

# Realize the Promise of In-Memory Computing

Tech Mahindra drives the best usage for SAP Hana 2.0 with Intel® Optane™ DC persistent memory across all private/public cloud platforms.

## Authors Executive Summary

**Intel:**  
**Narinder Sharma**  
Industry Technical Specialist

**Tech Mahindra:**  
**Abhinandan Chakraborty**  
Technical Alliance Manager – ICS



Application performance is a requirement for succeeding in today’s business climate. Organizations rely on sophisticated applications, analytics, and database-management systems to provide a positive customer experience and to increase employee productivity. Customers demand constant availability and instantaneous responses, whereas employees rely on fast application performance from corporate databases and applications like enterprise resource planning (ERP) tools to get their jobs done as efficiently as possible.

Intel continually pushes the performance envelope with each new generation of Intel® Xeon® processor, system memory and Intel® SSDs. As a result, Intel® Optane™ DC persistent memory, a completely new class of memory, promises not only to continue the tradition of pushing the performance envelope, but to revolutionize application performance and capabilities.

The SAP HANA platform altered the database landscape by providing an in-memory database-management system. While main memory is fast and provides SAP HANA a performance advantage, it is also volatile and doesn’t retain its contents during a server reboot, scheduled downtime, or power outage. When a server reboots, all SAP HANA data must be reloaded from storage, which can be time consuming. Intel is changing the game for in-memory databases with Intel Optane DC persistent memory, a higher density, non-volatile memory technology that retains its contents like traditional solid-state drives (SSDs) and spinning disks, yet provides speeds that are similar to main memory.

SAP HANA 2.0 SPS 03 is the first major database platform to support Intel Optane DC persistent memory. Intel Optane DC persistent memory represents an entirely new means of managing data for demanding workloads like the SAP HANA platform. Intel Optane DC persistent memory is non-volatile, so in-memory databases like the SAP HANA platform do not have to completely reload all data from persistent storage to memory, and it runs at near-DRAM speeds, maintaining today’s performance expectations.

## Table of Contents

Executive Summary ..... 1

Transforming In-Memory Computing with SAP HANA 2.0 ... 1

Why Intel Optane Persistent Memory for SAP HANA 2.0? ..... 2

Why Tech Mahindra for Digital Transformation? ..... 3

Conclusion ..... 4

## Transforming In-Memory Computing with SAP HANA 2.0

In a fast-evolving landscape, enterprises are shifting the services they need or provide to cloud servers to save investing on hardware infrastructure and optimizing the people and process landscape. A big task now facing enterprises is to assess the cloud-readiness of their on-premises application portfolios, and discover how best to make their data centers more efficient and agile engines of digital transformation through deployed solutions. In this context, the delivery of enterprise cloud platform through software and hardware integration is becoming a new trend in the private cloud, hybrid cloud and multi-cloud market. Through the implementation of hyper-converged infrastructure (HCI), enterprises can centrally deploy and manage computing, storage, network, management resources, etc., thereby simplifying management

and significantly shortening the cloud deployment cycle while achieving high performance, high reliability, flexible expansion, and low cost.

The SAP HANA® platform provides state-of-the-art database

and data-management technologies, advanced analytical capabilities, and intuitive application-development tools in an all-in-one, in-memory-first data platform. SAP HANA 2.0 SPS 03 brings several additional innovations to the SAP HANA platform for intelligent enterprises.

**SAP HANA 2.0 SPS 03 contains numerous innovations to analytics, development, and the management of data and databases on the SAP HANA platform:**

- **Database Transformed:** SAP HANA 2.0 SPS 03 provides built-in real-time data anonymization that helps protect sensitive data in real time on a single instance of the data.
- **Analytical Intelligence Transformed:** SAP HANA 2.0 SPS 03 provides improved performance through parallel processing improvements for the training and scoring of predictive models.
- **Application Development Transformed:** A Advanced model (XSA) in SAP HANA 2.0 SPS 03 provide more options in the application layer to increase productivity with industry-standard runtimes.
- **Data Management Transformed:** Streamlined planning to execution through collaboration, intuitive design, and development automation of conceptual to physical modelling for hybrid cloud and on-premises environments

## Why Intel Optane Persistent Memory for SAP HANA 2.0?

Intel's new Persistent Memory called Intel Optane DC Persistent Memory, is a new type of computer storage, combining speeds approaching those of dynamic RAM (DRAM) along with RAM's byte-by-byte address-ability, plus the permanence of solid-state disks (SSDs). It resides on the memory bus which allows Persistent Memory (PMEM) to have DRAM-like access to data, which means you can achieve almost the same speed with slightly higher latency of DRAM alongside non-volatility of SSDs.

Support for Intel Optane DC persistent memory in SAP HANA 2.0 SPS 03 is a revolution for enterprise computing because it will change the way IT organizations think about data tiering. Currently, tiering data comes down to an unappetizing choice between investing in more expensive memory or degrading performance by keeping more data in persistent storage. Intel Optane DC persistent memory represents an entirely new means of managing data for demanding workloads like the SAP HANA platform.

Memory for databases is currently small, expensive, and volatile; Intel Optane DC persistent memory is denser, more affordable, persistent, and performs at speeds close to memory. These features of Intel Optane DC persistent memory can help drive lower TCO through reduced downtime and streamlined data tiering. These same features can also make SAP HANA in-memory databases economically viable for a wider range of use cases: Intel Optane DC persistent memory provides near-DRAM in-memory computing speed in a form factor similar to dual in-line memory modules (DIMMs) at a lower price per gigabyte than DRAM memory. Support for Intel Optane DC persistent memory will be available with the next-generation Intel® Xeon® processor Scalable family.

### Less Downtime

Because it is non-volatile, Intel Optane DC persistent memory enables you to keep the data in the SAP HANA platform loaded in main memory, even when power is off. Because you don't have

to reload the data back into memory after downtime, restart time for the SAP HANA platform is greatly reduced.

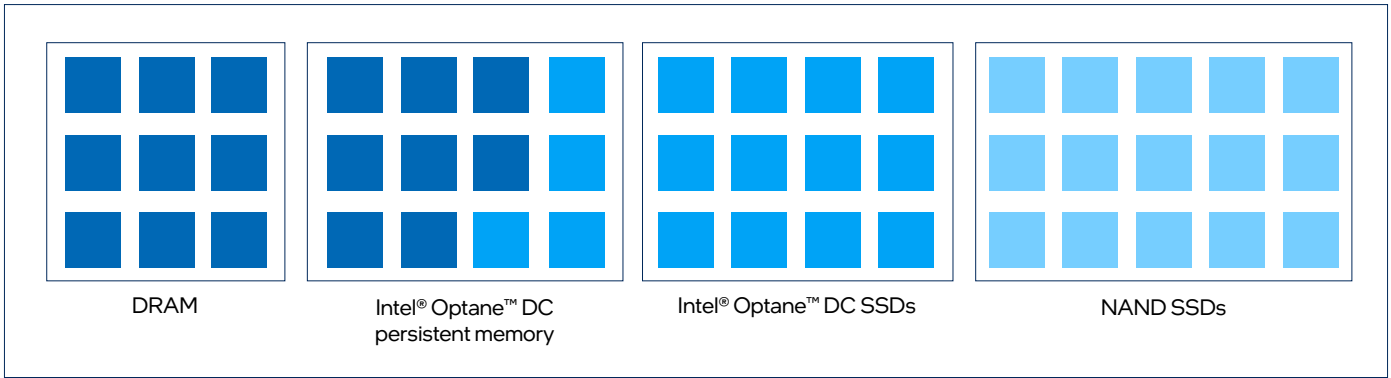
### Upending the Realities of Data Tiering

Due to the high cost of memory, traditionally only the most frequently accessed, most valuable hot data could be housed there. Less valuable warm data had to be stored outside of main memory at the price of lower performance. Intel Optane DC persistent memory and innovations in SAP HANA 2.0 SPS 03 change these economics, providing new options for data tiering.

Intel Optane DC persistent memory is available in a form factor called persistent memory modules (PMMs), similar to DIMMs, but with greater capacity than is available with conventional, volatile memory. Not only can using Intel Optane DC persistent memory reduce the overall price of memory for the SAP HANA platform, it also means that you can place your entire storage-area network (SAN)-based warm-data tier in data modules that act like main memory, which improves performance.

### Enabling a Flexible, Memory-Like Tier

Larger memory configurations in Intel Optane DC persistent memory have the potential to radically reshape data tiering and in-memory processing for your SAP HANA deployment. Larger, less-expensive memory configurations increase the economic viability of use cases that can benefit from the speed of insights available from in-memory databases but that cannot justify the high cost of investing in large amounts of DRAM. Combined with Intel Optane DC Solid State Drives (SSDs), Intel Optane DC persistent memory enables a new, flexible, memory-like tier: valuable data traditionally stranded in slow NAND storage can be activated and acted upon. Intel Optane DC persistent memory opens up new opportunities for companies to cost-effectively gain actionable insights from their data—and gain a competitive edge.



**Figure 1.** Example of updated data tiering using DRAM, Intel Optane DC persistent memory, and Intel Optane DC SSDs; hot data is represented in red, warm data is in orange, and cold data is in blue

### Revolutionary Innovations

SAP HANA 2.0 SPS 03 provides additional new enhancements for data tiering, one of which is native extension-node support. This native support in the SAP HANA platform enables a standard SAP HANA node to be used for warm data storage. Extension nodes are essentially scale-out nodes with relaxed processor and memory requirements for SAP HANA applications. Using extension nodes, a node from a scale-out cluster is set aside just for warm data; this extension node is allowed to store up to four times (as of today) as much data as a hot node in the cluster. Native extension-node support provides key advantages, such as fast in-memory performance when querying warm data, including support for all native data types and advanced analytics engines. This provides customers with an additional option for multi-temperature data management, helping optimize the performance/cost ratio of customers' SAP HANA systems. Features like this and Intel Optane DC persistent memory provide new, cost-effective means of deriving insights faster and from larger datasets.

### Why Tech Mahindra for Digital Transformation?

Tech Mahindra along with esteemed partners like Intel & SAP, helps enterprises accelerate their digital transformation with a comprehensive suite of offerings that covers the entire IT

infrastructure stack. With platform-enabled solutions powered by AI and Automation, Tech Mahindra helps enterprises accelerate their journey to a digital future focusing on helping organizations develop a future ready infrastructure aligned to their core business needs.

### Tech Mahindra fosters customer success through a 3-pillar strategy:

<b>Imagine</b>	with enterprises to define how their industries will change in the future and how they will proactively transform to be ready for that eventuality.
<b>Build</b>	systems and processes for enterprises so that they can handle revolutionary change.
<b>Run</b>	the business value of clients by helping them differentiate themselves in an increasingly crowded market.



## Leadership across verticals

 <b>Communications</b>	5G Slice to Price - Communications Value Chain Produce to Play - Media & Entertainment Value Chain	 <b>Retail</b>	Design to Delight - Retail & Consumer Goods Value Chain
 <b>BFSI</b>	Cradle to Grave - Banking, Financial Services Value Chain	 <b>Energy &amp; Utilities</b>	Wells to Wheels - Oil & Gas Value Chain Generation to Consumption utilities Value Chain
 <b>Manufacturing</b>	Concept to Customer Delight - Manufacturing Value Chain	 <b>Healthcare</b>	Delivering Future of Care Health & Lifesciences Value Chain

Learn more at <https://www.techmahindra.com/en-in/>

## Conclusion

The improvements and innovations in SAP HANA 2.0 SPS 03 can boost performance and help enterprises do more, more intelligently, with their data. These improvements include incremental steps for better analytics, faster application development, and more efficient data management. With these improvements also comes a radical step forward for database design and management with the SAP HANA platform: support for Intel Optane DC persistent memory.



### Disclaimers

Performance varies by use, configuration and other factors. Learn more at <https://intel.com/benchmarks>

Performance results are based on testing as of dates shown in configurations and may not reflect all publicly available updates. See backup for configuration details.

No product or component can be absolutely secure.

Your costs and results may vary.

Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel technologies may require enabled hardware, software or service activation.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries.

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