Improve performance and minimize latency for IO-intensive apps by pairing Intel NVMe SSDs with Intel Virtual RAID on CPU (VROC)

Compared to a Broadcom MegaRAID 9460-16i Tri-Mode Storage Adapter

A study commissioned by Intel Corp.

Keeping business-critical information safe from data loss is a critical part of data center management. By using RAID volumes (redundant array of independent disks), you can store copies across multiple drives in case a drive failure occurs. Dedicated RAID controller cards were once key to enabling data loss protection without slowing performance to a crawl, but they are no longer the only available option.

The new Intel® Xeon® Scalable processors offer Intel Virtual RAID on CPU (Intel VROC), a built-in way to manage RAID volumes that doesn't require a dedicated storage adapter. According to Intel, Intel Volume Management Device (VMD) technology, a new hardware architecture on the processor, enables NVMe SSDs to connect directly to the processor, which allows VROC to better utilize the PCIe bandwidth of NVMe SSDs.

Up to **7.3%** better performance vs. Broadcom using Intel SSD DC P4510 Series

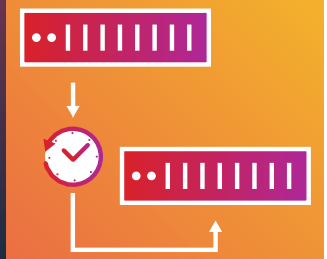

Upgrade drives from SATA to PCIe NVMe for **3.9x** the performance**

*Using the geomean of all user counts
**Comparison of Intel VROC w/ Intel SSD D3-S4510 Series vs. Intel VROC w/ Intel SSD DC P4510 Series.

In the PT data center, we explored two factors that can affect performance for IO-intensive apps: 1) RAID management and 2) drive type. We found that using Intel VROC gave a performance edge while running a large OLTP database over the Broadcom® adapter on all three drive types we tested, and that upgrading drives from Intel SSD D3-S4510 Series SATA drives to the latest Intel SSD DC P4510 Series NVMe SSDs nearly quadrupled the transactions per minute the system could process.

If you're running IO-intensive apps such as large MySQL™ databases, pairing premier Intel NVMe SSDs with Intel VROC could help you achieve higher performance and lower latency than using the Broadcom MegaRAID 9460-16i.

Get built-in RAID management with Intel VROC

Massive data growth drives modern business to seek faster and better ways to process and access that data, which makes powerful in-server storage vital. RAID volumes offer essential data protection from disk failures but require hardware or software that lets these drives act in concert. Traditionally, dedicated storage adapter cards have handled this task, but they can require additional expense and create complexity that could limit storage performance. With large IO-intensive datasets, the goal is to eliminate as many hurdles to fast performance and low latency as possible.

Intel VROC works with new Intel Xeon Scalable processors, combining hardware and software RAID in a hybrid approach that connects NVMe drives directly to the CPU without going through an additional storage adapter. Other potential benefits include the following:

- Intel VROC allows admins to connect and create a RAID volume with NVMe drives using the PCIe bus without the need for a dedicated RAID controller, potentially reducing cost.
- Intel VROC supports up to 12 CPU-connected and up to 48 switch-connected NVMe devices per processor, can span over multiple PCIe slots, and enables you to use the entire PCIe bandwidth of each drive instead of being limited to the link speed of a single PCIe RAID controller.
- Intel VROC lets you use multiple RAID configurations in a single drive set (e.g., RAID 5 and RAID 10 volumes on the same drives), which gives you the flexibility to use your storage most efficiently.

To learn more about what Intel VROC offers, visit https://www.intel.com/content/www/us/en/software/virtual-raid-on-cpu-vroc.html.

About the Intel SSDs we tested

Baseline: Intel SSD D3-S4510 Series SATA drives

If your data center still uses mostly hard disk drives, upgrading to the Intel SSD D3-S4510 Series can improve read performance because accessing the data doesn't require a spinning platter as in HDDs. The Intel SSD D3-S4510 Series, which uses the SATA 6Gb/s interface, comes in a 2.5-inch form factor and has a capacity of up to 7.68 TB.[1]

Upgrade: Intel SSD D5-P4320 Series NVMe drives

Intel SSD D5-P4320 Series drives are QLC 3D NAND drives that connect via the PCIe interface, which speeds performance compared to using SATA. According to Intel, these drives can "enable 3x storage consolidation compared to hard disk drives (HDDs), leading to lower operational costs."[2]

Upgrade further: Intel SSD DC P4510 Series NVMe drives

For heavily IO-intensive applications and top read performance, the Intel SSD DC P4510 Series PCIe NVMe drives offer capacity of up to 8 TB and built-in encryption for data security. Intel states that these drives, which use TLC and 3D NAND technology, are optimized for cloud storage, offering high capacity, density, and reliability features to minimize downtime.[3]

In our labs: Increased performance with Intel VROC vs. Broadcom

Making the most efficient possible use of your hardware investment makes good business sense. During our HammerDB tests running an OLTP workload with a range of user counts, the server using Intel VROC for RAID management improved overall performance over the Broadcom adapter, delivering a 7.3 percent better geometric mean of the transactions per minute (TPM) across all tested virtual user counts with the Intel SSD DC P4510 Series PCIe NVMe drives. The Intel SSD D5-P4320 Series NVMe and Intel SSD D3-S4510 Series SATA drives also showed slight performance advantages over the Broadcom adapter, with TPM geometric means of 5.3 percent and 9.9 percent improvement, respectively.

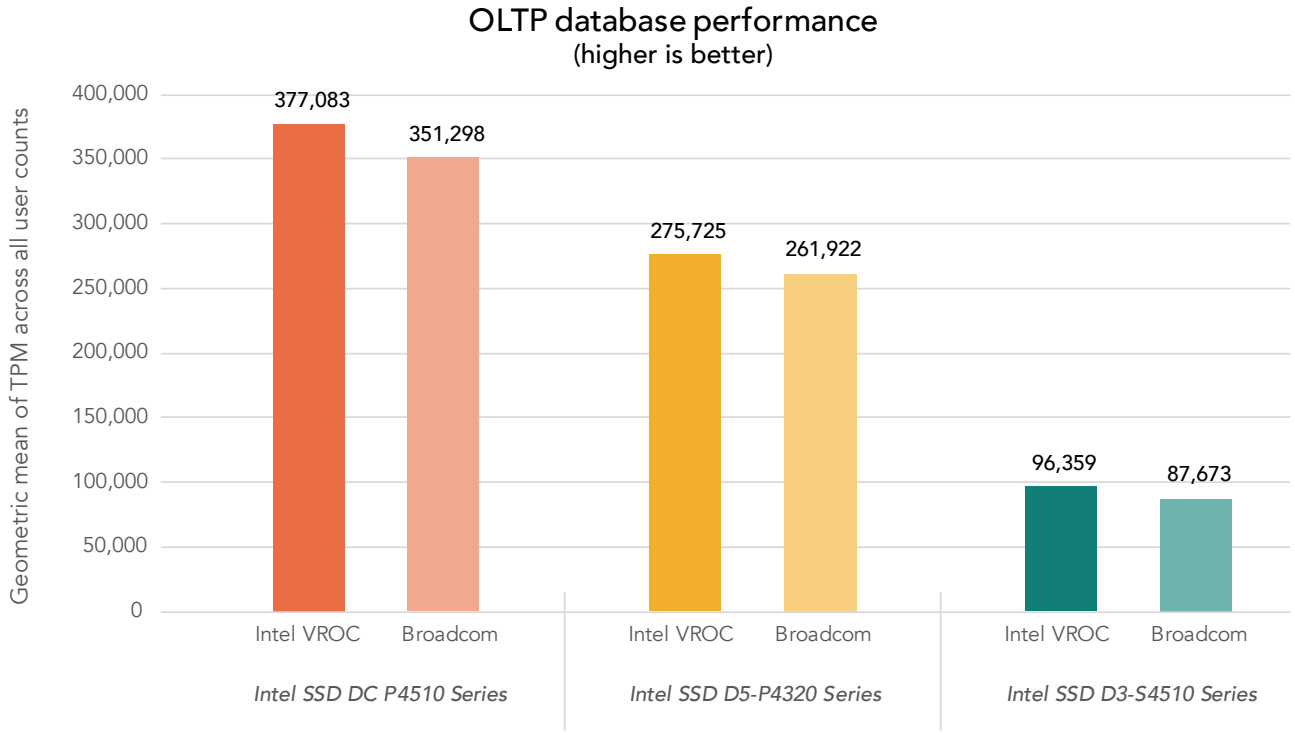

Drive choice also affects performance, particularly with workloads requiring high read I/O capabilities, like when many users require access to large databases. Using NVMe drives allows Intel VROC technology to deliver significant boosts in these read-heavy, high-I/O scenarios when compared to SATA. Intel offers NVMe drive options to help fit specific needs: the Intel SSD DC P4510 Series offers premium performance, while the Intel SSD D5-P4320 Series delivers some benefits of NVMe at a lower price point. Both NVMe drive options showed significant performance improvements compared to the Intel SSD D3-S4510 Series SATA drives, offering 3.9 times and 2.8 times the OLTP database performance (using the geometric mean of the TPM across user counts). Additionally, our tests showed that as we increased the number of virtual users, the performance differences in TPM between P4510, P4320, and S4510 drives continued to increase until each drive approached its respective performance limit. The Intel VROC configuration delivered the greater TPM on all three drives, with the Intel SSD DC P4510 Series drives delivering the best performance at 471,673 TPM, compared to 391,935 TPM and 113,079 TPM for the Intel SSD D5-P4320 Series and Intel SSD D3-S4510 Series drives, respectively.

Using geometric means

Because the difference in performance data points varied greatly over increasing numbers of virtual users, taking a simple average could mask lower virtual user data points and skew results unfairly toward higher virtual user data points. Instead, we took the geometric mean to fairly represent performance trends across all virtual user counts throughout the course of the tests.

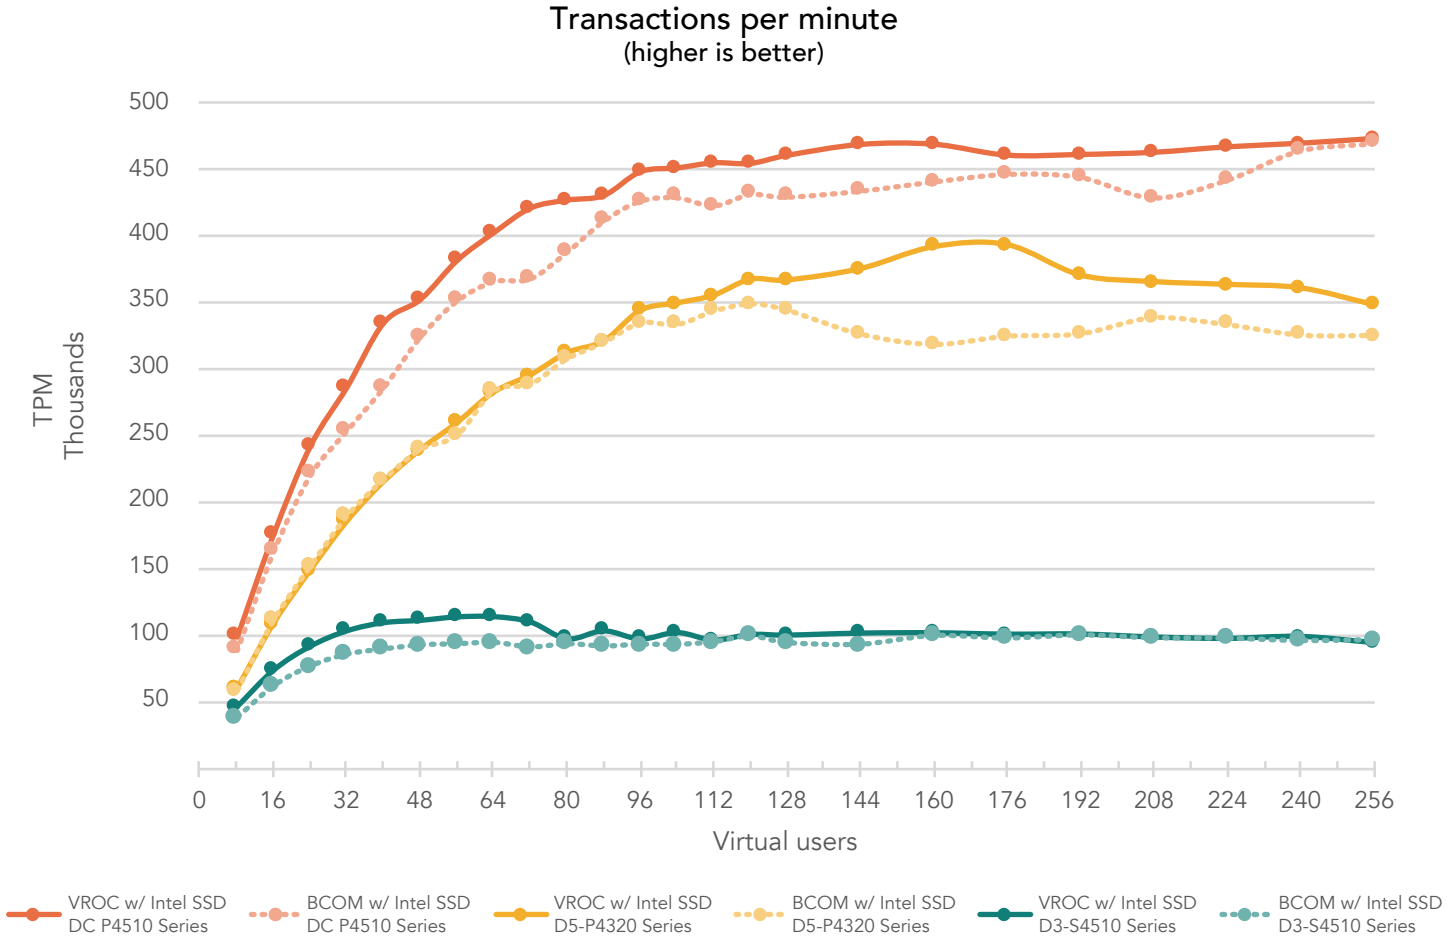

Alongside strict performance numbers comes latency, or wait times. The less latency, the faster the access to data. In certain latency-sensitive use cases, a requirement to maintain ultra-low latencies could force administrators to throttle I/O performance to keep latency at acceptable levels. When we sampled latency on the MySQL procedures measured by the test harness for a single virtual user across three virtual user counts in all drive types, Intel VROC delivered comparable or lower latency compared to the Broadcom solution.

The graph below shows the geometric mean latency for each RAID solution and drive combination at 64 virtual users. This calculation includes the average latencies of the five different functions that a virtual user performed throughout the test (giving them each equal weight). While we ran tests at many virtual user counts, for application latency measurements, we used the 64, 96, and 128 user counts to show a healthy sample for latency analysis at higher user counts. Across the three virtual user count latencies we captured, Intel VROC delivered between 12.6 and 14.4 percent lower latency on the Intel SSD D3-S4510 Series configuration and between 4.4 and 9.5 percent lower latency on the Intel SSD DC P4510 drive configuration. The two solutions showed comparable latencies on the Intel SSD D5-P4320 Series drive configuration, with a difference of less than 2.7 percent between the two. This was the only instance where latency was lower on the Broadcom solution; at both 96 and 128 users, latency was slightly lower on the VROC solution.

For more detailed performance data, including the geometric mean latencies at 96 and 128 virtual user counts we captured, see the science behind the report.

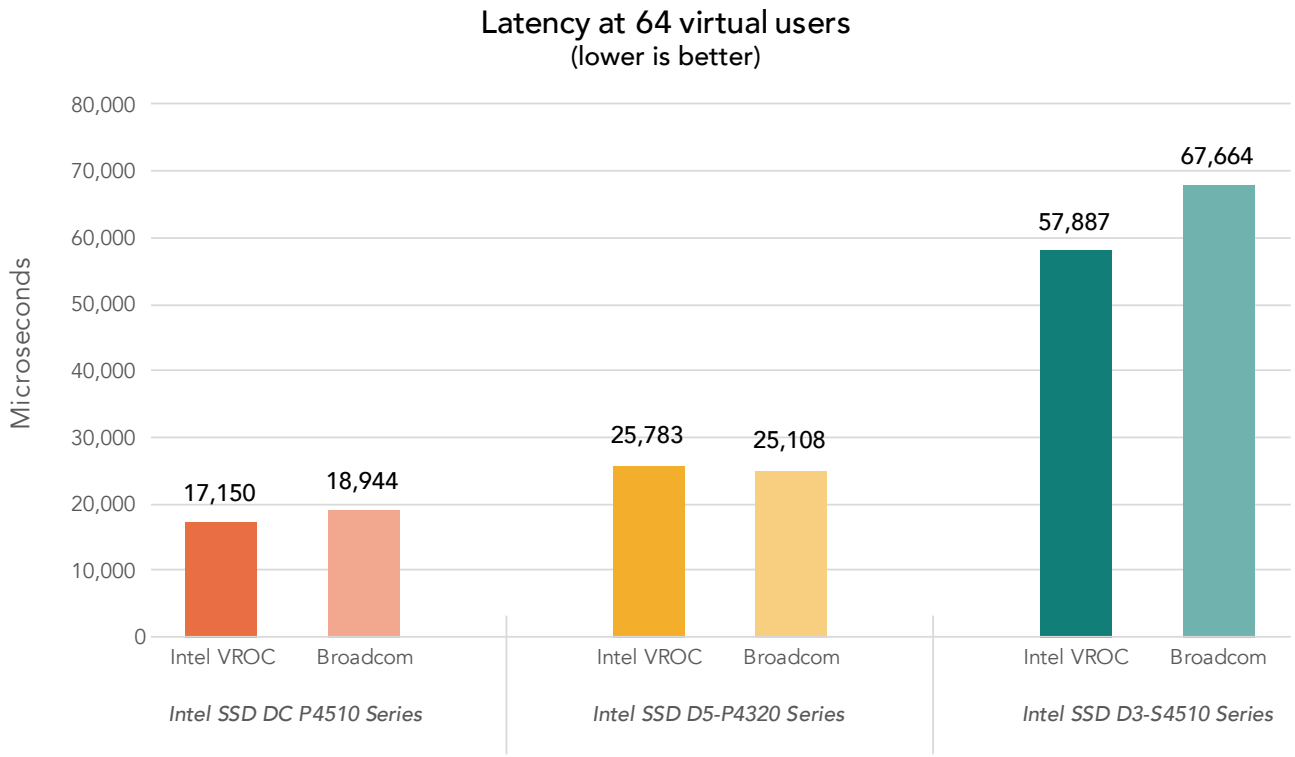

About our tests

To test the solutions, we used HammerDB, a TPC-C-like benchmark (see sidebar). TPC-C is a benchmark maintained by the Transaction Processing Council (TPC) that simulates an order-entry environment for a wholesale supplier. The simulated order application tasks in our tests included new orders, tracking delivery, processing payments, querying the order status, and checking stock levels. The selected warehouse count dictates the size of the database. As warehouses increase, the sales districts increase, and in turn, the number of customers increase. For more information on the TPC-C benchmark model, see the TPC site at http://www.tpc.org/tpcc/detail.asp. (Note: Our test was not an officially audited TPC benchmark and therefore cannot be compared to official TPC-C results.)

Our test configuration used 100,000 warehouses, resulting in a large database of nearly 9TB on the disks. We configured MySQL to use a buffer pool of 96GB out of the available 192GB of system memory. This resulted in a read/write IO profile of approximately 80 percent read and 20 percent write. We share results for our sample workload only; your workload parameters may differ.

We created two RAID volumes on each solution. On the VROC solution, we created a RAID5 volume for database data, and a RAID10 volume for database logs. Both volumes were distributed across the four disks; the VROC solution did not use a write back cache due to the nature of VROC itself.

Broadcom did not allow for different RAID types on the same underlying disks, so we configured two RAID5 volumes on Broadcom: one for database data, and one for database logs. We disabled the write cache on the Broadcom database volume, and we enabled the write cache on the Broadcom logs volume.

About HammerDB

We tested each solution with an OLTP workload from the HammerDB suite of benchmarks. The HammerDB transactional TPC-C-like benchmark gives results in transactions processed per minute (TPM). Our tests were not official TPC results and are not comparable in any manner.

To learn more about TPM and other information specific to our benchmark workloads, visit the HammerDB website at www.hammerdb.com.

Conclusion

If you're purchasing servers powered by new Intel Xeon Scalable processors, don't forget that a new feature exists that could improve performance for IO-intensive apps and reduce complexity by connecting RAIDed NVMe drives directly to the CPU: Intel Virtual RAID on CPU. On our large OLTP database workload, using Intel VROC improved performance and provided comparable or lower latency compared to using a Broadcom MegaRAID 9460-16i Tri-Mode Storage Adapter.

This was true for the one SATA drive type and two NVMe drive types we tested, with the premier Intel SSD DC P4510 Series PCIe NVMe drives with VROC accomplishing up to an average of 3.9 times as many TPM as the Intel SSD D3-S4510 Series SATA drives with VROC and 4.3 times the TPM of the Intel SSD D3-S4510 Series SATA drives with the Broadcom adapter. The new QLC-NAND-based Intel SSD D5-P4320 Series delivered up to 3.1 times as many TPM as the Intel SSD D3-S4510 Series SATA drives and came within 30 percent of TLC-NAND-based Intel SSD DC P4510 drives. Because it offers NVMe performance benefits over SATA SSDs at a lower price point than the premium Intel SSD DC P4510 Series, the Intel SSD D5-P4320 Series offers a balance of price and performance that many organizations may find compelling. By using Intel VROC and upgrading your internal storage, you could allow IO-intensive workloads to access data faster.

1 Intel SSD D3-S4510 and D3-S4610 Series Product Brief, accessed March 3, 2019, https://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/data-center-ssds/dc-d3-s4510-s4610-series-brief.html.

2 Intel SSD D5-P4320 Product Brief, accessed March 3, 2019, https://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/data-center-ssds/d5-p4320-series-brief.html.

3 Product Brief: Intel SSD DC P4510 Series, accessed March 3, 2019, https://www.intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/data-center-ssds/dc-p4510-series-brief.html.

Read the science behind this report at http://facts.pt/pjfyjap ▶

Principled Technologies®

Facts matter.®

This project was commissioned by Intel.