Solution Design Brief

Data Centers | Infrastructure Modernization

Intel Accelerated Solution

Boost Performance and Lower Latency with VMware vSAN 8 and 4th Gen Intel® Xeon® Scalable Processors

Business Challenge: Are your servers providing enough compute and storage performance? Do you need to consolidate your server footprint? Outdated hardware and software can struggle to support modern workloads; may increase compliance and security risks; and can raise management complexity and OpEx.

45%

Lower TCO

Lower total cost of ownership by more than 45% with CapEx and OpEx savings.

Solution Overview and Summary

Solution: Simultaneously upgrading your data center's hardware and software can help solve two of the greatest IT challenges that drive modernization: infrastructure cost and OpEx along with sub-par legacy infrastructure performance. Latest-generation VMware vSphere and VMware vSAN 8 software running on 4th Generation Intel® Xeon® Scalable processors make it easier to transition to a modern hybrid/multicloud environment and can provide operational, performance, and cost efficiencies.

Higher-performance compute and storage in a hyperconverged infrastructure (HCI) environment can help you achieve compelling business outcomes, including:

- Support memory- and compute-hungry workloads and modern applications.
- Enhance security and reduce risk through fast encryption as well as other data protection features.
- Enhance customer experience by responding to customer requests quickly.
- Consolidate servers to reduce data center footprint and power consumption for a more sustainable infrastructure.
- Uncover in-depth insights that enable innovation and better decision making.
- Lower total cost of ownership by more than 45% with CapEx and OpEx savings.[1]

The testing highlighted in this document was conducted in February and March 2023 by Intel engineers to demonstrate the performance enhancements enabled by the combination of the newest Intel® technology with the new Express Storage Architecture (ESA) introduced in VMware vSphere/vSAN 8. ESA is an optional, alternative architecture in vSAN designed to process and store data with elevated levels of efficiency, scalability, and performance.

Results: Upgrading hardware and software together provided up to 6x higher throughput and up to 7x lower latency, compared to solutions running on previous generations of Intel-based servers and vSAN software.[2]

Test Methodology

We used VMware Hyperconverged Infrastructure Benchmark (HCIBench[3]) to measure throughput and latency. HCIBench fully automates the end-to-end process of deploying test virtual machines (VMs), coordinating workload runs, aggregating test results, and collecting necessary data for troubleshooting purposes. We tested with a range of block sizes, read versus write ratios, and random versus sequential operations to simulate a range of popular workloads that can run on vSAN (see "Workload Profiles" for details).

HCIBench Setup Details

- The HCIBench appliance was hosted outside of the cluster under test.
- Virtual disk was zeroed-out before testing.
- Test duration was 30 minutes with a 300-second warmup.
- FIO latency target mode was used to measure throughput. FIO finds the maximum point at which the given workload will run while maintaining a latency below our target of 10ms. For latency measurement, we used a generated FIO static controlled workload.

Results

Figure 1 illustrates the performance gains provided by upgrading from older infrastructure to newer. We tested four configurations. Configurations 1–3 used VMware vSphere/vSAN 7 with 1st, 2nd, and 3rd Gen Intel Xeon Scalable processors. Configuration 4 used vSphere 8 and the new vSAN 8 ESA with 4th Gen Intel Xeon Scalable processors. The hardware-plus-software advancements achieved up to 6x higher I/O throughput and up to 7x lower latency per vSAN cluster than older hardware and software configurations.[4] These two metrics are examples of the overall performance boost you can expect from up-to-date, modern infrastructure.

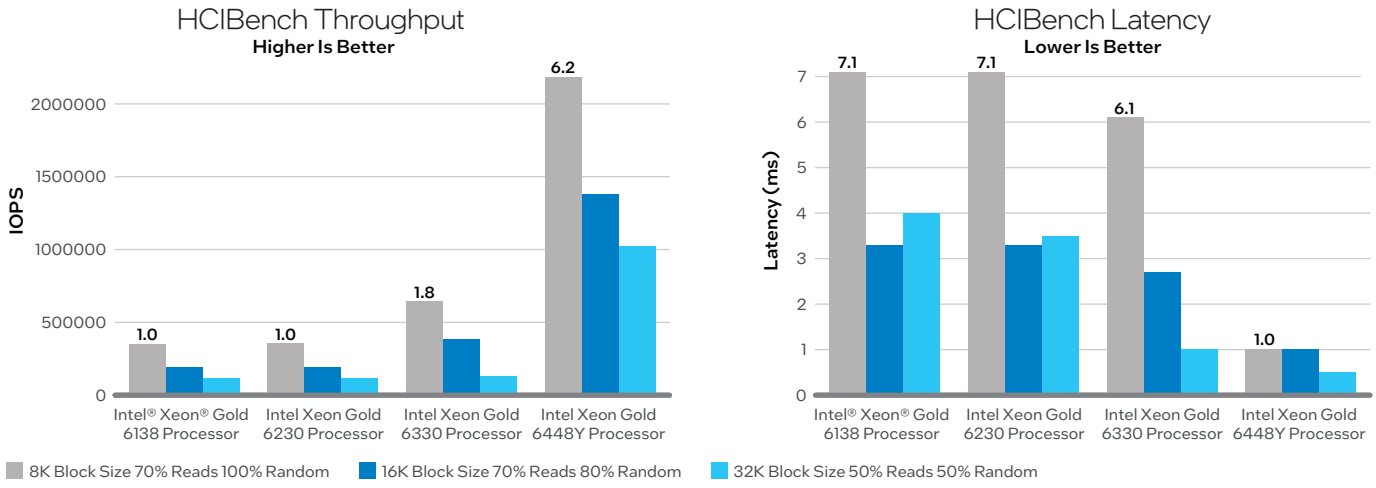

Figure 1. Obtain significant performance improvements by upgrading both hardware and software.

Configuration Details

The following tables provide information about components and settings of the infrastructure used during testing. These tables reflect typical types of configurations used for each generation of processor.

Component	Hardware Configurations			
	Configuration 1 VMware vSphere/vSAN 7 with 1st Gen Intel® Xeon® Scalable Processor	**Configuration 2 VMware vSphere/vSAN 7 with 2nd Gen Intel® Xeon® Scalable Processor**	**Configuration 3 VMware vSphere/vSAN 7 with 3rd Gen Intel® Xeon® Scalable Processor**	**Configuration 4 VMware vSphere 8/vSAN ESA with 4th Gen Intel® Xeon® Scalable Processor**
Server	4x Dell PowerEdge R640	4x Dell PowerEdge R640	4x Intel® Server Board M50CYP2SBSTD	4x Intel Server Board M50FCP2SBSTD
Processor (per node)	2x Intel Xeon Gold 6138 processor (20 cores, 2.0 GHz base/2.7 GHz All-Core Turbo[a])	2x Intel Xeon Gold 6230 processor (20 cores, 2.1 GHz base/2.8 GHz All-Core Turbo[a])	2x Intel Xeon Gold 6330 processor (28 cores, 2.0 GHz base/2.6 GHz All-Core Turbo[a])	2x Intel Xeon Gold 6448Y processor (32 cores, 2.1 GHz base/3.0 GHz All-Core Turbo[a])
Memory	384 GB (12x 32 GB DDR4 2666 MHz)	384 GB (12x 32 GB DDR4 2933 MHz)	512 GB (16x 32 GB DDR4 2933 MHz)	512 GB (16x 32 GB DDR5 4800 MHz)
Storage	**Cache:** 2x Intel® Optane™ SSD DC P4800X 375 GB **Capacity:** 6x Solidigm SSD DC-P4510 2 TB	**Cache:** 2x Intel Optane SSD DC P4800X 375 GB **Capacity:** 6x Solidigm SSD DC-P4510 2 TB	**Cache:** 2x Intel Optane SSD P5800X 400 GB **Capacity:** 6x Solidigm D7-P5510 SSD 3.84 TB	**Capacity (flat tier):** 9x Solidigm D7-P5510 SSD 3.84 TB
Network Card	1x Intel® Ethernet Network Adapter X710 Dual-Port 10 GbE	1x Intel® Ethernet X710 Network Adapter Dual-Port 10 GbE	1x MT27800 Family, Dual-Port at 25 GbE with RDMA over Converged Ethernet version 2 (RoCEv2)	1x Intel Ethernet Network Adapter E810-CQDA2, Dual-Port at 100 GbE with RoCEv2
Network Switch	Arista DCS-7060CX-32S-R	Arista DCS-7060CX-32S-R	Arista DCS-7170-64C-R	Arista DCS-7170-64C-F

[a] **All-Core Turbo Boost frequency range:** Intel® Turbo Boost Technology automatically provides opportunistic performance improvement by allowing individual cores to operate at a higher frequency. This results in higher frequency in both single- and multi-threaded applications when headroom is available. For more information on the All-Core Turbo Boost frequency range, see section 2.1 and 2.3 at https://builders.intel.com/docs/networkbuilders/power-management-technology-overview-technology-guide.pdf.

Software Versions	Configurations 1-3	Configuration 4
Hypervisor	vSphere ESXi 7.0U3G, 20328353	vSphere ESXi 8.0GA, 20513097
vSAN Mode	OSA, default RAID1 (2-disk group) storage policy	ESA, default policy RAID5
Benchmark Tools	HCIBench 2.8, storage engine FIO 3.3	HCIBench 2.8, storage engine FIO 3.3

Accelerator Technologies Enabled
Intel® Speed Select Technology Performance Profile 2.0[a]
Intel® Hyper-Threading Technology
Intel® Turbo-Boost Technology

[a] Refer to https://builders.intel.com/docs/networkbuilders/power-management-technology-overview-technology-guide.pdf.

Important System Settings	
Number of Nodes	4
Power and Performance Policy	Performance
Frequency Governor	Native (OS control)
Max C-State	c0/c1
Prefetchers	L2 HW, L2 Adj., DCU HW, DCU IP
NUMA	Enabled (no sub-NUMA clustering)
IRQ Balance	Enabled
RDMA	Configurations 1-2 Not Applicable Configurations 3-4 RoCEv2 Enabled

Workload Profiles

The following table describes the various workload profiles used during testing.

Workload	# of VMs	vCPUs	RAM	I/O Size	# Data Disks	Size of Disk	Read Percentage	Random Percentage	Latency Target Mode
Throughput Mode									
	16	4	8 GB	8K	4	50 GB	70%	100%	10 ms
	16	4	8 GB	16K	4	50 GB	70%	80%	10 ms
	16	4	8 GB	32K	4	50 GB	50%	50%	10 ms
Latency Mode									
	16	4	8 GB	8K	8	50 GB	70%	100%	n/a
	16	4	8 GB	16K	8	50 GB	70%	80%	n/a
	16	4	8 GB	32K	8	50 GB	50%	50%	n/a

Hardware + Software Upgrade = Win-Win for Data Centers

Improved workload efficiency through faster performance can enable server consolidation and lower power, to help meet overall data center efficiency and sustainability goals. The performance improvements revealed by our testing are a direct result of the optimizations from both Intel® hardware and VMware software, as described here:

Hardware

- **Processor.** Each generation of Intel® Xeon® Scalable processor not only offers more cores, but improvements in per-core performance, enabling more parallel processing of workloads. That is, new generations can process more instructions per clock cycle. Applications running in VMs on vSphere 8 can use the newest Intel® accelerators for AI, compression, encryption, and confidential computing. In addition, upgrading the software can help ensure regulatory compliance.
- **Memory.** Memory bandwidth has increased, and older systems support only slower speeds of DRAM (such as DDR4). With the newer systems, you have access to current memory technologies, such as DDR5, which delivers more MT/s.
- **Storage.** Over time, I/O bandwidth has increased and versions of the PCIe standard have advanced. Legacy systems work only with older generations of SSDs and other devices, while newer systems support the most recent generations of PCIe, such as Gen4 and Gen5 (however, we did not include Gen5 devices in our testing).
- **Network.** Legacy Ethernet controllers and cards are limited to 10 or 25 GbE. With the latest generation of Intel® Ethernet products, you can achieve up to 100 GbE, while intelligent offloads improve performance and VM scalability. You can also enable Remote Direct Memory Access (RDMA) over Converged Ethernet, version 2 (RoCEv2) to accelerate and streamline traffic between nodes in the vSAN cluster. RDMA can increase data center efficiency by enabling nodes in a vSAN cluster to exchange data in main memory without involving the CPU, thereby lowering network packet latency and reducing CPU utilization.

Software

- **vSphere.** vSphere 8 has new software capabilities and performance enhancements, including support for more VMs per cluster, additional ESXi hosts that can be managed by Lifecycle Manager, and more VM DirectPath I/O devices per host. These enhancements enable vSphere to support increasingly powerful workloads or applications and complete tasks quickly.
- **vSAN.** vSAN 8 and prior versions offer vSAN in its Original Storage Architecture (OSA). A new optional offering of vSAN 8 is Express Storage Architecture (ESA), which reimagines how data can be processed more efficiently in the storage stack. Some of the new ESA features that help to reduce write amplification and write latency include:
  - More performant, flexible, scalable and resilient architecture
  - Unlocks the capabilities of modern HW, including improved CPU efficiency
  - Simplified provisioning and operations

Conclusion

Modernizing the data center is a crucial component of the necessary ongoing digital transformation that helps companies compete in today's fast-paced business environment. Combining vSphere/vSAN 8 with 4th Gen Intel Xeon Scalable processors can help achieve the following business goals:

- Improve customer experiences, which in turn can help increase revenue.
- Drive innovation and sustainability.

Further Information

- 4th Gen Intel® Xeon® Scalable processors
- Intel® Ethernet 800 Series
- VMware vSphere 8
- VMware vSAN with ESA

Accelerate Digital Transformation

Support for Modernization

Future Ready Infrastructure

Enable Remote Workforce

Empower Developers

Reduce Complexity

Authors

Ewelina Kamyszek
Cloud Solutions Engineer
DCAI/CESG Intel Group

Patryk Wolsza
Cloud Solutions Architect, vExpert
DCAI/CESG Intel Group

Learn more about the Intel and VMware Partnership and Data Center solutions.

Contact your Intel representative to learn more about this solution.

Solution Provided By:

intel®

[1] https://www.vmware.com/content/dam/digitalmarketing/vmware/en/pdf/docs/vmw-lower-your-data-center-costs-with-vmware-vsan.pdf

[2] **Intel® Xeon® Gold 6138 processor configuration:** 4 nodes, 2x Intel Xeon Gold 6138 processor per node, 4x Dell PowerEdge R640 servers, total memory 384 GB (12x 32 GB DDR4 2666 MHz), Intel® Hyper-Threading Technology: Enabled, Intel® Turbo Boost Technology: Enabled, NUMA no SNC, Intel® VMD: N/A, BIOS: 2.17.1 (microcode:0x2006e05), storage (cache): 2x 375 GB Intel® Optane™ DC SSD P4800X Series, storage (capacity): 6x 2 TB Solidigm SSD DC P4510, network devices: 1x Intel® Ethernet Network Adapter X710, network speed: 10 GbE, OS/Software: VMware 7.0U3G, 20328353, vSAN OSA – default policy (RAID-1, 2DG). Test by Intel as of 03/13/2023 using VMware HCIBench 2.8, FIO 3.3. Throughput test mode: 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB); 16K profile (I/O size 16K, read percentage 70%, random percentage 80%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB); 32K profile (I/O size 32K, read percentage 50%, random percentage 50%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB). Latency test mode: 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB); 16K profile (I/O size 16K, read percentage 70%, random percentage 80%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB); 32K profile (I/O size 32K, read percentage 50%, random percentage 50%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB).

Intel Xeon Gold 6230 processor configuration: 4 nodes, 2x Intel Xeon Gold 6230 processor per node, 4x Dell PowerEdge R640 servers, total memory 384 GB (12x 32 GB DDR4 2933 MHz), Intel Hyper-Threading Technology: Enabled, Intel Turbo Boost Technology: Enabled, NUMA no SNC, Intel VMD: N/A, BIOS: 2.17.1 (microcode:0x4003303), storage (cache): 2x 375 GB Intel® Optane™ DC SSD P4800X Series, storage (capacity): 6x 2 TB Solidigm SSD DC P4510 Series, network devices: 1x Ethernet Network Adapter X710, network speed: 10 GbE, OS/Software: VMware 7.0U3G, 20328353, vSAN OSA – default policy (RAID-1, 2DG). Test by Intel as of 03/13/2023 using VMware HCIBench 2.8, FIO 3.3. Throughput test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB). Latency test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB).

Intel Xeon Gold 6330 processor configuration: 4 nodes, 2x Intel Xeon Gold 6330 processor per node, 4x Intel® Server Board M50CYP2SBSTD, total memory 512 GB (16x 32 GB DDR4 2933 MHz), Intel Hyper-Threading Technology: Enabled, Intel Turbo Boost Technology: Enabled, NUMA no SNC, Intel VMD: Enabled, BIOS: SE5C620.8 6B.01.01.0006.2207150335 (microcode:0xd000375), storage (cache): 2x 400 GB Intel® Optane™ SSD P5800X Series, storage (capacity): 6x 3.84 TB Solidigm SSD D7-P5510 Series, network devices: 1x MT27800 Family, at 25 GbE RoCE, network speed: 25 GbE, OS/Software: VMware 7.0U3G, 20328353, vSAN OSA – default policy (RAID-1, 2DG). Test by Intel as of 03/13/2023 using VMware HCIBench 2.8, FIO 3.3. Throughput test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB). Latency test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB).

Intel Xeon Gold 6448Y processor configuration: 4 nodes, 2x Intel Xeon Gold 6448Y processor per node, 4x Intel Server Board M50FCP2SBSTD, total memory 512 GB (16x 32 GB DDR5 4800 MHz), Intel Hyper-Threading Technology: Enabled, Intel Turbo Boost Technology: Enabled, NUMA no SNC, Intel VMD: Enabled, BIOS: SE5C741.8 6B.01.01.0002.2212220608 (microcode:0x2b000161), storage (data): 9x 3.84 TB Solidigm D7-P5510 Series, network devices: 1x Intel Ethernet E810-CQDA2 at 100 GbE RoCE, Network speed: 100 GbE, OS/Software: VMware 8, 20513097, vSAN ESA – Optimal default policy (RAID-5, flat). Test by Intel as of 03/13/2023, using VMware HCIBench 2.8, FIO 3.3. Throughput test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, latency target mode <10ms, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 4, size of disk 50 GB). Latency test 8K profile (I/O size 8K, read percentage 70%, random percentage 100%, #VMs per cluster 16, vCPU 4, vRAM 8, # data disks per VM 8, size of disk 50 GB).

[3] Source: https://flings.vmware.com/hcibench

[4] See endnote 2.

Performance varies by use, configuration, and other factors. Learn more on the Performance Index site.

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

Your costs and results may vary.