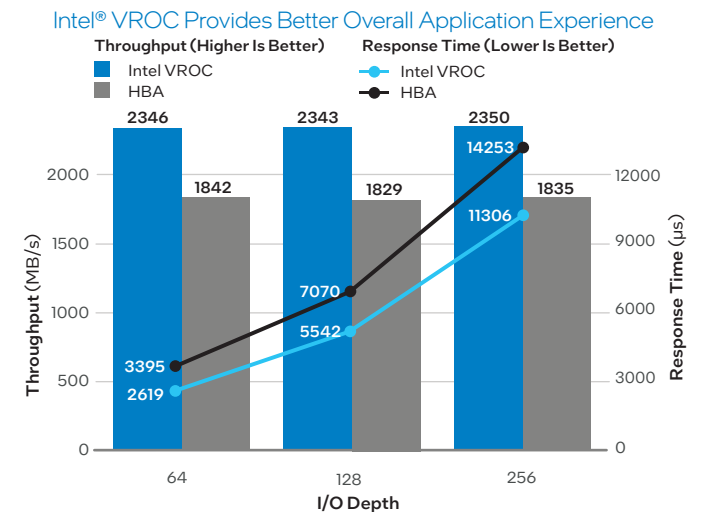


Figure 4. Intel® VROC delivers up to 25% more throughput with much better response time.²

Table 1. Comparison of CPU Utilization and Throughput²

I/O Depth	Available CPU Cycles	Throughput (MB/s)	Response Time (usec)
VROC			
64	96.54%	2,346	2,619
128	96.56%	2,343	5,542
256	96.22%	2,350	11,306
HBA			
64	98.08%	1,842	3,395
128	98.07%	1,829	7,070
256	97.79%	1,835	14,253

Up to 16% More Transactions per Second with Similar CPU Utilization

Similarly, we also collected performance data from the MySQL OLTP database. This real-world test showcases how Intel VROC can make a demonstrable difference in the quality of business experience and services, since they depend heavily on database application performance. As shown in Figure 5 and Table 2, Intel VROC demonstrates a similar or higher performance in transactions per second (TPS) over HBA while maintaining similar or smaller amount of CPU overhead. For example, when the thread count reaches 256 (equivalent to the number of user requests in parallel), Intel VROC delivers 16% more TPS at the cost of only 2% extra CPU resources.³

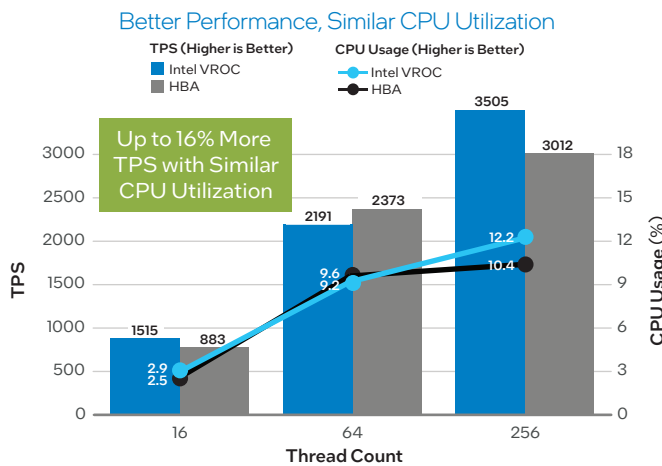


Figure 5. Intel® VROC shows similar or higher performance in TPS over HBA and maintains a comparable amount of CPU overhead.³

Table 2. Comparison of CPU Utilization and TPS³

Thread Count	Transaction Per Second (TPS)		CPU Usage (%)	
	VROC	HBA	VROC	HBA
16	883	774	2.9	2.5
64	2191	2373	9.2	9.6
256	3505	3012	12.2	10.4

Conclusion

Compared to RAID HBAs, Intel VROC’s integrated approach to RAID provides both improved performance and RAS capabilities that scale efficiently with today’s demanding workloads from edge to cloud. Replacing HBAs with Intel VROC helps organizations take full advantage of high-performing NVMe SSDs, which can improve business responsiveness with better application experience.

Learn More

You may find the following resources helpful:

Intel VROC Information

- [Web page](#)
- [Product brief](#)
- [Detailed Comparison to RAID HBA](#)
- [SKUs and Licensing](#)
- [User Guide](#)
- [FAQ](#)
- [Video of how to configure Intel VROC](#)

Other Intel Resources

- [Intel® VMD home page](#)
- [Intel® Xeon® Scalable processors](#)

Find the solution that is right for your organization. Contact your Intel representative or visit intel.com/VROC.



¹ Performance results are based on testing by Intel as of June 25, 2021 and may not reflect all publicly available updates. **System configuration:** Intel® Server Board M50CYP2SB2U/ M50CYP2SBSTD (chassis M50CYP2UR208BPP), 2x Intel® Xeon® Platinum 8358 processor (2.60 GHz, 32 cores each), DRAM 128 GB. **BIOS version (released April 2, 2021):** SE5C6200.86B.0020.P24.210402081 (microcode 0x0D000280). **OS:** Red Hat Enterprise Linux 8.1, kernel-4.18.0-147.el8.x86_64, mdadm - v4.1 - 2018-10-01, Intel® VROC Pre-OS version 7.5.0.1152. **Storage:** Both configurations used 4x 400 GB Intel® Optane™ SSD P5800X PCIe Gen4 U.2 (model: SSDPF21Q 400 GB, firmware: L0310100) connected to backplane, which is connected via SlimSAS cables directly to a Broadcom 9560-16i (x8) card on Riser 2, PCIe slot 1 on CPU2. **BIOS setting:** SpeedStep (Enabled), Turbo (Enabled), ProcessorC6 (Enabled), PackageC-State (C0/C1 State), CPU_PowerAndPerformancePolicy (Performance), HardwareP-States (NativeMode), Workload configuration (I/O Sensitive). **RAID configuration:** 4-disk RAID0/5/10 and 2-disk RAID1 with Intel VROC and Broadcom MegaRAID 9560-16i. **Workload generator:** FIO 3.25, 16-thread 16-I/O depth.

² 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® Server Board M50CYP2SB2U/ M50CYP2SBSTD (chassis M50CYP2UR208BPP), 2x Intel® Xeon® Platinum 8358 processor (2.60 GHz, 32 cores each), DRAM 128 GB. **BIOS version (released March 22, 2021):** SE5C6200.86B.0022.D08.2103221623 (microcode 0x0D000280). **OS:** Red Hat Enterprise Linux 8.1, kernel-4.18.0-147.el8.x86_64, mdadm - v4.1 - 2018-10-01, Intel® VROC Pre-OS version 7.5.0.1152. **Storage:** Both configurations used 4x 3.84 TB SSD D7-P5510 Series (model: SSDPF2KX03BTZ, Firmware: JCV10016) connected to internal backplane. With Intel VROC config, backplane connects directly to CPU2 via SlimSAS. With RAID HBA, backplane connects to RAID HBA on Riser 2, PCIe slot 1 on CPU2. **BIOS setting:** SpeedStep (Enabled), Turbo (Enabled), ProcessorC6 (Enabled), PackageC-State (C0/C1 State), CPU_PowerAndPerformancePolicy (Performance), HardwareP-States (NativeMode), Workload configuration (I/O Sensitive). **RAID configuration:** 4-disk RAID0/5/10 and 2-disk RAID1 with Intel VROC and Broadcom MegaRAID 9560-16i. **Workload generator:** FIO 3.25.

³ Performance results are based on testing by Intel as of February 3, 2022 and may not reflect all publicly available updates. **System configuration:** 2x Intel® Xeon® Platinum 8358 processor (2.60 GHz, 32 cores each)—XCC CPUs, QDF: QWMU, Stepping: D1, DRAM 128 GB (16x 8 GB DDR4 3200 MT/s). **BIOS version:** SE5C620.86B.01.01.0003.2104260124 (microcode 0x0D000280). **OS:** Red Hat Enterprise Linux 8.1, kernel: 4.18.0-147.el8.x86_64. **Storage:** 4x 15.36 TB P5316 QLC U.2 SSDs (Broadcom 9560-16i HBA Card) and 4x 15.36 TB SSD D5-P5316 QLC U.2 (Intel® Virtual RAID on CPU); Preferred VROC/mdadm driver version: RHEL8.1 inbox mdadm driver (mdadm-4.1-9.el8.x86_64). **RAID configuration:** Broadcom 4DR5 (chunk size: 256 KB (default) + 8 GB integrated cache on HBA card, Intel VROC 4DR5 (chunk size: 8 KB). **Workload generator:** Sysbench (sysbench-1.0.20-5.el8.x86_64); DBMS edition/ version: MySQL Community Edition 8.0.23.