SOLUTION BRIEF

Data Center
High-Performance Computing

Explore New Frontiers with Flex Memory for HPC

Control cluster cost and dynamically choose between extended memory or fast storage with Intel® Optane™ DC SSDs with Intel® Memory Drive Technology—and start solving your most pressing HPC problems

Solution Benefits

- **Flexibility.** Affordable memory or very fast storage on demand, using a software-defined infrastructure.
- **Cost-per-performance.** With lower cost, but up to 2x more memory or storage capacity, every infrastructure dollar invested works harder.[1]
- **Affordability.** Aggregating specialized nodes into flex memory systems can reduce memory or storage costs by up to 25 percent.[2]

Executive Summary

Some high-performance computing (HPC) workloads have extreme memory requirements; others need low-latency access to local storage. It can be expensive to provision fat memory nodes with ample DRAM along with separate nodes that provide fast I/O.

HPC data centers can aggregate these specialized nodes with a software-defined flex memory solution based on Intel® Optane™ DC SSDs with Intel® Memory Drive Technology. Without any physical changes to the hardware or to applications, you have access to a flexible but unified, on-demand infrastructure:

- **Memory-hungry workloads.** The job scheduler can choose to load Intel Memory Drive Technology prior to booting the OS, transparently integrating the Intel Optane DC SSD into the memory subsystem and making it appear like DRAM to the OS and applications.
- **Storage latency-sensitive workloads.** The job scheduler boots the OS without loading Intel Memory Drive Technology, making the entire Intel Optane DC SSD capacity available for low-latency I/O transactions.

Deploying flex memory can provide up to 2x greater storage or memory capacity[1] while lowering storage or memory costs by as much as 25 percent.[2]

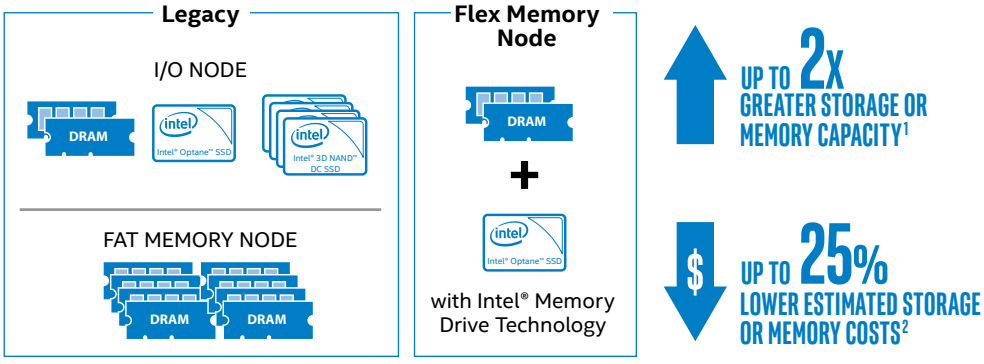

Figure 1. Intel® Optane™ DC SSDs with Intel® Memory Drive Technology can lower high-performance computing (HPC) costs, increase available storage or memory capacity, and run workloads that were previously impossible.

Business Challenge: Meeting Memory and Storage Needs without Breaking the Bank

Data volumes are exploding across all industries, such as scientific modeling and simulation, computational chemistry, astrophysics, genomics analysis, gas and oil exploration, and many more. High-performance computing (HPC) data center clusters are expanding to keep up. While standard compute clusters can handle the typical HPC workload, certain applications demand high memory capacities, while others require fast access to local storage. In fact, Intel estimates that about 10 percent of a legacy HPC infrastructure is dedicated to these "islands" of special-needs hardware.[3]

But budgets are not necessarily expanding at the same rate as data, and adding high-end I/O nodes and fat memory nodes to clusters can be expensive, especially if extending memory requires an upgrade to a multi-socket platform (see Figure 2). What's more, modern in-core HPC workloads can quickly outstrip even the largest legacy memory configurations. If an application cannot fit the entire dataset into memory (which is typical for legacy applications that have not been optimized), the application must either use a distributed implementation or spill data to disk. The former approach can compromise accuracy because block operations may be performed with partial updates; the latter approach can negatively affect application performance.

Today's HPC data centers need a solution that alleviates both of these types of overhead by providing a way to affordably provision high memory capacity or fast storage capabilities, as needed. And Intel® Optane™ DC SSDs with Intel® Memory Drive Technology can do just that.

Solving Today's Scientific Questions with Innovative Intel® Technology

HPC customers are using Intel Optane DC SSDs with Intel Memory Drive Technology in a wide variety of applications that have large memory needs and/or low-latency local storage requirements. These applications include computational chemistry and biology, fluid dynamics, astrophysics, pharmaceutical research, medical imaging, atmospheric research, oceanography, and genomics assembly. With more memory available, researchers take advantage of innovations such as virtual reality (VR) visualization technology that can enhance analysis.[4]

For example, through the Texas Advanced Computing Center (TACC), the University of Texas at Dallas (UTD) is using this technology to run quantum out-of-core workloads. Physicists have used Intel Optane DC SSDs with the Houdini* application to significantly decrease render time for a complex physics model. And a paper published by the Dell EMC HPC and AI Innovation Lab states that the performance of Intel Optane DC SSDs with Intel Memory Drive Technology is "impressive."[5]

Another emerging use case for Intel Optane DC SSDs with Intel Memory Drive Technology is containerization. Executing HPC workloads on a single system enables a transition to cloud-based HPC and high-throughput computing (HTC) as an alternative to HPC clusters that use Message Passing Interface (MPI).[6]

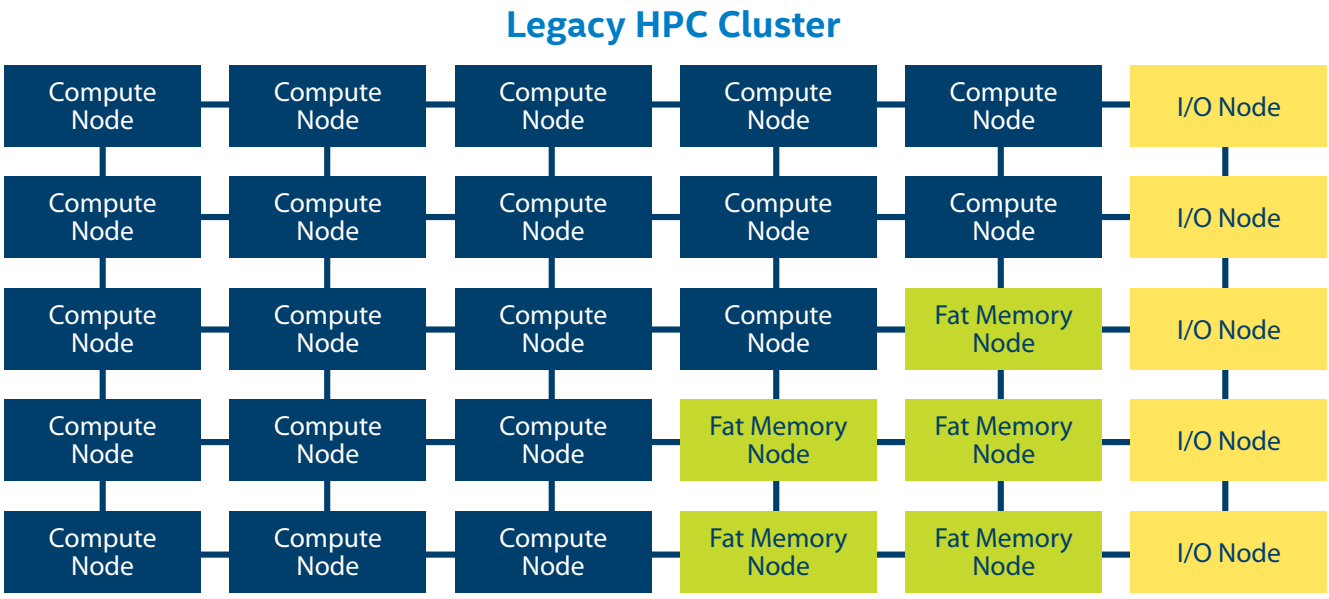

Figure 2. The need for both fat memory and I/O nodes in a legacy high-performance computing (HPC) cluster can drive up costs and hinder cluster flexibility.

Solution Value: Affordable, Universal HPC Cluster Architecture

Imagine an HPC cluster that can convert from a fat memory configuration to a low-latency local storage (I/O) configuration—on demand and without having to make any changes to the hardware or to existing applications. That is exactly how Intel Optane DC SSDs with Intel Memory Drive Technology work. Such flexibility enables a unified, software-defined flex memory system design that can do the work of both special configurations, as well as run HPC workloads that don't need either fat memory or high-performance storage.[7] The flex memory system can be integrated with cluster management tools and workload schedulers, which can use automation to choose and provision the appropriate configuration during a reboot cycle.

With ever-increasing data volumes, provisioning enough DRAM to hold massive datasets in memory isn't even feasible. But with flex memory, you can affordably allocate 2 TB, 6 TB, or even 12 TB of memory for your most demanding HPC workloads. And then, when those jobs are complete, the next job can use the same hardware to perform low-latency I/O if necessary. There is no capital expense involved in switching back and forth between memory and storage modes—all it takes is a quick reboot of the server.

In addition to this on-demand flexibility, Intel Optane DC SSDs with Intel Memory Drive Technology can also lower cluster costs. By aggregating the fat memory and I/O node islands into a single, unified infrastructure, you can provision flex memory for 10 percent of your servers for up to 25 percent less cost than provisioning more DRAM for just the fat memory nodes.[8] Plus you get a 2x increase in memory or storage capacity.[9] With the savings, you can deploy more flex memory nodes, which can improve the cost-per-performance ratio for your HPC data center.

Software developers now have a choice about how to develop applications that process massive datasets. They may keep all the data within addressable memory, since Intel Optane DC SSDs with Intel Memory Drive Technology help to reduce the constraints on memory capacity. Or, if memory bandwidth is an issue, developers can design their applications to run off of local storage, taking advantage of the low latency of Intel Optane DC SSDs.[10]

Solution Architecture: Software-defined Memory and Storage Extension

Intel Memory Drive Technology is software that is loaded on top of Intel Optane DC SSDs. It transparently converts some of the SSD capacity into memory. In memory mode, Intel Optane DC SSDs with Intel Memory Drive Technology create a transparent memory pool that can be as large as you need it (up to 12 TB per socket, which would mean 24 TB of memory on a two-socket system). Or, in storage mode, use the Intel Optane DC SSDs alone for fast storage access (see Figure 3). Intel Optane DC SSDs deliver an impressive combination of low latency, high endurance, quality of service, and high throughput.

No hardware or OS changes are necessary—all that is required is a server reboot with instructions to either load Intel Memory Drive Technology from BIOS or UEFI before booting the OS (for memory mode) or not (for storage mode). With pre-scheduled workloads, the job scheduler can be programmed to automatically choose the appropriate mode for each workload.

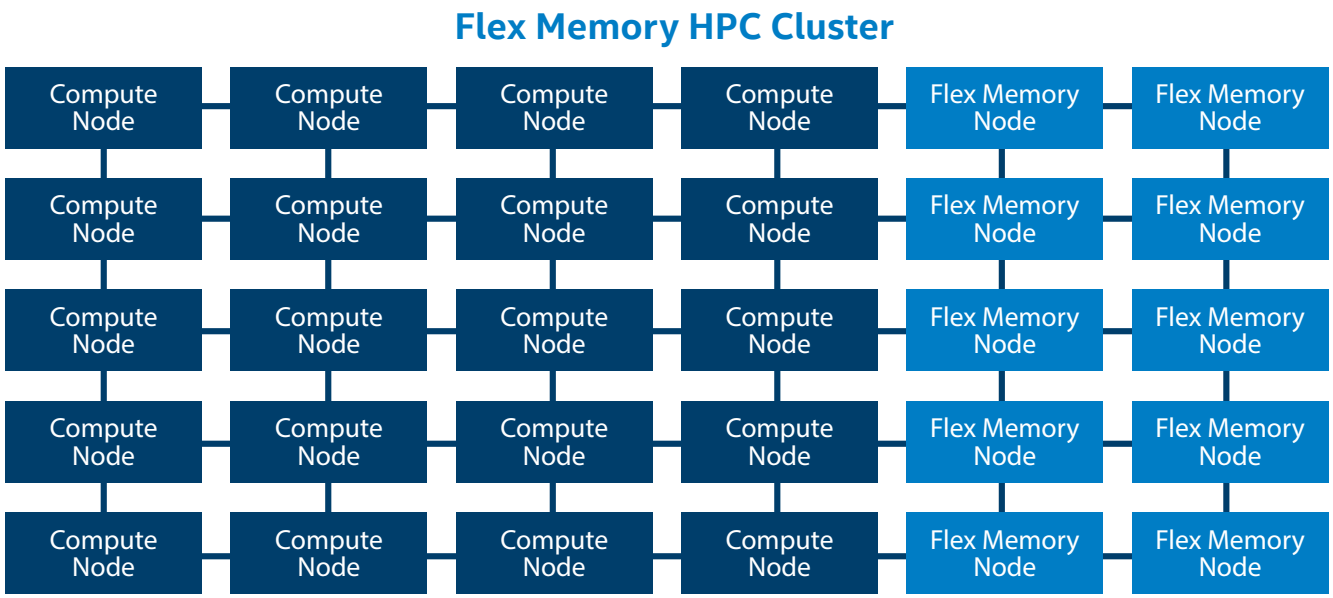

Figure 3. A flex memory system uses Intel® Optane™ DC SSDs with Intel® Memory Drive Technology to provide a flexible, unified infrastructure that can satisfy both extreme memory and I/O-intensive high-performance computing (HPC) workloads.

Intel® Optane™ technology offers several avenues for increasing memory and storage. For HPC users who are running mid-range (below 6 TB) memory systems, a flex memory solution is a first step. A cluster can then be upgraded to Intel® Optane™ DC persistent memory if higher performance or persistency is required. For extreme memory capacity situations (over 6 TB), flex memory is currently the only way to provide the necessary memory.

Conclusion

More data. Bigger problems to solve. And tight IT budgets. Sound familiar? Solve your memory and storage woes by deploying Intel Optane DC SSDs with Intel Memory Drive Technology. This flex memory solution can provide extra memory (without the DRAM cost burden) or deliver low-latency storage—all in a single solution that can actually lower HPC infrastructure costs. Flexibility and affordability powered by innovation from Intel enable scientists and researchers to explore new HPC horizons, running in-memory workloads that were never before possible.

Find the solution that is right for you. Contact your Intel representative or visit **intel.com/optane.**

Learn More

You may find the following resources helpful:

- Intel® Memory Drive Technology
- Intel® Optane DC SSDs
- Benefits of Intel Memory Drive Technology for Scientific Applications
- Intel® HPC Products Portfolio

For further reading, explore the following papers and blogs:

- Evaluation of Intel Memory Drive Technology Performance for Scientific Applications
- IMDT for Containers
- Toward Footprint-Aware Power Shifting for Hybrid Memory Based Systems
- SPAdes assembler test with Intel® Optane™ DC P4800X and Intel® Memory Drive Technology
- Studies on the energy and deep memory behavior of a cache-oblivious, task-based hyperbolic PDE solver

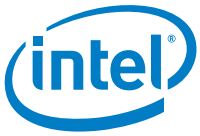

[1] 2x greater storage or memory based upon flexibility to configure an Intel® Optane™ SSD DC P4800X with Intel® Memory Drive Technology as either 100 percent storage or 100 percent memory, from a respective baseline of 50 percent for each memory and storage.

[2] Source: Memory4Less: Pricing as of May 21, 2019.
Baseline cost for fat memory nodes and I/O nodes: Fat memory nodes: Five nodes * 16x Samsung M393AAK40B42-CWD70* 128 GB PC4-21300 DDR4-2666 MHz ECC Registered @ $1422.76/each: memory4less.com/samsung-128gb-ddr4-pc21300-m393aak40b42-cwd70 + 16x Samsung M393A2K40DB2-CTD* 16 GB PC4-21300 DDR4-2666 MHz ECC Registered @ $186.89/each: memory4less.com/samsung-16gb-ddr4-pc21300-m393a2k40db2-ctd. I/O nodes: Intel: manufacturer's suggested retail price (MSRP) as of May 21, 2019 for Intel® SSD DC P4610 1.6 TB. $659/each * 5 nodes. ***Total cost = about $119,000.***
New configuration: Intel and Memory4Less: 10 nodes * Intel® Optane™ SSD DC P4800X MSRP as of May 21, 2019: 2x Intel® Optane™ SSD DC P4800X 750 GB with Intel® Memory Drive Technology configured to 640 GB each = $6480; DRAM configuration: 12x Samsung M393A2K40DB2-CTD* 16 GB PC4-21300 DDR4-2666 MHz ECC Registered @ $186.89/each: memory4less.com/samsung-16gb-ddr4-pc21300-m393a2k40db2-ctd. ***Total cost = about $88,000.***

[3] Based on Intel internal estimates.

[4] A flex memory solution based on Intel® Optane™ DC SSDs with Intel® Memory Drive Technology is not well suited for HPC applications that are bound by memory bandwidth; these applications can better benefit from Intel® Optane™ DC persistent memory in memory mode.

[5] Testing by Dell EMC HPC and AI Innovation Lab as of March 2019. For the full paper, visit dell.com/support/article/us/en/04/sln316587/spades-assembler-test-with-intel-optane-dc-p4800x-and-intel-memory-drive-technology.
All-DRAM configuration: Dell EMC PowerEdge* R940 server equipped with the following: 4x Intel® Xeon® Platinum 8168 processor (24 cores @ 2.70 GHz), 1.5 TB DRAM (48x 32 GB @2666 MHz), logical processor and virtualization technology enabled, OS = Red Hat Enterprise Linux* (RHEL*) v7.4, kernel = 3.10.0-693.el7.x86_64, BIOS system profile = performance optimized.
Intel® Memory Drive Technology configuration: Dell EMC PowerEdge* R740xd server equipped with the following: 2x Intel® Xeon® Platinum 8168 processor (24 cores @ 2.70 GHz), 24x 16 GB @2666 MHz with 3 TB software-defined memory (provided by 12x Intel® Optane™ SSD DC P4800X 375 GB), logical processor and virtualization technology enabled, OS = Red Hat Enterprise Linux* (RHEL*) v7.4, kernel = 3.10.0-693.21.1.el7_lustre.x86_64, BIOS system profile = SDM/default.
Workload: SPAdes* v3.10.1, Python* 2.7.13. The data used for the tests is a paired-end read, ERR318658 which can be downloaded from the European Nucleotide Archive (ENA). The reads are generated from a blood sample as a control to identify somatic alterations in the primary and metastatic colorectal tumors. This data contains 3.2 billion beads with a read length of 101 nucleotides.

[6] For full details, read the paper from the University of Pisa at itc.unipi.it/wp-content/uploads/2019/01/ITC-TR-01-18.pdf.

[7] If Intel® Optane™ DC SSDs with Intel® Memory Drive Technology are installed in a configuration, but Intel Memory Drive Technology is not loaded during boot, there is not a negative impact on DRAM performance.

[8] See endnote 2.

[9] See endnote 1.

[10] For a technical explanation of this concept, see the paper, “Toward Footprint-Aware Power Shifting for Hybrid Memory Based Systems,” at oaciss.uoregon.edu/icpp18/publications/pos118s2-file2.pdf.