SOLUTION BRIEF
Data Center
Virtualization

Boost Memory Capacity and Lower TCO in Virtualized Data Centers

Increased VM density, enabled by Intel® Optane™ DC persistent memory and related innovative technologies from Intel and VMware, allows IT organizations to increase resource utilization

Solution Benefits

- Intel Optane DC persistent memory can offer **greater memory capacity per socket** than DRAM for virtualized data center infrastructures, leaving more headroom for virtualizing future workloads requiring larger memory capacity rather than having to run those demanding workloads on bare metal.
- Implementing solutions with Intel Optane DC persistent memory can **improve overall TCO**. This can free up the budget to refresh a portion of IT virtualized infrastructure and/or increase IT innovation.
- Accessing the storage subsystem can incur unacceptable latencies. Keeping more data in larger system memory can often **improve workload performance**.

Executive Summary

As data grows, the number of virtual machines (VMs) running databases and the number of VMs per server also likely increases. Enterprises need to increase the VMs and databases per node to meet demand, and build VMs to handle larger datasets. Without sufficient memory resources, VM and application performance can suffer. It's no surprise that application owners request high levels of memory allocation to protect their apps against possible workload surges and bottlenecks. While this is a conservative approach, it does not yield right-sized VM memory allocations, leaving organizations still searching for ways to boost memory capacity while lowering TCO.

Ideally, DRAM would be cheap and abundant, allowing database workloads to run in memory. But DRAM can be expensive, and the capacity of DRAM modules is not keeping up with application demand.

2nd Generation Intel® Xeon® Scalable processors with Intel® Optane™ DC persistent memory modules (DCPMMs) offer an affordable, higher-capacity alternative to conventional DRAM. Intel Optane DCPMMs place low-latency, high-performance Intel Optane media on a DDR4-format memory module, enabling higher total server memory capacities than DRAM alone. Greater memory capacities facilitate higher VM counts per server and server consolidation. Moreover, smart resource monitoring can lead to more right-sized allocation from these larger memory pools. These optimizations can lead to lower costs for hardware and licensing along with lower TCO, particularly for virtualized database solutions such as Microsoft SQL Server running under VMware vSphere. The low-latency, high-performance nature of Intel Optane DC persistent memory allows for these benefits while still maintaining nearly the same service and performance levels found in DRAM-only configurations.

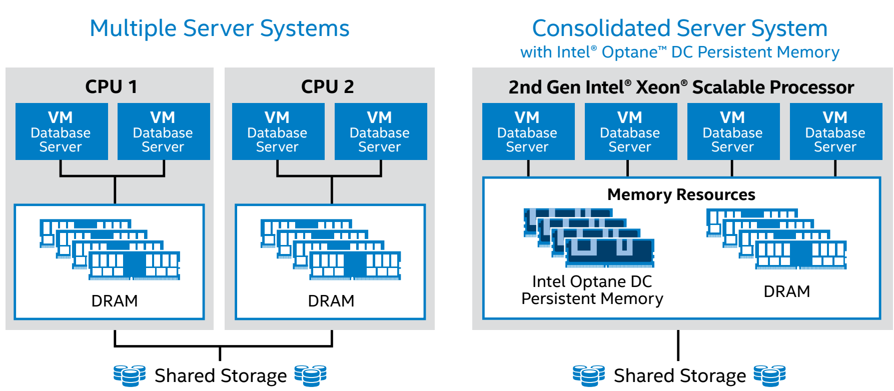

Figure 1. Replacing some DRAM with high-capacity Intel Optane DCPMMs is an affordable solution for significantly extending system memory.

Business Challenge: Rein in VM Application "Super-Sizing"

According to Flexera/RightScale, waste abounds in cloud costs. The RightScale 2017 State of the Cloud Report noted that almost one-third of cloud spend was waste, but subsequent thinking put that number even higher. The biggest waste center of all was compute instances of virtual machines (VMs), with Flexera noting that "typically 40 percent of instances are one to two sizes bigger than needed for their workloads, which wastes 50-75 percent of the spend on each oversized instance."[1]

Admittedly, one analyst's "waste" may be an application owner's "safety margin." IT managers often try to gauge consistent use levels, add a significant percentage to the top end of that range for burst loads, then pad further still for possible future growth. This conservative strategy does well for avoiding the performance impacts of disk swapping if and when workloads should exceed available memory, even if the ongoing requests and discussions between application owners and IT management over-defending allocation decisions can consume an inordinate amount of time.

However, the practice of leaving a super-sized safety margin may grow increasingly nebulous in an age of exponentially increasing data sizes and frustratingly vague (or complex) tools for measuring true utilization levels. How much safety is enough today—or 18 months from now? "Just in case" provisioning margins have made outsized DRAM allocations common and, as a result, yielded fleets of underutilized CPUs. A 2018 report by ParkMyCloud, encompassing "tens of thousands of VMs across hundreds of customers" found that the average peak CPU utilization was just 12.5%."[2] Underutilized hardware ultimately costs the industry billions of wasted dollars, and virtualized database platforms stand at the center of this loss.

Now is the perfect time realize the benefits of right-sizing VMs. Improving TCO through better resource utilization, and possibly consolidation, is one factor. Another may be the pressing need to update database applications, in part from the rapidly rising tide of IoT and other data sources and also from the recent ending of mainstream support for Microsoft SQL Server 2008. IT and business managers remain under constant pressure to tighten budgets, despite application users vying for more compute resources that might be needed in the future. Enterprises running VMware and similarly virtualized environments need affordable, non-disruptive ways to accommodate all these needs while maintaining the performance applications demand.

New Efficiencies and Possibilities Through Memory Right-Sizing

In cases of critical, exceptionally memory-hungry workloads, IT may run databases directly on bare metal to achieve maximum processing performance and steer clear of VM hypervisor and resource sharing overhead. More often, though, workloads aren't that memory hungry. Application owners may require a large amount of memory, or experience occasional load bursts to high memory levels, but they tend to approach higher capacities as a means to meet performance needs by avoiding disk swapping. This is an especially common scenario with database applications, such as Microsoft SQL Server, which can avoid memory over-subscription in virtualized environments and provide better performance when memory is abundant.

When applications and/or workloads are provisioned with ample amounts of memory in a virtualized setting, a scenario like the one shown in Figure 2 may develop. This snapshot shows a 30-day use chart for VM memory resources. Admins provisioned nearly 16.78 TB (the consumed line) of memory for the VMs, but in reality the applications demanded an average of only 2.39 TB. Even accommodating over 50 percent above this level for load spikes and a safety buffer, a recommended amount of DRAM for the VM's active data would be 3.73 TB, leaving almost 11.5 TB of DRAM provisioned to the VM completely idle.

Intel Optane DC persistent memory can replace overprovisioned DRAM. Because nearly all active data remains in DRAM that has been right-sized for the workloads, the workloads continue to perform at virtually unchanged levels. Meanwhile, Intel Optane DCPMMs can assist during peak application workloads while leaving application owners with the normal workload memory amounts they desired. With Intel Optane DCPMMs shipping in sizes that are multiple times bigger than today's available DRAM modules (enabling up to 6 TB of Intel Optane DC persistent memory per 2-socket system), admins may also process larger workloads than were possible with the previous capacity constraints of DRAM, even while increasing the number of VMs per server. Moreover, that bare-metal database workload scenario mentioned earlier can now be run more cost-effectively with virtualization.

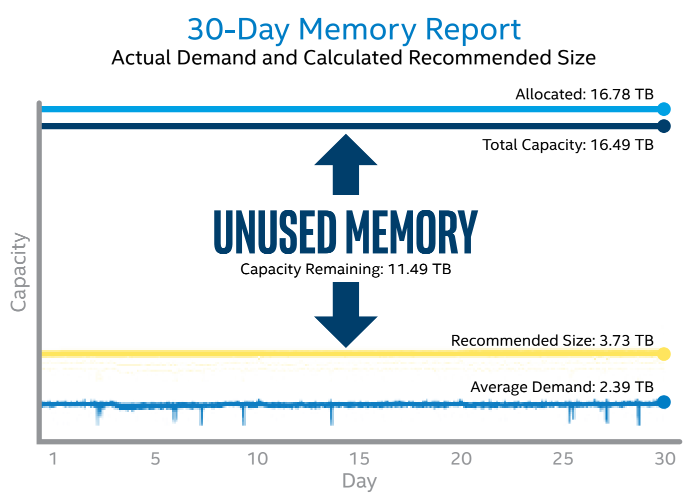

Figure 2. This 30-day profile of a cluster's cumulative VM memory utilization shows the area reflecting active data use below the yellow Recommended Size line. The area above this denotes dormant DRAM that is underutilized and negatively impacting TCO.

Solution Value: More Memory, Lower TCO for VMware and Microsoft SQL Server

By dramatically increasing the physical memory pool with Intel Optane DC persistent memory, system and database administrators have new options. Especially when dealing with large workloads, they will have the ability to add more VMs to each server, which in turn will drive up CPU utilization from those sub-optimal averages ParkMyCloud noted. Having more VMs per server will allow admins to deploy fewer servers, saving on infrastructure and power costs and driving down TCO. Simultaneously, admins can future-proof their infrastructure against rising database sizes with an upgrade process built on a 2nd Generation Intel Xeon Scalable processor accompanied by Intel Optane DCPMMs that adhere to the same form factor and slot design as conventional DDR4 DRAM modules.

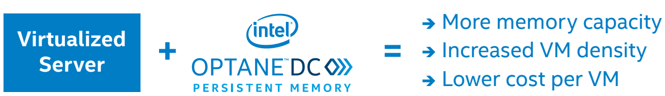

The value of Intel Optane DC persistent memory can continue beyond these initial benefits. Expanding the memory capacity of one server with Intel Optane DCPMMs helps enable more VMs to run from the system, thus increasing the VM density per CPU. Lower socket and system counts helps to lower licensing needs. The end result is lower cost per VM, lower licensing expense, and significant hardware savings.

While a DRAM-only two-socket platform will hit a ceiling at 3 TB of memory, Intel Optane DC persistent memory can already scale to 6 TB. These higher memory capacities can be very beneficial in cases where CPU resource demands may be low while memory needs are high. Such situations yield very low CPU utilization for the entire system, yet admins can't increase VM counts due to memory capacity limits. Extending that memory capacity allows for the addition of more VMs, which allows higher, more efficient use of physical CPUs. And again, the cost reductions on total memory made possible by Intel Optane DCPMMs allows for lower licensing costs and lower TCO for a Microsoft SQL Server platform.

Now, let's take a different look at this picture through the lens of virtualizing a large memory workload. As workloads require increasingly large memory footprints, and performance requirements in the form of service level agreements (SLAs) make overflowing from memory into storage prohibitive, virtualized systems need a solution that will maintain performance while accommodating these expanding datasets. Shown in Figure 3, Intel testing demonstrated the business value of Intel Optane DCPMMs by running side-by-side VMware vSphere systems, one with 768 GB of DRAM (24 x 32 GB) and one with 192 GB of DRAM (12 x 16 GB RDIMMs) and 1 TB of Intel Optane DC persistent memory (8 x 128 GB modules).

In both cases, the systems ran a modified version of VMmark for Incremental Memory benchmark designed to simulate outsized memory workloads. During the test, the number of VMs in each configuration was increased, seeking the point where VMmark's test suite SLAs could no longer be maintained. The DRAM-only platform supported 152 VMs at SLA-compliant performance levels, while the platform bolstered by Intel Optane DCPMMs supported 190 VMs—a 25 percent increase in VMs per node—at the same performance level.[3]

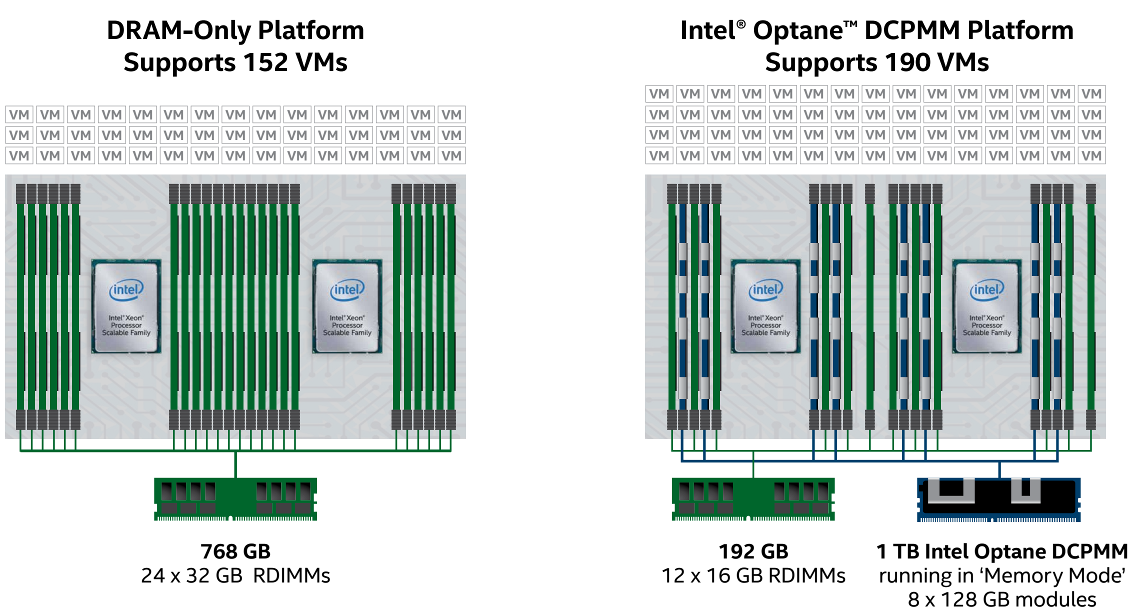

Figure 3. Example of VM support using 2x 2nd Gen Intel® Xeon® Platinum 8270 processors with and without Intel Optane DC persistent memory.[4]

The second configuration also realized up to 17 percent lower hardware cost per VM. The significant costs saved will help free up budget that can be applied to IT innovation for enterprise-wide growth.

In short, adoption of Intel Optane DCPMMs into VMware-virtualized and similar environments offers numerous benefits, including:

- Reduced IT costs per VM
- Ability to maintain SLA quotas due to Intel Optane DCPMMs operating at near-DRAM-level performance
- Greater VM density, yielding more customers per server

Conclusion

VMware-based or similar environments, especially those hosting database applications such as Microsoft SQL Server, can experience a wide array of benefits from expanding memory capacities beyond DRAM limitations with Intel Optane DC persistent memory. At a time when technology trends demand VMs to support ever-larger sizes with no sacrifice in performance, Intel's memory breakthrough provides a once-in-a-generation opportunity to surpass traditional architecture limitations, consolidate server resources (if desired), and significantly reduce TCO liabilities.

Intel Optane DC persistent memory is an affordable solution that drops right into existing DDR4 memory slots. IT can right-size DRAM allocations for the great majority of active datasets while using persistent memory modules to allow for workload surges and service an application owner's memory requests. IT can expand total memory per server, accommodate much larger workloads, increase VM density, and dramatically improve resource utilization.

BREAKTHROUGH DENSITY

DRAM cache plus high-performance Intel® Optane™ DCPMMs deliver outsized memory capacity for up to 25% higher VM density with VMware.[3]

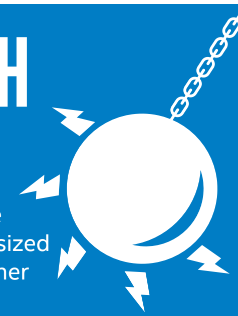

These benefits fit perfectly with enterprises in the midst of their digital transformation. Companies working to modernize their infrastructure and increase their efficiency to meet the mounting wave of data demands and analyze it all need the new capabilities and benefits that Intel Optane DC persistent memory can bring.

Learn more about **Intel Optane DC persistent memory**. See VMware and Intel's partnership in action in the **Intel Optane DC persistent memory video**.

More Memory for Virtualized Environments

Intel® Optane™ DC persistent memory uses a novel form of non-volatile media rather than DRAM or the NAND flash found in conventional SSD storage. Intel Optane technology has the data-retention properties of NAND but endurance and performance properties much closer to those of DRAM. These combined characteristics are unique in the data center world and offer exciting opportunities to extend functionality and platform value in ways never before possible. When combined with Intel® firmware, driver, and memory/storage management software, Intel Optane DC DCPMMs enable some compelling use cases.[5]

Broadly speaking, Intel Optane DC persistent memory can be configured into three modes: Memory Mode, App Direct Mode, and Dual Mode. Memory Mode provides a single volatile memory pool in which the DRAM installed in the system acts as a cache to deliver DRAM-like performance for DCPMM as the high-capacity main memory. While DRAM remains invisible to the user, all of the persistent memory modules' capacity appears to the operating system as addressable memory.

The hypervisor abstracts these resources and makes them available as virtualized RAM to VMs. Thus, in Memory Mode, the process of VMware interacting with Intel Optane DC persistent memory resources in a VMware vSphere ESXi hypervisor host is straightforward. Overall, Memory Mode helps drive down IT costs, increase VM density, and improve platform scalability.

App Direct Mode makes both DRAM and Intel Optane DCPMMs visible to the host, providing a larger potential space for system memory. This mode requires application enabling to decide which memory type to use for certain tasks, which can yield to higher performance. Also, data on the DCPMM remains persistent in App Direct Mode, which can help accelerate reboot/reload times, improve application availability, and drive down data center server footprint, which helps lower power consumption, infrastructure costs, and TCO.

Dual Mode allows some Intel Optane DCPMMs to be provisioned in Memory Mode and the rest in App Direct Mode. Dual Mode still provides excellent memory performance without incurring the latencies inherent in disk swapping across the I/O bus.

Learn More

You may also find the following resources useful:

- Intel Optane DC Persistent Memory
- Intel Optane Technology for Data Centers
- 2nd Generation Intel Xeon Scalable Processors
- Intel Optane DC Persistent Memory "Memory Mode" Virtualized Performance Study

Solution Provided By:

[1] Flexera blog, November 13, 2017. "Where Is the $10B Waste in Public Cloud Costs?" flexera.com/blog/cloud/2017/11/where-is-the-10b-in-waste-in-public-cloud-costs

[2] ParkMyCloud blog, December 4, 2018. "10,000 Years of Data Says Your Server Sizing is Wrong" parkmycloud.com/blog/server-sizing

[3] **Baseline configuration (DRAM):** 2x 2nd Gen Intel® Xeon® Platinum 8270 processor @ 26 cores/socket; BIOS: C2030.BS.1C03.GN1; Memory: 768 GB (24x 32 GB DDR4 @ 2666 MHz); Network: Intel® X520 SR2 (10 Gbps); Storage: 8x Samsung PM963M.2 960 GB, 4x Intel® SSDs S4600 (1.92 TB).

New configuration (DCPMM): 2x 2nd Gen Intel Xeon Platinum 8270 processor @ 26 cores/socket; BIOS: C2030.BS.1C03.GN1; Memory: 192 GB (12x 16 GB DDR4 @ 2666 MHz) + 1 TB Intel® Optane™ DC persistent memory (8x 128 GB @ 2666 MHz); Network: Intel® X520 SR2 (10 Gbps); Storage: 8x Samsung PM963M.2 960 GB, 4x Intel SSDs S4600 (1.92 TB).

Performance results are based on testing as of January 31, 2019 and may not reflect all publicly available security updates. Hypervisor: VMware vSphere ESXi 6.7.0 U1 10764712/10-16-2018. Datacenter management software: VMware vCenter 6.7.0.20100 Build 11338176/12-20-2018 with Intel driver for Intel X722: VMW-ESX-6.7.0-intel-nvme-vmd-1.6.0.1003.

[4] DRAM/DCPMM configurations can vary by ratio, capacity, and other factors as detailed in the Intel Optane DC Persistent Memory Implementation Guide.

[5] Intel support information, intel.com/content/www/us/en/support/articles/000023994/memory-and-storage/intel-optane-memory.html

The VMmark benchmark tests described in this paper were performed under VMmark's Academic and Research rules. They are not compliant VMmark tests and can therefore not be compared with published VMmark results.

Intel® Optane™ DC persistent memory pricing and DRAM pricing referenced in TCO calculations is provided for guidance and planning purposes only and does not constitute a final offer. Pricing guidance is subject to change and may revise up or down based on market dynamics. Please contact your OEM/distributor for actual pricing. Pricing guidance as of July 28, 2019.

 1219/CWAN/KC/PDF 341249-001US