Solution Brief

Data Center | Software-Defined Storage

Transform Data Center Efficiency with Technology Innovation

Process larger datasets in less time, drive data center efficiency and modernization, and optimize your infrastructure investment with Intel® Optane™ technology and Microsoft Azure Stack HCI

Microsoft Azure

Solution Benefits

Microsoft Azure Stack HCI, powered by 2nd Gen Intel® Xeon® Scalable processors, uses Intel® Optane™ technology to help organizations reimagine memory and storage, and accelerate data-centric transformation. Organizations can expect to:

- **Access your hottest data faster.** The solution's robust capabilities allow users to handle large datasets and the most demanding workloads.
- **Do more with less.** The solution's high-performance design helps increase IOPS while requiring fewer nodes for greater consolidation.
- **Help ensure reliability.** Microsoft Azure Stack HCI solutions use Microsoft-validated hardware for optimal performance and reliability.

Executive Summary

Businesses of all sizes and across industry verticals must optimize and modernize their data centers to remain competitive. Through hyperconverged infrastructure (HCI), Microsoft Azure Stack HCI solutions can enable organizations to simplify deployments, scale operations, increase reliability and manageability, and maximize resource utilization—all of which helps transform data center efficiency and help reduce total data center costs.

By combining highly virtualized compute, memory and storage, and networking on industry-standard servers and components, Microsoft Azure Stack HCI makes it possible to run virtualized applications on premises as well as connect to Microsoft Azure for cloud services.

Microsoft Azure Stack HCI is optimized for 2nd Generation Intel® Xeon® Scalable processors, a workload-optimized platform designed to deliver agility with enhanced performance and advanced capabilities. Adding Intel® Optane™ technology to Microsoft Azure Stack HCI can help organizations increase data throughput, reduce latency, affordably increase memory capacity, and quickly extract value from large datasets for timely, actionable insights, all while potentially consolidating workloads on a smaller data center footprint.

Intel has created multiple reference architectures to accelerate infrastructure decisions and solve storage efficiency and memory capacity issues. Organizations can look to these modern solutions to help meet today's storage and memory requirements across a wide variety of use cases.

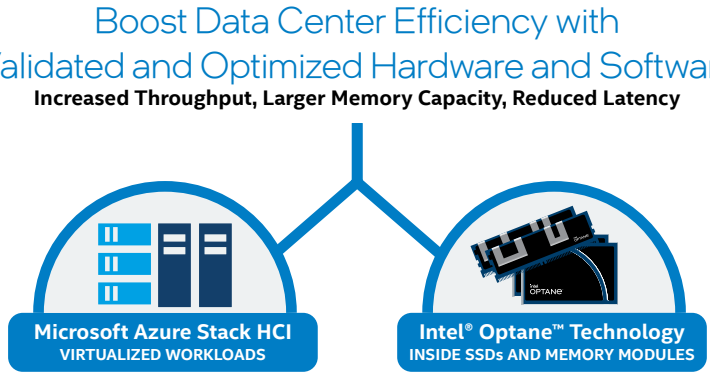

Figure 1. Intel Optane technology helps businesses using Microsoft Azure Stack HCI optimize their data centers.

Business Challenge: Choose the Right Infrastructure for Data-Hungry Workloads

Constantly expanding datasets can be overwhelming, but they also offer tremendous opportunities to organizations that know how to squeeze the most value from their data. Companies of all types and across all industries now understand the importance of taking advantage of the growing volume of data to gain actionable insights that help inform their business decisions and keep them competitive.

Rapidly mining data requires an infrastructure tuned to a specific workload's memory and storage requirements. Today's organizations need a highly flexible, scalable solution that provides affordable handling of high-performance, high-density data. HCI—tightly integrated and virtualized compute, memory and storage, and networking—combined with Intel® Optane™ technology can help organizations meet these challenges.

Performance-Optimizing Use Cases

Today's organizations are looking to create efficiencies, improve performance, and protect their infrastructure investments as much as possible. Here are some use cases that illustrate how Microsoft Azure Stack HCI can help businesses find operational efficiencies and boost performance.

- **Virtual desktop infrastructure (VDI).** Azure Stack HCI suits users that want to implement remote desktop virtualization on a large scale. VDI provides user desktops through a virtual desktop broker. The virtual desktops connect back to VMs and central storage on the Azure Stack HCI cluster. VDI delivers client desktops on a range of devices so users don't have to store data locally or upload data from local devices, helping to enhance security.
- **Trusted enterprise virtualization.** Azure Stack HCI works in general business scenarios in which customers want to serve any applications hosted on VMs with high security and high availability. Azure Stack HCI provides a security-enabled infrastructure for workloads through virtualization-based security (VBS), using the Hyper-V hypervisor to create and isolate the Virtual Secure Mode (VSM) from the normal operating system. Enabling VBS allows security-sensitive operations to occur in the isolated memory, independent of the host operating system.
- **High-performance SQL Server.** Running Microsoft SQL Server in an Azure Stack HCI cluster allows users to run SQL Server and associated applications with the resiliency virtualization provides. Azure Stack HCI also offers the benefit of a single vendor for its hypervisor, host operating system, and database server when used to host SQL Server. Adding Azure Backup provides comprehensive database backup management. While Azure Stack HCI deployments built to support database servers are typically performance-optimized, users can boost performance with high-bandwidth PCIe/NVMe-based Intel® SSDs. Users can further increase performance by using Intel® Optane™ SSDs for caching and SATA-based SSDs or HDDs for capacity.
- **Branch office and edge.** Azure Stack HCI helps businesses like retail stores, branch offices, field sites, and other edge sites that want affordable, highly available, and resilient storage for business-critical applications and new edge workloads built on containers. Azure Stack HCI solutions designed for this use case offer cost-effective fault tolerance and resilience. Intel SSDs can be used in a single tier, or as a cache tier to support HDDs in the capacity tier.
- **Scale-out storage.** Businesses that require file serving with high scalability, performance, and availability can look to Azure Stack HCI to provide storage performance on validated hardware that can be optimized for density, speed, or performance-to-cost ratio. This requires an affordable file server and a small hardware commitment. This use case draws on the capabilities of Microsoft Storage Spaces Direct, which creates a pool of highly available and highly scalable storage from locally attached drives at a cost that can be lower than traditional SAN or NAS arrays. Intel SSDs support the cost-savings and performance goals of scale-out storage servers. Using high-bandwidth PCIe/NVMe-based Intel SSDs for all storage devices will help optimize performance. Using Intel Optane SSDs as cache drives and SATA-based SSDs or HDDs as capacity drives can deliver great performance at a low cost.

intel OPTANE™

Organizations that use Microsoft Azure Stack HCI can optimize their data centers with Intel® Optane™ technology.

Intel® Optane™ Persistent Memory (PMem)

Intel® Optane™ persistent memory gives enterprises the ability to extract more from larger datasets by combining more capacity and native persistence in a DIMM form factor. Data can be accessed, processed, and analyzed in near real time to deliver deep insights, improve operations, and create new revenue streams.

Intel® Optane™ SSDs

Intel Optane SSDs help remove data bottlenecks to accelerate transactions and time to insights, so users get what they need, when they need it. With high quality of service and at least 6x faster performance than NAND SSDs at low queue depths, Intel Optane SSDs deliver fast, predictable performance even in the most demanding environments.[1]

Solution Value: Get Speed *and* Scale

Microsoft Azure Stack HCI is a hyperconverged Microsoft Windows Server 2019 cluster that uses validated hardware to run virtualized workloads on-premises. Microsoft Azure Stack HCI solutions combine highly virtualized compute, memory and storage, and networking on industry-standard x86 servers with local-attached drives to create highly available, highly scalable SDS. The components of the solution include Microsoft Storage Spaces Direct, which has built-in support for PCIe/NVMe-based Intel SSDs, Intel Optane SSDs, and Intel® Optane™ persistent memory (PMem), as well as Microsoft Hyper-V, which serves as the hypervisor. Azure Stack HCI dramatically simplifies deployment, while caching, storage tiers, and erasure coding,

together with the latest hardware, deliver exceptional efficiency and performance.

The solution is powered by 2nd Generation Intel Xeon Scalable processors, which provide compute capabilities to the VMs, as well as enabling I/O and storage efficiency technologies such as deduplication, compression, and erasure coding. Adding Intel Optane technology to Microsoft Azure Stack HCI can result in high performance and low latency. With Intel Optane technology, organizations don't need to choose between speed and scale. Various aspects of Intel Optane technology work together to enable organizations to process, store, and move larger and more complex datasets. The technology bridges critical gaps in the storage and memory hierarchy to deliver a combination of persistent memory, large system memory pools, rapid caching, and fast storage, depending on which Intel Optane technology is in use. Overall, combining 2nd Gen Intel Xeon Scalable processors with Intel Optane PMem and Intel Optane SSDs can help organizations efficiently access more data, consolidate servers, and significantly reduce costs by moving data closer to the CPU while delivering greater agility and value to challenging data center environments.

Increase Caching Speed

Low-latency, high-performance, high-endurance, and reliable Intel Optane SSDs in the cache tier help to break through storage bottlenecks and minimize I/O wait times. Reducing I/O wait helps to recover CPU cycles, which can help to increase workload density, and that ultimately can enable a reduction in node count. Fewer nodes not only mean a smaller hardware investment, but can also lead to lower software licensing costs. On a workload simulating a multi-VM environment on Microsoft Storage Spaces Direct, adding Intel Optane SSDs to a server cluster resulted in 26.6 percent faster response times, and improved IOPS by up to 52.9 percent using one less server, compared to a solution using only SATA SSDs (see Figure 2).[2]

Improve Memory Capacity

For memory-intensive workloads, Intel Optane PMem can provide benefits beyond using DRAM alone. The technology introduces a new flexible tier within the memory/storage hierarchy, which is immediately applicable to workloads across cloud, in-memory computing, and storage. This new technology improves system speed and efficiency, increases data availability, and delivers rapid data insights.

Intel Optane PMem in Memory Mode allows users to expand volatile system memory, thereby increasing VM density and the amount of memory available to applications. By affordably expanding system memory with Intel Optane PMem in Memory Mode, the number of VMs per node can increase by up to 36 percent. This increases VM density in the Microsoft Azure Stack HCI cluster and can help to decrease costs (see Figure 3).[3]

Increase Bandwidth and Memory

Intel Optane PMem in App Direct Mode can serve as non-volatile (persistent) cache. Persistence not only protects data in the event of a power loss, but can shorten application restart times and increases application performance, as memory no longer has to be continuously flushed to storage devices.

In App Direct-Dual Mode (Memory Mode and Storage Over App Direct Mode), this technology can serve both as non-volatile (persistent) cache, as well as fast, volatile system memory, with a percentage of the memory assigned to each mode. This configuration can increase VM density by up to 41 percent and boost throughput by up to 76 percent (see Figure 4).[4]

Achieve Low-Latency Ethernet

Traditionally, storage workloads have run over a dedicated fabric. However, HCI environments distribute data across nodes in the cluster using a standard Ethernet network, making network performance an important factor in the HCI solution. Microsoft Azure Stack HCI has built-in support for high-performance, low-latency Intel® Ethernet network adapters that support RDMA. RDMA improves the throughput for traffic between nodes, enabling a low-latency, high-throughput direct memory-to-memory data communication between applications over a network.

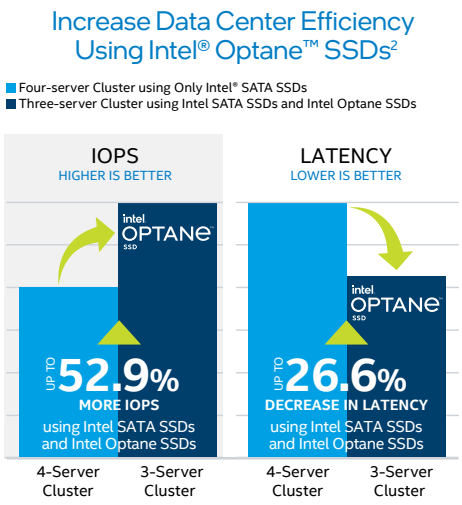

Figure 2. Get more I/O performance on fewer servers with Intel Optane SSDs in the cache tier.

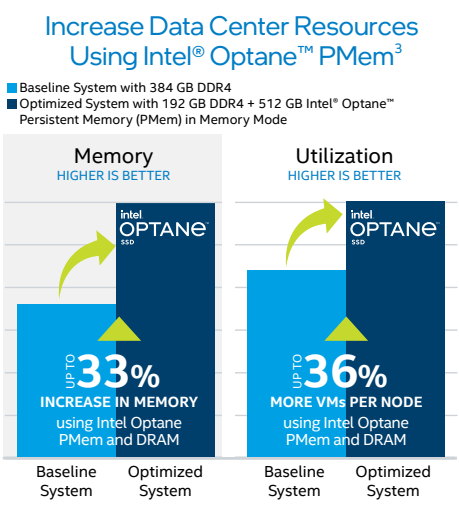

Figure 3. Adding Intel Optane PMem resource utilization when used in Memory Mode.

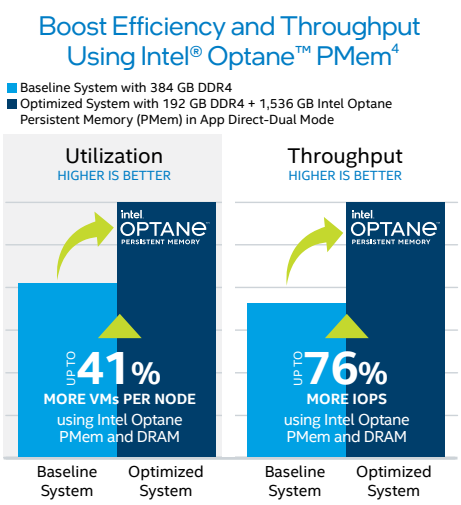

Figure 4. Adding Intel Optane PMem can boost efficiency when used in App Direct-Dual Mode.

Organizations can use these Intel® technologies to solve their memory and storage issues and accelerate applications, resulting in excellent overall data center efficiency.

Solution Architecture: Workload-Driven HCI Designs

Intel simplifies infrastructure decision making by defining reference designs, optimized for specific use cases for organizations that need faster caching, more memory—or both.

Figure 5 illustrates the reference architecture for workloads that can benefit from using Intel Optane technology. Each of the options uses Intel Optane SSDs; Options 2 and 3 add in Intel Optane PMem. Here are the details:

- **Option 1.** Use Intel Optane SSDs as cache, plus SATA-based Intel SSDs for the capacity tier, to speed caching and increase VM density—leading to server consolidation.
- **Option 2.** Workloads that need more memory can benefit from Intel Optane PMem in Memory Mode, in addition to the Intel Optane SSDs in the cache tier.
- **Option 3.** Workloads that not only need additional memory, but also need extremely low latency can combine Intel Optane PMem in App Direct-Dual Mode, where the cache layer uses a two-tier architecture that allows for much faster cache allocation and frees up drive bays for more capacity. The App Direct Mode persistent memory replaces Intel Optane SSDs in the cache tier. This configuration is ideal for high-bandwidth SQL Server use cases.

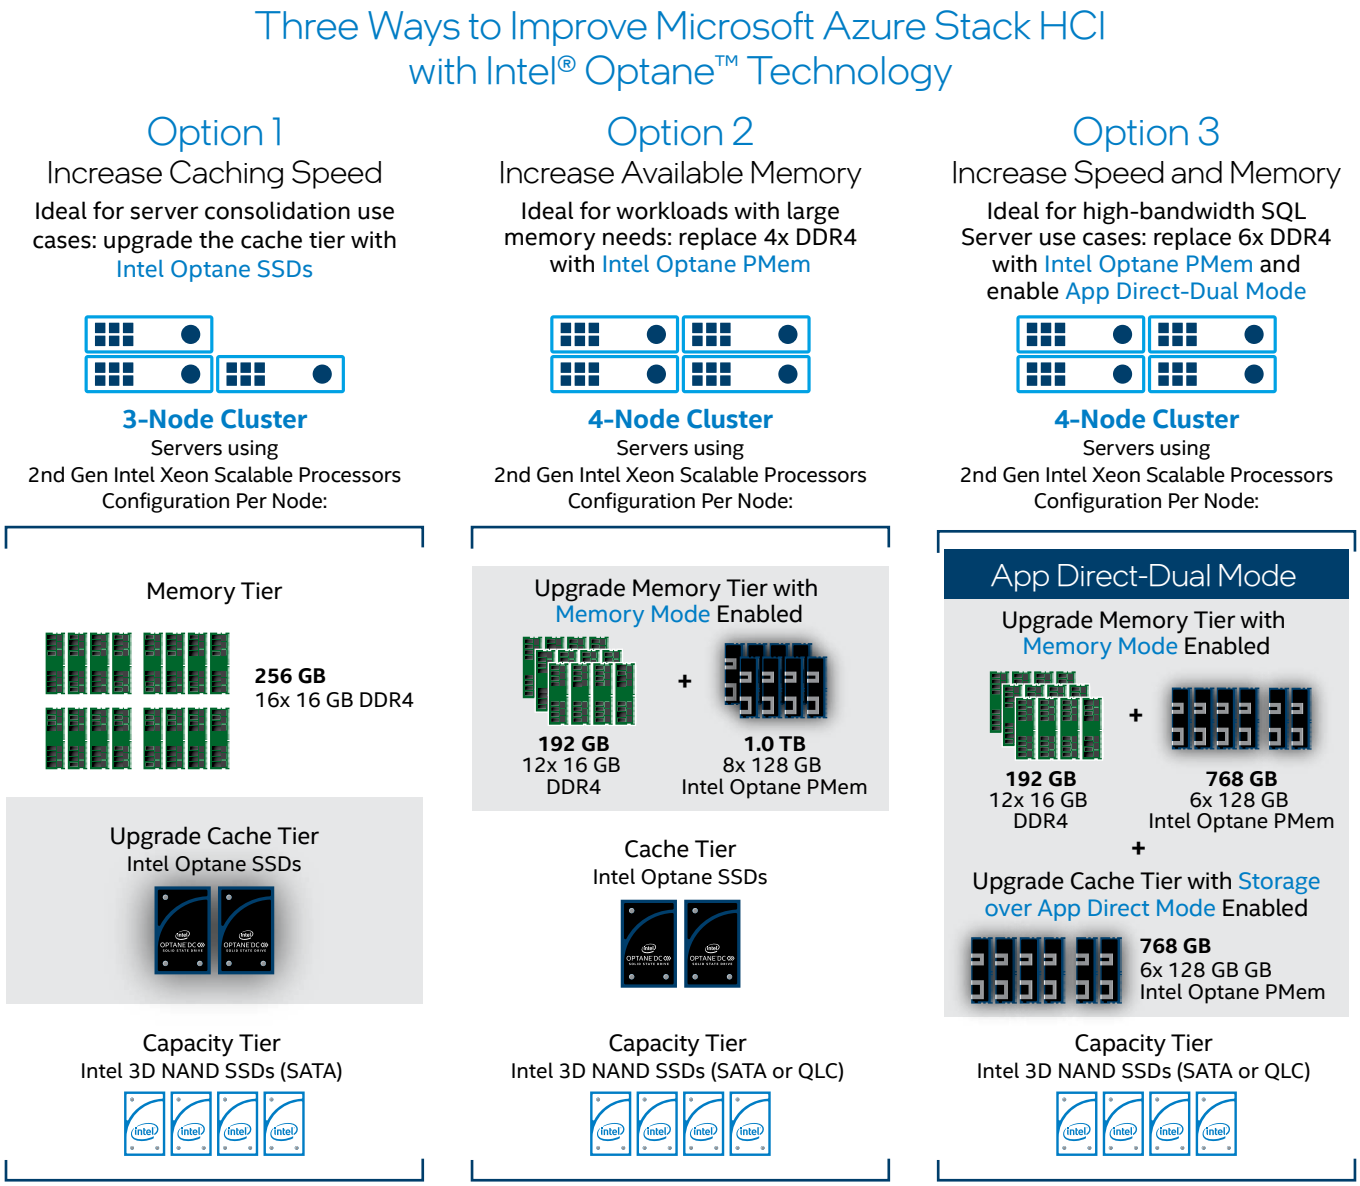

Figure 5. Adding Intel Optane technology can speed caching, increase VM density, and improve throughput for a variety of workloads.

Build a Balanced Architecture Design

It is a best practice to increase all subsystems by the same factor, thereby avoiding creating bottlenecks by lack of one or more subsystems. With the increase in options to design HCI solutions, it is imperative that all resources support one another in a balanced approach. The key is not to under-provision nor over-provision any subsystem. Overall, every subsystem needs to grow when one grows, otherwise bottlenecks are highly likely.

Conclusion

Optimizing compute, memory, and storage is important for organizations that want to perform fast data analysis to inform business decisions that can improve time-to-market and lead to competitive advantage. Enterprises need an infrastructure that's tuned to a specific business scenario or that suits their requirements. Companies can use Microsoft Azure Stack HCI, with a foundation of 2nd Gen Intel Xeon Scalable processors and Intel Optane technology, to transform data center efficiency and save on hardware and licensing costs. The solution allows users to either get more functionality from fewer nodes by enabling more VMs per server, add more servers to do exponentially more work, or extend system memory and memory caching to improve throughput and increase VM density at the same time.

Replacing standard SSDs with Intel Optane SSDs at the cache tier can help businesses get insights faster. And by including Intel Optane PMem, organizations can access their hottest data faster, optimize their infrastructure investment, and modernize their data centers with confidence. Together, these technologies can give organizations of every size and type the ability to modernize their infrastructure to get the most value from their data so they can grow and thrive.

To get started with Microsoft Azure Stack HCI, visit the Azure Stack HCI product catalog, which offers solutions from 20 hardware suppliers. To simplify deployment even further, choose an Intel® Select Solution for Azure Stack HCI. Intel Select Solutions are verified configurations that can speed selection and deployment of data center and communications network infrastructure. The solutions are developed from deep Intel experience with industry solution providers, as well as extensive collaboration with the world's leading data center and service providers.

Learn More

You may also find the following resources useful:

- Intel® Optane™ technology
- Intel® Optane™ SSDs
- Intel® Optane™ PMem
- 2nd Generation Intel® Xeon® Scalable Processors
- Intel® Ethernet Technology
- Microsoft Azure Stack HCI
- Intel® Select Solutions for Azure Stack HCI

Find the solution that is right for your organization. Contact your Intel representative or visit **Intel Optane Technology for Data Centers**.

Solution Provided By:

Microsoft intel

[1] Intel-tested as of November 15, 2018. 4K 70/30 read/write performance at low QD. Measured using FIO 3.1. Common configuration: Intel 2U Server System, OS: CentOS 7.5, kernel 4.17.6-1.el7.x86_64, 2 x Intel(R) Xeon(R) Gold 6154 processor at 3.0 GHz (18 cores), 256 GB DDR4 RAM at 2,666 MHz. Configuration: 375 GB Intel Optane SSD DC P4800X compared to 1.6 TB Intel SSD DC P4600. Intel microcode: 0x2000043; system BIOS: 00.01.0013; Intel ME firmware: 04.00.04.294; BMC firmware: 1.43.91f76955; FRUSDR: 1.43. The benchmark results may need to be revised as additional testing is conducted.

[2] Testing by Principled Technologies as of August 7, 2019. For more information, visit principledtechnologies.com/Hpe/Intel-Optane-HPE-ProLiant-Storage-Spaces-Direct-0919.pdf and principledtechnologies.com/Hpe/Intel-Optane-HPE-ProLiant-Storage-Spaces-Direct-science-0919.pdf.

Common configuration: 2x Intel® Xeon® Gold 6154 processor @ 3.0 GHz (18 cores); 12 x 32 GB DDR4-2666 (total memory = 384 GB); OS drive = 1x Intel® SSD DC S3700 400 GB; Intel® Hyper-Threading Technology = ON; Intel® Turbo Boost Technology = ON; BIOS = U30 v1.46 (10/02/2018); BIOS setting = Performance; OS = Windows Server 2019 Build 1809 (patched 8/2/19); Power management policy = Static High Performance Mode; NIC = 2x Intel® Ethernet Adapter XXV710 (25 GbE). **All-SATA configuration:** four-node cluster; 4x Intel® SSD D3-S4510 3.84 TB. *Results: IOPS = 387,092; Latency = 6.0 ms.* **SATA plus Intel® Optane™ SSD configuration:** three-node cluster; 4x Intel® SSD D3-S4510 3.84 TB and 2x Intel Optane SSD DC P4800X 375 GB. *Results: IOPS = 592,173; Latency = 4.4 ms.* Workload: VMFleet/DISKSPD 2.0.21a.

[3] Performance results are based on testing by Intel as of January 15, 2019 and may not reflect all publicly available security updates. See the configuration disclosure for details. **Common configuration:** Intel® Xeon® Gold 6230 processor @ 2.10 GHz. **All-DRAM configuration:** 384 GB DDR4 DRAM memory. **DRAM + Intel® Optane™ persistent memory configuration:** 192 GB DDR4 DRAM memory + 512 GB Intel Optane persistent memory. Benchmark Setup: VMFleet Test: Each VM with 1 core, 8 GB; memory, 40 GB VHDX; Test Setup: threads = 2; buffer size = 4 KB; pattern: random, duration = 300 seconds; queue depth = 16, 30% write; OS: Windows Server 2019 Standard (desktop) with updated patch.

[4] Testing by Intel as of February 8, 2019. **All-DRAM cofiguration:** 2x Intel® Xeon® Gold 6230 processor with 384 GB DDR4 DRAM. Benchmark Setup: VMFleet Test: 18 VMs/node, each VM with 4 cores, 8 GB Memory, 40 GB VHDX, testfile: 10 GB. Test setup: Threads=4 , Buffer Size= 4 KB ,Pattern: Random , Duration = 300 Seconds, Queue Depth=16, 30% write, OS: Windows Server 2019 Standard (Desktop) with updated patch. **DRAM + Intel Optane PMem configuration:** 2x Intel Xeon Gold 6252 processor with 192 GB DDR4 DRAM and 1,536 GB Intel Optane PMem in App Direct-Dual Mode; PMem as cache: 2x 512 GB segments. Benchmark Setup: VMFleet Test: Each VM with 1 Core, 8 GB Memory, 40 GB VHDX. Test setup: Threads=2 , Buffer Size= 4KiB ,Pattern: Random , Duration = 300 Seconds, Queue Depth=16, 30% write OS: Window Server 2019 Standard (Desktop) with updated patch.

 1220/KMAH/KC/PDF 341107-002US