BUSINESS BRIEF

Business Insight at the Speed of Thought

A paradigm shift in data processing that will change your business

accenture

Advanced analytics and the efficiencies of Hybrid Cloud computing models are radically altering the way businesses operate and compete. The use of advanced analytics, increasingly enabled by in-memory databases like SAP HANA that pair operational database and analytics database into one pool of real-time data that's current and actionable, is enabling businesses to improve the accuracy and velocity of the decision making. At the same time, hybrid cloud models are bringing efficiency to IT organizations by placing lighter weight workloads into the public cloud, virtualizing other on-premise workloads for rapid deployment, and allowing IT to focus their internal operations on high-value compute services such as advanced analytics.

When IT delivers advanced analytics capabilities, consider what's possible when the full spectrum of business data is available for instantaneous processing. Enterprises can have nearly immediate, holistic visibility into business.

A Flexible, Cost-Effective Solution to Turn Enterprise Data Into Value

Competitive business advantages	• Lower infrastructure and licensing costs • Access more data to improve decision making • Deepen business and operational insight • Deliver new services • Improve customer experiences

In-Memory Technology Enabling Action at the Speed of Thought

For advanced analytics to be actionable, the data set needs to be holistic and presented to the user at speeds synched with the pace of human thought. This is the promise of *in-memory analytics*, as outlined by Professor Dr. Alexander Zeier in his book *In-Memory Data Management; Technology & Applications*. "In-memory analytics allows analytics to be run on operational data, simplifying both the software and the hardware landscape, leading ultimately to lower overall cost." This defines the technology and TCO benefit, but Dr. Zeier goes further in defining the tipping point where actionable data is considered to be delivered in "real-time" to the decision maker. And it has actual numbers.

Human "Real-Time" Reaction Times

220 milliseconds	Average human reaction time from receiving a piece of stimulus to forming a response
384 milliseconds	Recognition reaction time. Response that includes understanding and comprehension
550-750 milliseconds	Addition of complex context. Generally considered at "the speed of thought"

This brings us to a fundamental rule. If the speed of thought is measured in milliseconds, then the latency of the system components providing the data must be measured in nanoseconds for the combined latency to be near the minimum threshold of 550 milliseconds.

You can then expand this line of thought out to workloads such as machine learning and training, the building blocks of artificial intelligence (AI). For the broad range of future AI applications to be considered "human," they'll need response times at the speed of thought or faster to be considered useful to humans.

In-Memory Technology Today

The building blocks of in-memory analytics technologies today fall into two categories. First, the real-time transactional applications capable of analyzing terabytes of data in real-time (SAP HANA for example) and second, the servers that run the transactional applications. For the purposes of this discussion, we break the later into the three major subsystems:

1. The microprocessor (less than 1 ns latency). Today, Intel® Xeon® processors power more than 90% of the world's data centers. The most sophisticated of these processors have up to 28 cores. Modern database application extract performance from these cores by disaggregating tasks and spreading them among the cores to process them in parallel. This simple example of processor "parallelism" has delivered such spectacular leaps in performance that the microprocessor is limited only by the rate of data delivered to it from system memory.
2. Memory (less than 100 ns latency). Mainstream memory solutions today are DRAM based and deliver data at 2666 MHz rates and in aggregate capacities measured in terabytes. But they're also volatile, meaning the data is lost after power cycles. This leads to latency increases when data needs to be reloaded into memory.
3. Storage (more than 1,000,000 ns latency; often measured in milliseconds). Storage today is a combination of spinning disks, 3D NAND, and a new type of media called 3D XPoint™ technology which powers Intel® Optane™ solid state drives and delivers dramatic reductions in latency. Storage capacity can be scaled up to petabytes of storage. It's also persistent, meaning the data remains in place after a power cycle.

As you can see from the latency numbers above, keeping data in memory and close to the CPU keeps latency in the nanosecond range. But once a request is made to storage, latency goes from nanoseconds to milliseconds. This is where Intel engineers have focused their efforts. How do we create a technology that delivers memory-class latency with the capacity and cost of storage?

Intel Persistent Memory Solution

The Intel persistent memory solution turns memory from a limited resource to an ample one that is fast, non-volatile, and affordable.

Alper Ikbahar, Intel's Vice President of Memory Solutions, sums up the tremendous advantage of Intel's efforts,

> "Memory has been thought of as a small, expensive, and volatile resource. And we're going to change all of those assumptions with Intel persistent memory. This new memory is going to be big, affordable, and persistent."

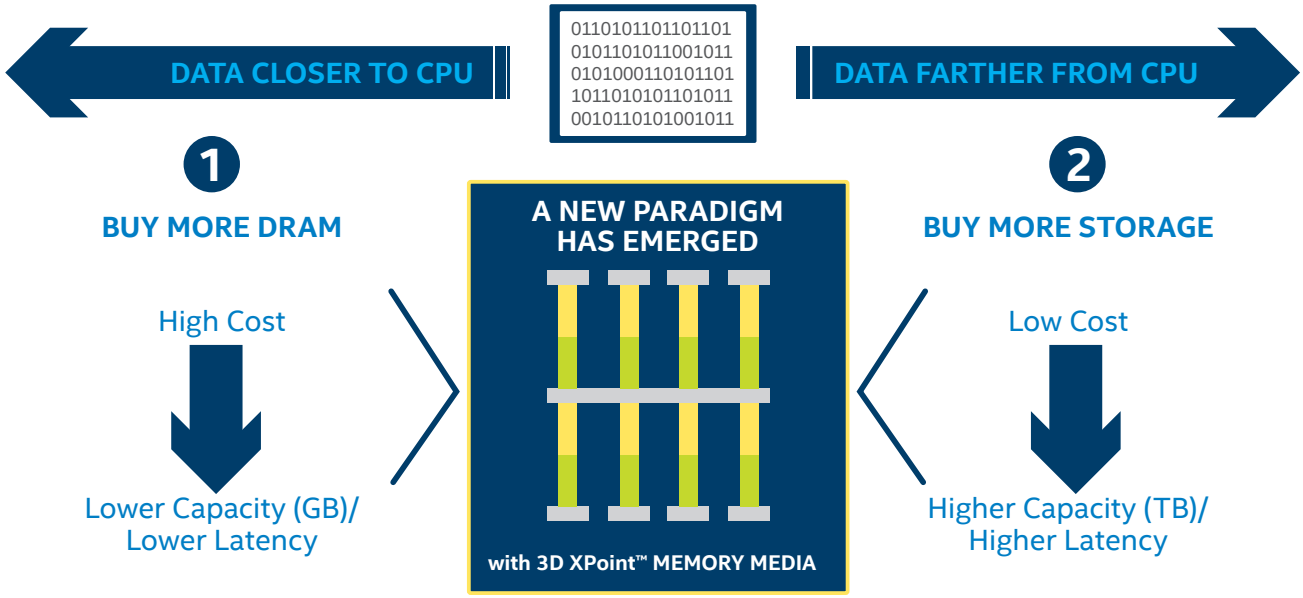

Intel Persistent Memory: This innovative new solution, launching in-market in the second half of 2018, will move larger amounts of data closer to the processor, so it can be accessed, processed, and analyzed in real-time (without first being retrieved from storage). The expectation is that this will speed time-to-insight from hours and days to seconds.

It fundamentally changes how much memory can cost-effectively be placed close to the processor. In turn, this fundamentally changes how much data can be processed in real time to accelerate effective decision making.

A 2017 article in Forbes claims that persistent memory will revolutionize computing.[1]

Sample Use Cases for Intel Persistent Memory

Because persistent memory technologies have the potential to enable faster processing of large data sets, they are poised to impact a wide range of industries, including scientific research, manufacturing, finance, oil and gas, retail, and healthcare. While we've focused on in-memory analytics, IT managers can deploy Intel persistent memory for workloads which are both memory intensive and I/O intensive as shown on the next page.

"Persistent memory is the missing link—as enterprises reshape infrastructure and business models to take advantage of the ubiquitous connectivity and intelligence of IoT."

MULTIPLE USE CASES IN VALUE ADD WORKLOADS

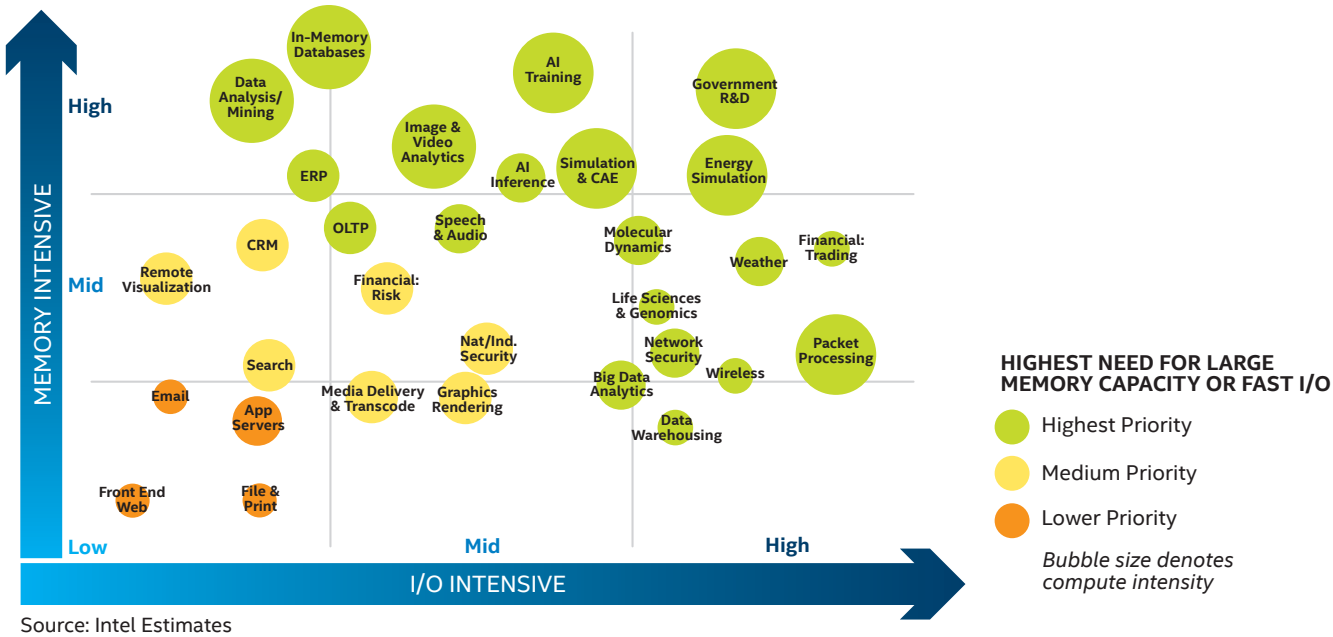

Source: Intel Estimates

IT managers should look for data center workloads that are both memory intensive and I/O intensive. These include a range of analytics workloads such as in-memory databases, enterprise resource planning (ERP), online transaction processing (OLTP), and artificial intelligence training and inference.

Tapping the Potential

While persistent memory might be a missing link, it's not a solution in and of itself. It won't replace traditional DRAM, nor will it replace storage. Rather, IT decision makers will need to strategically think about the right combination of traditional DRAM, persistent memory, and traditional SSDs to enable the real-time system.

Accenture and Intel can help enterprises to prepare for the shift to persistent memory technologies and the considerable advantages this will bring. This may include strategic planning and testing of new architecture to assess the capabilities and use cases that can help transform business.

As you build a foundation for the future to offer competitive services to customers, investment in Intel persistent memory technology will be crucial.

Businesses should consider:

- The kind of insight that would support business growth
- The size of data sets to place in-memory
- The platform that best suits your requirements

Let's talk about your current technology configuration and the options ahead for reducing costs and maximizing opportunity.

What's Available Now?

For in-memory analytics solutions, the software and hardware IT managers need to deploy is available today. Solutions like SAP HANA running on Intel® Xeon® processor-based platforms provide a robust framework from data analytics, but the TCO equation often doesn't add up. Provisioning terabytes worth of DRAM to support real-time analytics is too cost-prohibitive for the majority of business.

- **New Intel® Xeon® Scalable processors** bring a huge leap in performance for a wide range of workloads. The new processors feature 50% more memory bandwidth than previous generation Intel Xeon processors and new Intel® AVX-512 instructions that exploit data parallelism.
- **Intel® Optane™ SSDs** based on 3D XPoint™ technology offer high-performance storage for enterprise data centers.
- **(Coming in 2H 2018) Intel persistent memory** based on 3D XPoint technology will offer high-performance, high-capacity memory at lowered costs for enterprise data centers.

accenture

About Accenture

The fast pace of technology is creating seismic shifts for industries worldwide. Accenture helps customers to stay ahead of the digital curve and power transformation through new technology. With expertise across more than 40 industries and all business functions, Accenture partners with more than three-quarters of the Fortune Global 500.

Learn more

Explore Intel solutions for enterprise storage at **intel.com/optane** or contact us at **datacentermemory@intel.com**.

Find out how Accenture can help move your business forward by visiting **accenture.com** or contact us at **aic.sap.hana@accenture.com**.

Intel Persistent Memory Resources for Your IT Professionals

pmem.io

software.intel.com/pmem

blogs.intel.com

intel®

[1] forbes.com/sites/tomcoughlin/2017/01/24/persistent-memory-will-revolutionize-computing/#5d2377321912.

Technology claims are based on comparisons of latency, density, and write cycling metrics amongst memory technologies recorded on published specifications of in-market memory products against internal Intel® specifications. Results have been estimated or simulated using internal analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure. Check with your system manufacturer or retailer, or to learn more at intel.com/iot.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

* Other names and brands may be claimed as the property of others.

0318/SC/HBD/PDF ♻ Please Recycle 337214-001US