Storage Class Memory Improves the Performance of Key Workloads

By Peter Burris

September 6, 2019

As enterprises consider long term IT strategies, they are considering how best to modernize the high-value traditional applications (HVTAs) that run the business and still consume 40%-50% of IT spending in large enterprises. Storage Class Memory (SCM) and NVMe-oF hardware technologies will make HVTA modernization approaches simpler.

New types of digital experiences, analytics workloads, business automation, and market integration increase transaction and data loads and the pathway complexity on high-value traditional applications (HVTAs). These new capabilities are generating new requirements to modernize high-value traditional applications (HVTAs) with technologies that offer superior performance and scale options.

Technologies intended to increase the productivity of HVTA infrastructure, like virtualization, workload parallelism, and in-memory DBMSs, indeed are increasing hardware utilization and aggregate throughput. However, they can also generate unintended complexity and cost challenges within the HVTA infrastructure, especially storage resources. The performance and reliability of each of these technologies are very sensitive to I/O performance within the storage hierarchy, from external devices to system memory. New storage device types, like flash memory-based solid state drives (SSD) and NVM Express over Fabrics (NVMe-oF) are helping to enhance system I/O balance, but the latency gap between I/O and memory is widening as distributed computing demands increase.

Consequently, increased HVTA loads are generating a "legacy butterfly effect," whereby small changes in data latency within the HVTA (and other applications) can have far-reaching effects on the performance of distributed systems serving customer experience, smart system, and business automation activities (see Figure 1). This is especially true as applications are refactored to selectively increase business agility and exploit emerging applications are refactored to selectively increase business agility and exploit emerging cloud technologies, thereby increasing the count and distribution of independent software machines participating in crucial business work. Moreover, as enterprises move to container and related technologies, the legacy butterfly effect may only get worse.

STORAGE RELIEF FOR LATENCY

Today's IT professionals who own and operate the high-value traditional applications (HVTAs) running the business face daunting decisions. The business wants a cloud experience because it promises greater options for scale, operating costs, and digital value creation. However, the throughput demands of increasingly complex distributed systems are pushing system architects to undermine cloud ideals and sustain traditional approaches to locating data across system resources including frequently accessed data cached in memories and volume data organized for bulk transfer on storage devices. Given that new classes of applications like data-first AI systems) can require frequent read access to bulk-like data (which exacerbates data latency problems), traditional

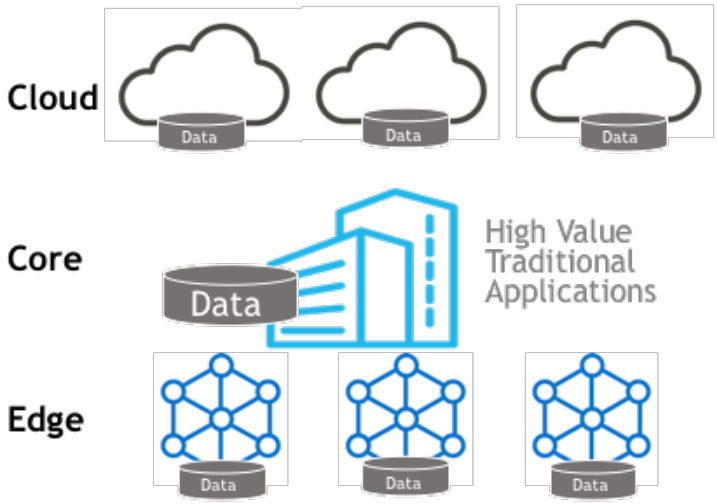

Figure 1 - Greater distribution of Data and Computing Increases Sensitivities to HVTA latencies

approaches to mapping physical data schemas to storage hierarchies are fraying.

Today's IT professionals who own and operate the high-value traditional applications (HVTAs) running the business face daunting decisions. The business wants a cloud experience because it promises greater options for scale, operating costs, and digital value creation. However, the throughput demands of increasingly complex distributed systems are pushing system architects to undermine cloud ideals and sustain traditional approaches to locating data across system resources including frequently accessed data cached in memories and volume data organized for bulk transfer on storage devices. Given that new classes of applications (like data-first AI systems) can require frequent read access to bulk-like data (which exacerbates data latency problems), traditional approaches to mapping physical data schemas to storage hierarchies are fraying.

In addition, HVTAs typically aren't candidates for simple migration to public cloud environments. This is, in part, because HVTAs have been finely tuned to minimize the impacts of data latency, sometimes across a decade of technology implementations, thereby presenting practical migration barriers. Our research shows that enterprises are pursuing a variety of HVTA modernization approaches. However, we expect that one approach will be especially fruitful for at least the next 3-5 years: Investing in infrastructure that improves cloud affinity.

A variety of innovative technologies are capable of undergirding modern cloud-oriented system and application software to better serve the data requirements of data-first AI workloads. One, in particular, is Storage Class Memory (SCM). SCM adds a new tier to the storage hierarchy, one that features memory-like performance at storage-like cost.

SCM reduces data latencies by providing memory-like byte-addressable access to data, thus reducing the need to read and perform garbage collection on large blocks of data. However, unlike system memories (and like storage devices), SCM is persistent. Large amounts of byte-addressable data can be mapped across CPU-side resources and accessed using load and store instructions. The lengths of I/O paths to crucial data is dramatically reduced when that data is placed into SCM. Testing performed by Dell and Intel using Dual Port Intel Optane SSD devices shows a 10x drive-level latency improvement when compared to NAND flash drives. This is a dramatic improvement that can push access times into the low hundreds of nanoseconds, a performance range that will help ameliorate the HVTA butterfly effect while opening up new possibilities for adding cloud-like experiences to HVTAs.

NEXT STOP: SCM IN STORAGE SYSTEMS – AND PRACTICES

Introducing a powerful, new storage technology like SCM into existing operations requires a degree of integration that exploits the technology without exploding complexity and compromising data security. Over the last few years, storage solutions have been marrying SSD devices and NVMe, resulting in a relatively seamless approach to enhancing storage performance and flexibility. But SCM will require an even deeper level of integration to ensure its performance and scale are easily within reach of applications, application developers, and system engineers.

We believe Dell EMC's PowerMax family of purpose-built, high-performance storage systems will be the first offering to holistically engineer Dual Port Intel Optane SCM as a persistent tier of storage, SSD, and NVMe-oF together. When announced in September 2019, our expectation is that PowerMax will be able to demonstrate storage systems capable of delivering up to 15 million IOPS, aggregate bandwidth of 350GB/sec, 50% better data response times for modern workloads, and a major step up in environmental factors. To simplify the use of SCM technology, PowerMax will offer user-defined service levels for predictable and consistent performance.

Consequently, IT organizations will have access to new approaches for marrying tried and true as well as new applications – and practices – together. The PowerMax family should expand options for consolidating file and block storage formats on a common storage subsystem, notably improving the fit between HVTA storage and other classes of application storage. While this will have the impact of simplifying storage administration, it also should lead to new levels of integration between data, opening up brand new application vectors for improving customer experience, automation, and digital business models.

ACTION ITEM

Storage Class Memory (SCM) technology begins to combine memory-class performance and storage-class pricing and ease-of-use. This combination will be essential to economically and strategically integrate high-value traditional applications (HVTAs) that run businesses and emerging data-first applications, like AI, that will set them apart. CIOs who want to deliver a cloud experience for HVTA extensions should look to solutions like Dell EMC's PowerMax family as an early enabler of SCM technologies like Intel Optane SSDs.

Peter manages Wikibon and directs research at SiliconANGLE Media. His own research focuses on digital business transformation, data value, and cloud optimization. He has over 30 years of experience as an IT practitioner, analyst, research leader, and executive in technology and business. Prior to joining Wikibon, Peter ran B2B CMO and CIO research teams at Forrester, was co-Research Director at META Group (now Gartner) where he drafted the seminal Adaptive Infrastructure Strategies program and has worked as an IT and marketing leader at large, global enterprises. He is a well-regarded speaker on technology infrastructure trends and digital business transformation. He graduated from Yale University and now lives in Silicon Valley, CA.

Peter Burris
@plburris
peter@siliconangle.com