Best Practices Guide
Data Center | Hybrid Cloud

intel.

Deploying a Cost-Efficient, High-Performance vSAN Cluster

Discover how to improve storage performance, lower latency, improve scalability, and reduce total costs with a VMware vSAN-based solution optimized for Intel® Optane™ persistent memory and Intel® Optane™ SSDs

vmware®

Author

Flavio Fomin
Solutions Architect, Intel Corporation

John Hubbard
Solutions Architect, Intel Corporation

Table of Contents

Introduction

Software-defined storage (SDS) abstracts storage software from the storage hardware. By providing a shared pool of storage capacity that can be used across service offerings, SDS eliminates storage silos and helps improve utilization ratios. Intelligent, automated orchestration can reduce operating costs and speed provisioning from several weeks to a few minutes.

IT organizations can no longer afford legacy storage architecture that limits their ability to scale cost-effectively. Intel hardware coupled with VMware software is an effective alternative for modernizing and future-proofing storage. Intel® Optane™ SSDs combined with 3D NAND SSDs are a key differentiator for this architecture, delivering scalability, predictability, and beneficial cost/performance.

VMware vSAN is a vSphere-native SDS solution that powers industry-leading hyperconverged infrastructure (HCI) solutions in the hybrid cloud. Intel and VMware have worked closely to develop reference architectures and best practices for deploying a vSAN platform that is optimized to run on technology from Intel. Examples include the latest generation of Intel® Xeon® Scalable processors, Intel Optane SSDs, and Intel® Optane™ persistent memory (Intel® Optane™ PMem).

The steps and recommendations in this best practices guide can help you build a storage platform that can keep up with current needs as well as scale into the future.

Solution Overview

A VMware vSAN-based deployment consists of the vSAN software, cache, and capacity drives organized into disk groups.

vSAN Architecture Overview

vSAN abstracts and aggregates locally attached disks in a vSphere cluster to create a storage solution that can be provisioned and managed from vCenter and the vSphere Web Client. Each host in a vSAN cluster can contribute storage capacity to the cluster, as well as consume data from the cluster. All the storage devices combine to create a single vSAN datastore.

vSAN is an HCI solution, where VM storage and compute resources are delivered from the same x86 server platform running the hypervisor. As such, vSAN reduces the need for external shared storage, and simplifies storage configuration and VM provisioning. It integrates with the entire VMware software stack—vSAN is a distributed layer of software included in the VMware ESXi hypervisor. You can use VM-level policies to control VM storage provisioning and storage service-level agreements (SLAs). And you can set these policies per VM and modify them dynamically.

vSAN is a common solution for hybrid cloud deployments that need scale-out and virtualized compute, storage, and networking resources.

How Storage Works in vSAN

VMware vSAN is a powerful HCI platform that serves as a critical building block for the software-defined data center. Organizations can deploy vSAN to take advantage of the solution's distinctive scalability, security, and performance features for today's most demanding, storage-intensive data center workloads.

vSAN aggregates all local capacity devices into a single datastore shared by all hosts in the vSAN cluster. The datastore can be expanded by adding capacity devices or hosts with capacity devices to the cluster. vSAN works best when all ESXi hosts in the cluster share similar or identical configurations across all cluster members.

Figure 1 shows how data moves into and out of vSAN storage.

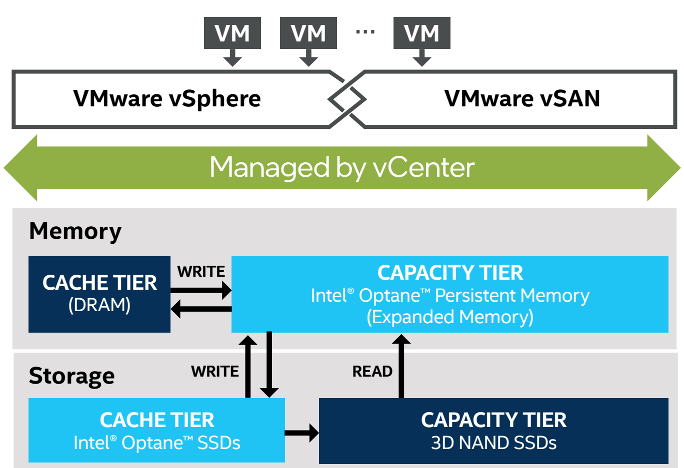

Figure 1. Data moving from memory into the vSAN cache tier and capacity tier.

Business Value

Intel Optane SSDs deliver excellent business value to vSAN deployments by offering high performance, low latency, low acquisition costs and footprint, and high reliability and availability.

Consistently High Performance

Because Intel Optane SSDs are a write-in-place media, they provide a significantly higher write throughput (measured in MB per second) than NAND flash SSDs. This technology difference enables up to 80 percent higher write throughput compared to typical NVMe NAND drives,[1] even though both the Intel Optane and Intel NVMe drives use PCIe Gen 4 (see upper portion of Figure 2). Higher throughput enables higher scalability. An SSD with a higher write throughput adds meaningful value to the vSAN two-tier architecture because all writes must go through the cache device.

Predictably Low Latency

Intel Optane SSDs use bit-level (not page-level) write operations, which significantly reduces time-consuming garbage collection and drives down latency (see lower portion of Figure 2).

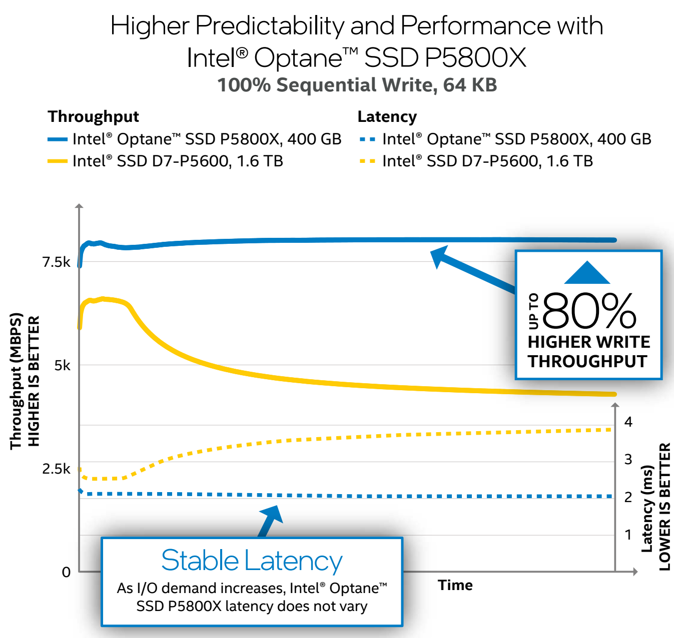

Figure 2. Intel® Optane™ SSD P5800X sustains performance under write pressure, enabling greater scalability and predictable performance.[1]

Response time remains low even when workload I/O demands spike unexpectedly. The predictable response time of Intel Optane SSDs means that even with volatile workloads, or more workloads per VM, organizations can continue to meet SLAs (Figure 3).

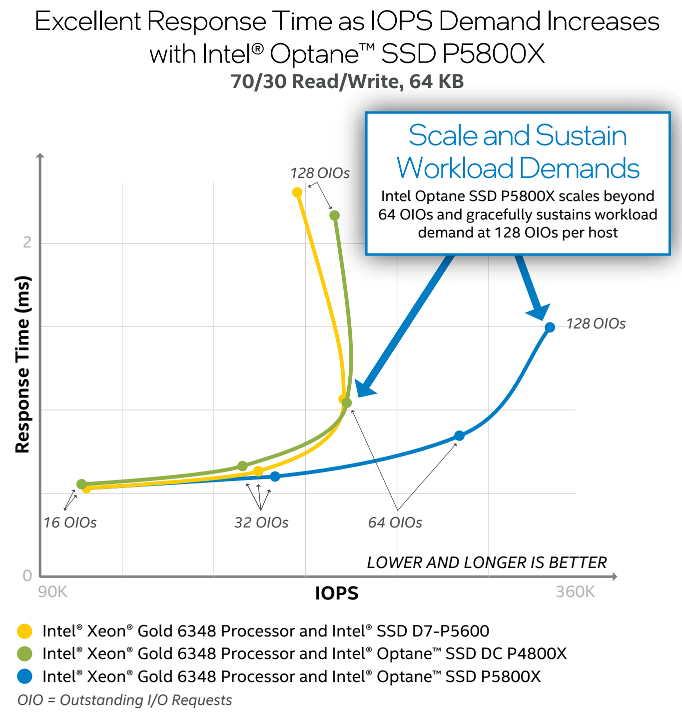

Figure 3. For 64 KB 70/30 read/write workloads, the Intel® Optane™ SSD P5800X can keep up with steadily growing IOPS demand—long after the Intel® SSD D7-5600 and even the Intel Optane SSD DC P4800X response times increase exponentially.[2]

Lower Acquisition Costs and Footprint

Using Intel Optane SSDs as cache enables a higher VM density per host. Fewer hosts per cluster help reduce hardware and software acquisition costs and consumed footprint. Intel Optane SSDs can reduce node count by up to 30 percent, when comparing the combination of a 3rd Generation Intel® Xeon® Scalable processor and the latest-generation Intel Optane SSD to the previous-generation processor and SSD.[3]

Higher Reliability and Availability

With an endurance of 100 DWPD, Intel Optane media does not run the risk of wearing out before the 5-year warranty, avoiding the exposure to data loss and unavailability related to cache drive replacements.

Solution and System Requirements

Table 1 provides the performance specifications for the Intel Optane SSD P5800X Series.

Table 1. Intel® Optane™ SSD P5800X Series Performance Specifications[4]

Feature	Specification
Capacity	400/800/1600 GB
Interface	PCIe x4
DWPD	100
Throughput	
– Sequential Read	Up to 7.4 GB/sec
– Sequential Write	Up to 7.4 GB/sec
– Random 4K Read (IOPS)	Up to 1.55 million
– Random 4K Write (IOPS)	Up to 1.6 million
– Random 4K 70/30 (IOPS)	Up to 2.0 million
– Random 512B Read (IOPS for metadata)	Up to 5.0 million
QoS	
– 4KRR, QD=1, 99%	< 6µs
– 4KRR, RW, Mixed QD=1, 99.999%	< 66µs

Memory and I/O Layout

Each 3rd Gen Intel Xeon Scalable processor supports eight memory channels. Each memory channel supports up to two DIMMs. Based on NUMA settings and memory population, these channels will be interleaved channels of a two-socket system. The system can have access to a maximum of 6 TB of memory, combining DRAM and Intel Optane PMem at 3200 MHz. The PCIe subsystem provides up to 128 lanes of high-speed I/O.

Memory Configurations

In an HCI environment, having enough memory is important to increase CPU utilization and drive up VM density. Memory capacity can be configured as DRAM-only or can be a tiered memory configuration, which is a combination of DRAM and Intel Optane PMem, where the DRAM capacity serves as a cache and the Intel Optane PMem provides the capacity consumed by the VMs. Intel Optane PMem 200 series in Memory Mode can provide up to a 20 percent platform-level cost reduction[5] and can provide comparable performance to DRAM-only[6] when the host's active memory fits in the DRAM capacity.

For optimal performance, populate all eight memory slots of the processors with DRAM DIMMs, and four or eight slots with Intel Optane PMem, according to the desired capacity. Refer to the following for memory population guidelines provided by Intel, VMware, and current-generation OEM solutions:

- Intel: Boost VMware vSphere Efficiency with Intel Optane Persistent Memory
- Dell: Dell EMC PowerEdge R650 Installation and Service Manual
- HPE: Intel Optane Persistent Memory 100 Series for HPE Guide
- Lenovo: Intel Optane Persistent Memory 200 Series Product Guide
- Cisco: Configuring and Managing Intel Optane Data Center Persistent Memory Modules

Required BIOS Settings

Intel® Virtualization Technology (Intel® VT). Intel VT-x extensions add migration, priority, and memory-handling capabilities to a wide range of Intel® processors and provide virtualization features that allow for more efficient execution of VMs. Intel VT-x also allows for the creation of 64-bit guest systems under VMware.

Intel® Virtualization Technology for Directed I/O (Intel® VT-d). This technology makes it possible for guest systems to directly access a PCIe device, with help from the provided Input/Output Memory Management Unit (IOMMU). This allows a local area network (LAN) card to be dedicated to a guest system, which makes it possible to attain increased network performance beyond that of an emulated LAN card. Once such a direct access system has been implemented, live migration of the guest system is no longer possible. VMware vSphere can be configured for use with an activated Intel VT-d system using VMware VMDirectPath for direct access to PCIe cards.

BIOS Settings for Best Performance

Modern data centers face growing concerns about power consumption, both from the perspective of total cost of ownership and total cost to the environment. Intel processors support BIOS settings that help conserve energy. However, in certain use cases, higher performance may be deemed more important than saving energy. The following BIOS options provide optimal performance on a variety of virtual applications on vSphere installations. However, they may not achieve additional objectives, such as optimal power consumption.

- **P-State.** If enabled, the CPU (all cores on a specific NUMA) will go to "sleep" mode if there is no activity. This mode is similar to C-State, but it applies to the whole NUMA node. In most cases, it saves power when the CPU is idle. However, for performance-oriented systems, when power consumption is not an issue, it is recommended that P-State be disabled.
- **C-State.** To save energy, it is possible to lower the CPU power when the CPU is idle. Each CPU has several power modes called C-states. If the BIOS is set to a performance profile, these operations are not suitable when aiming for optimal performance. Therefore, they should be disabled.
- **Turbo Mode.** Intel® Turbo Boost Technology automatically runs the processor core faster than the noted frequency. The processor must be working in the power, temperature, and specification limits of the TDP. Both single- and multi-threaded application performance is increased.
- **Hyper-threading.** Intel® Hyper-Threading Technology allows a CPU to work on multiple streams of data simultaneously to improve performance and efficiency. In some cases, turning hyper-threading off can result in higher performance with single-threaded tasks. Typically, Intel Hyper-Threading Technology should be enabled. In cases where the CPU is close to 100 percent utilization, hyper-threading might not help and can even harm performance. Therefore, in such cases, hyper-threading should be disabled.
- **VMware ESXi.** The hypervisor operating system can influence power profiles at an operating system level only if the BIOS power profile is set to "OS control mode" or equivalent. Exact options available will depend on the server platform used and the option exposed in the BIOS. Using alternative power profiles in the BIOS, such as "Performance," prevents ESXi from influencing power profiles within the operating system.

Other BIOS Considerations

- OEM BIOS software sometimes has profiles that initialize several of the BIOS settings to match application requirements. To improve performance, choose the profile that favors performance in virtualized environments.
- Platforms often present a trade-off between power savings and performance. Turning off power-saving schemes in the BIOS may result in improved performance of certain workloads at the expense of additional power consumed.
- Providing additional cooling to the CPU helps drive higher workloads more effectively. Set BIOS options that enhance cooling, such as changing the fan settings from Acoustic (low) mode to Performance (high) mode. A side-effect may be increased power consumed by the server.

Optimizing Latency

vSAN is a latency-sensitive application. You can reduce latency on a vSAN cluster with the following settings:

- **Use the highest frequency DIMM speeds** that the processor supports:
  - 2666 MHz for most 1st Gen Intel Xeon Scalable processors
  - 2933 MHz for most 2nd Gen Intel Xeon Scalable processors
  - 3200 MHz for most 3rd Gen Intel Xeon Scalable processors
- **Choose a BIOS performance profile** that optimizes performance.
- **Use a 100 GbE network interface card (NIC)** or at least a 25 GbE NIC, such as the Intel® Ethernet Network Adapter E810, to help avoid any networking I/O bottlenecks. Ensure that the NIC supports RDMA and RDMA over Converged Ethernet (RoCE) V2. For more information, visit Intel's Ethernet Adapters webpage.
  - Enable priority flow control for traffic class 3 on the physical switch.
  - Consider enabling RDMA to offload some of the vSAN CPU consumption to the NIC. In recent tests, enabling RDMA resulted in an up to 10 percent increase in IOPS and up to 8 percent lower CPU utilization when using the latest RDMA drivers with a 3rd Gen Intel Xeon Scalable processor.[7] For more information, read Network Requirements for RDMA over Converged Ethernet and Configure Remote Direct Memory Access Network Adapters.
- We recommend **using a Maximum Transmission Unit (MTU) of 9000** to increase overall vSAN throughput. Follow the best practices regarding latency in the vSAN Planning and Deployment guide.

- Always **provide adequate power** to the server by plugging in all the redundant power supply units on the server. This also helps keep the server running if one of the power supply units fails.
- **Evenly divide cache-tier storage devices and capacity-tier storage devices between CPUs.** This is not always possible with configure-to-order OEM servers, but do this where possible. We have seen best results when storage is divided evenly between CPU Socket 1 and Socket 2. This may require specific placement of storage controllers in PCIe slots and an understanding of the platform's interconnects between PCIe slots and CPU sockets. For more information, check the technical product specification guide: Intel® Server Board M50CYP2SB1US2600WF Product Family Technical Product Specification.

Recommended Hardware

CPU

- Intel Xeon Scalable processor
- Performance is driven by frequency
- VM density is increased with core count
- For more information: Intel Xeon Scalable Processors

Memory

Memory can be configured with a single DRAM tier or with a DRAM tier plus an Intel Optane PMem tier.

- **For best performance,** we recommend using all DRAM channels in a socket (eight slots). For instance, for 128 GB DRAM capacity per socket, use 8x 16 GB DDR4 DIMMs.
- **For tiered memory,** we recommend the same eight DRAM DIMMs plus four or eight Intel Optane PMem DIMMs per socket.

We also recommend obtaining the vSphere active memory data for the workloads that will be running in the host when determining the DRAM-to-PMem ratio. When no data is available, a best practice has been to use a 1:4 ratio.

For more details, please refer to the Boost VMware vSphere Efficiency with Intel Optane Persistent Memory.

Network

- **For small deployments,** we recommend a minimum of 10 GbE.
- **For most workloads,** we recommend 25 GbE or 100 GbE.

With a 100 GbE NIC, customers can expect a lower packet latency and more than double the read throughput on a vSAN cluster, when compared to previous-generation NICs.[8]

Storage

Cache tier: As discussed in the Solution Overview section, the performance characteristics of the cache drive make a significant difference to vSAN I/O performance. But drive endurance is another critical factor to be considered. Use Intel Optane SSDs for higher write throughput and endurance. A NAND-based 800 GB drive with an endurance rating of three DWPD will last five years only if the sustained write workload is lower than 2,400 GB per day or 28.4 write MB/s.[9] Put another way, even a light sustained I/O load such as 450 32-KB write IOPS activity would fill the drive three times per day, assuming the use of RAID 1 as data protection. In contrast, an Intel Optane SSD P5800X 400 GB drive has an endurance rating of 100 DWPD. It can sustain 40,000 GB written per day or 470 MB/s written per day for five years.[10] **That is a 16X increase in endurance.** If customers use a three-DWPD drive for the cache tier, they may increase their risk of costly downtime or may need to add more cache drives to the cluster to increase longevity, which can increase configuration costs.

For more details about sizing the cache tier, refer to the following VMware blogs: Extending All Flash vSAN Cache Tier Sizing Requirement for Different Endurance Level Flash Device; Do I Need a Bigger Write Buffer?

Capacity tier: Choose between SATA SSDs (lower cost) or NVMe SSDs (higher performance) like the Intel® SSD D7-P5510 Series.

Benchmarking Your Cluster

vSAN uses HCIBench as the benchmark to measure the IOPS and latency of different I/O workloads. The settings suggested in this guide have used HCIBench to provide optimal results on Intel Xeon Scalable processors.

Note that HCIBench is a synthetic workload. Hence, settings that result in improved results in the benchmarks may not necessarily translate to similar gains on real-world workloads.

For more details on performance using the settings in this guide, refer to the vSAN Planning and Deployment guide.

Experience the Difference

Intel® Optane™ technology can be used for delivering large memory pools or in applications requiring fast caching or fast storage. Intel Optane technology is available in a variety of products and solutions, including VMware vSAN and many others. For more information, go to Intel Optane Technology for Data Centers.

Installation and Configuration

Cache Tier Considerations

As mentioned earlier, it is recommended that all ESXi hosts in the cluster share similar or identical configurations across all cluster members, including similar or identical storage configurations. This consistent configuration balances VM storage components across all devices and hosts in the cluster.

In all-flash disk groups, vSAN uses the cache tier SSD as a write buffer. One-hundred percent of a cache device's capacity, up to a maximum of 600 GB, is used for the cache tier, while any remaining capacity is used for endurance.

When choosing a cache tier device, look for four key characteristics:

- **High random and sequential write throughput capability** to enable predictable, consistent I/O performance even when there are sudden variations in the aggregated cluster workload.

  The Intel Optane SSD P5800X provides nearly 8x higher write throughput per GB than NAND flash SSDs.[11]
- **Ability to sustain a low read latency** even when processing a high concurrent write activity.
  - *The Intel Optane SSD P5800X can maintain a low read latency with almost any concurrent write load.*[12]
  - *The Intel Optane SSD P5800X provides significant advantage over NAND when the workload shifts to a 50/50 mixed read/write, which is typical for a cache tier.*[13]
- **High read and write I/O capabilities** to keep up with today's demand and still provide sufficient headroom for future growth.
- **High endurance,** so that the SSD can handle heavy write activity without wearing out too fast.
  - *The Intel Optane SSD P5800X provides 100 DWPD, compared to three DWPD for a NAND SSD.*[14]
  - *The Intel Optane SSD P5800X supports a 73,000 TBs written (TBW) rating, compared to the Intel® SSD D7-5600 drive's 8,760 TBW rating.*[15]

Intel Optane SSDs are ideal for the cache tier because they satisfy all four characteristics. For more information about caching with vSAN, read the VMware blog.

Capacity Tier Considerations

When writes are being performed, the data is initially staged in the cache tier. As the cache tier fills, the data is de-staged to the disks in the capacity tier. This means that:

- The cache tier must be capable of dealing with both the incoming writes from the VMs and the reads from the de-staging activity.
- The capacity tier must be capable of processing the write requests from the de-staging activity.

The VM read requests in an all-flash vSAN cluster are performed directly on the capacity tier—most of the time—so the speed/throughput of the capacity tier is important to properly serve the cache de-staging activities and the VM read requests. This capacity tier throughput and speed depends on the characteristics of the drives used and on the number of drives.

While a configuration with just one capacity drive is possible, two or more are a better option for performance and cost. Similarly, SATA drives can satisfy several use cases, but as more workloads are aggregated in the vSphere/vSAN cluster and new data-intensive workloads are added, the throughput of the individual drives increases in relevance, affecting both individual VMs and overall cluster throughput. To maximize cluster performance, we recommend using PCIe 4.0 x4, NVMe drives, such as the Intel SSD D7-P5510 Series.

vSphere Sizing Recommendations

Appropriately configuring your VMware vSphere cluster, including the vSAN cache and capacity tiers, is crucial to achieving optimal performance (see Table 2).

Table 2. High-performance Host Configuration Recommendation

Component	Recommendation
CPU	2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz)
Memory	256 GB (16x 16 GB)
Intel® Optane™ PMem	1024 GB (8x 128 GB)
Storage Capacity	8x Intel® SSD D7-P5510 3.84 TB
Cache Drives	2x Intel® Optane™ SSD P5800X 400 GB
Network	1x Intel® Ethernet Adapter E810C 100 GbE NIC *For some deployments, 25 GbE may be sufficient*

Setup Summary

At a high level, setting up a vSAN environment consists of three steps, detailed below.

Install ESXi

1. Download VMware vSphere 7.0 and create a bootable media using VMware vSphere 7.0U2 or higher.
2. Install ESXi 7.0 to the boot device:
   - Enable SSH if planning to benchmark vSAN cluster.
   - Customize networking settings.

Install vCenter

1. Download VMware vCenter (VCSA) 7.0.
2. Mount the .iso and execute the installer:
   - Provide ESXi host information to deploy VCSA.
   - Provide networking/subnet details if using static IP addressing.
   - Select vCenter sizing dependent on data center needs.
3. Navigate to the URL associated with your vCenter Server Appliance for further configuration, if required.

Configure vSAN

1. Log in to the vCenter instance.
2. Create a data center, if required.
3. Create a cluster, if required.
4. Create a distributed virtual switch:
   - Assign hosts to the distributed virtual switch and assign uplinks.
   - Create port groups for vSAN, vMotion, and VM traffic.
   - Create vmkernel adapters for vSAN and vMotion port groups and apply this to all hosts that will be a part of the vSAN cluster.
5. Enable the vSAN service on the cluster:
   - Choose data-efficiency features such as duplication and compression.
   - Select cache tier devices and capacity tier devices.

Summary

The significant operational benefits continue to drive the consolidation of workloads on HCI/hybrid cloud platforms. And an increase in VM density continues to be the go-to option to control costs, energy use, and footprint. But as the number of VMs per cluster increases, so does the demand for I/O to the storage systems and network, requiring data center architects to rethink the components that they use in these HCI platforms.

Using Intel Optane SSDs as cache tier devices, along with higher-performance capacity drives and a wider network bandwidth, enables a significant increase in the overall cluster I/O throughput. This makes the configuration more scalable and enables higher VM density, creating good opportunities for cost and footprint reduction.[16] In addition, the higher write throughput of Intel Optane SSDs provides a higher level of predictability; in other words, it enables the vSAN system to sustain the I/O latency even when some of the VMs that share the cluster resources cause sudden variations in the I/O load. This facilitates the management of workloads with stricter service-level objectives.

As the size of the hosts used in the vSphere and vSAN clusters continues to increase, the cost of memory takes a larger portion of the total host cost. The use of Intel Optane PMem in Memory Mode creates significant opportunities to reduce infrastructure costs.

By using the recommendations contained in this best practices guide, you can create a vSAN storage environment that can scale as storage needs increase, while meeting today's performance and operational efficiency requirements.

[1] Testing by Intel as of May 10, 2021.

Intel® Optane™ SSD Configuration: 4 nodes, 2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/16 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD P5800X (cache) 400 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, HCIbench 2.5.3.

Intel® SSD D7-5600 Configuration: 4 nodes, 2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/16 GB/3200 MT/s), Intel Hyper-Threading Technology = ON, Intel Turbo Boost Technology = ON, 2x Intel® SSD D7-P5600 (cache) 1.6 TB and 8x Intel SSD D7-P5510 3.84 TB (capacity), 1x Intel Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 0x8d055260), VMware vSphere 7.0U2, HCIbench 2.5.3.

[2] Testing by Intel as of May 10, 2021.

Intel® Optane™ SSD P5800X Configuration: 4 nodes, 2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/16 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD P5800X (cache) 400 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, HCIbench 2.5.3.

Intel Optane SSD DC P4800X Configuration: 4 nodes, 2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/16 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD DC P4800X (cache) 375 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, HCIbench 2.5.3.

Intel® SSD D7-5600: 4 nodes, 2x Intel® Xeon® Gold 6348 processor (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/16 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® SSD D7-P5600 (cache) 1.6 TB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, HCIbench 2.5.3.

[3] Testing by Intel as of May 10, 2021. Based on 280 VMs - 4 vCPUs per VM, 8 GB MEM, 125 GB usable storage capacity, up to 1,500 IOPS per VM running a 70/30 32 KB I/O load. Overheads and optimal utilization levels were considered in calculations. Results may vary.

New Configuration: 4 nodes, 2x Intel® Xeon® Gold 6348 processor, (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/32 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD P5800X (cache) 400 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 0x8d055260), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3, 8x VMs per host, 2x 150 GB vDisks per VM, 100% WSS.

Baseline Configuration: 4 nodes, 2x Intel® Xeon® Gold 6248 processor (20 cores, 2.5 GHz), total memory = 384 GB (12 slots/32 GB/2933 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD DC P4800X (cache) 375 GB and 8x Intel SSD D7-P5510 3.84 TB (capacity), 1x Intel Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3, 8x VMs per host, 2x 150 GB vDisks per VM, 100% WSS.

[4] Intel® Optane™ SSD P5800X Series: intel.com/content/www/us/en/products/docs/memory-storage/solid-state-drives/data-center-ssds/optane-ssd-p5800x-p5801x-brief.html

[5] CPU cost was estimated. Pricing varies over time. dell.com/en-us/work/shop/cty/pdp/spd/poweredge-r750/pe_r750_14794_vi_vp?configurationid=b605e5ac-c8b9-4578-b0e2-7d9b15772b04.

[6] Claim [3] at Intel® Optane™ Persistent Memory 200 Series - 1 - ID:615781 | Performance Index

[7] Testing by Intel as of May 10, 2021. Workloads shown: Random Read 4KB, Random Write 64KB, 70/30 R/W 32KB. Drivers: ICEN: 1.5.5.0-1OEM.700.1.0.15843807; Irdman: 1.3.3.0-1OEM.700.1.0.15843807.

2nd Generation Intel® Xeon® Scalable Processor Configuration (no RDMA): 4 nodes, 2x Intel® Xeon® Gold 6248 processor (20 cores, 2.5 GHz), total memory = 384 GB (12 slots/32 GB/2933 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD DC P4800X (cache) 375 GB and 8x Intel SSD D7-P5510 3.84 TB (capacity), 1x Intel Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3.

3rd Generation Intel® Xeon® Scalable Processor Configuration (no RDMA): 4 nodes, 2x Intel® Xeon® Gold 6348 processor, (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/32 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD P5800X (cache) 400 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 0x8d055260), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3.

3rd Generation Intel® Xeon® Scalable Processor Configuration (with RDMA): 4 nodes, 2x Intel® Xeon® Gold 6348 processor, (28 cores, 2.6 GHz), total memory = 256 GB (16 slots/32 GB/3200 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD P5800X (cache) 400 GB and 8x Intel® SSD D7-P5510 3.84 TB (capacity), 1x Intel® Ethernet Adapter E810C 100 GbE, BIOS = 2.1 (ucode = 0x8d055260), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3.

[8] Testing by Intel as of May 10, 2021

New Configuration: 4 nodes, 2x Intel® Xeon® Gold 6248 processor (20 cores, 2.5 GHz), total memory = 384 GB (12 slots/32 GB/2933 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD DC P4800X (cache) 375 GB and 8x Intel SSD D7-P5510 3.84 TB (capacity), 1x Intel Ethernet Adapter E810C at 100 GbE data rate, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3.

Baseline Configuration: 4 nodes, 2x Intel® Xeon® Gold 6248 processor (20 cores, 2.5 GHz), total memory = 384 GB (12 slots/32 GB/2933 MT/s), Intel® Hyper-Threading Technology = ON, Intel® Turbo Boost Technology = ON, 2x Intel® Optane™ SSD DC P4800X (cache) 375 GB and 8x Intel SSD D7-P5510 3.84 TB (capacity), 1x Intel Ethernet Adapter E810C limited to 25 GbE data rate, BIOS = 2.1 (ucode = 05003003), VMware vSphere 7.0U2, vSAN 7.0U2, HCIbench 2.5.3.

[9] Calculation: 28.44 MB/s * 3600 secs * 24 hours = ~2400 GB

[10] Calculation: 470 MB/s * 3600 secs * 24 hours = ~40,000 GB

[11] See endnote 1.

[12] See various claims at https://edc.intel.com/content/www/us/en/products/performance/benchmarks/intel-optane-ssd-p5800x-series.

[13] See endnote 1.

[14] Intel® Optane™ SSD P5800X: https://ark.intel.com/content/www/us/en/ark/products/201859/intel-optane-ssd-dc-p5800x-series-1-6tb-2-5in-pcie-x4-3d-xpoint.html and Intel® NAND SSD: https://www.intel.com/content/www/us/en/products/sku/202707/intel-ssd-d7p5600-series-1-6tb-2-5in-pcie-4-0-x4-3d3-tlc/specifications.html

[15] TBW rating calculation: Size of drive in TB x DWPD x 365 days x 5 years.

[16] See endnotes 1-3.

Performance varies by use, configuration and other factors. Learn more at 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. 1221/ACHO/KC/PDF 337370-002US