Solution Brief

Solving the Big Data Analytics Riddle

How solutions from Intel and SAP work together to enable dynamically tiered data storage and analytics

For decades, enterprises have relied on structured data stored in relational databases to drive business decisions. But the growth of new data-collection methods—in addition to ever increasing processing and storage capabilities—has opened up new avenues for gathering and analyzing immense data sets. Enterprises can collect data at far greater volumes and velocities than ever before, and they can then use that data to make critical real-time decisions.

However, older methods of analyzing the data stored in traditional disk-based relational databases can hamper this real-time agility. As data volumes grow, traditional analysis tools can bring databases to a crawl, which can result in the applications that rely on these databases becoming unresponsive. Enterprises require solutions that can scale and provide reporting and near-real-time, ad-hoc analysis on data sets that are hundreds of terabytes or even petabytes in size.

To address these challenges, Intel and SAP worked together to build an innovative, proof-of-concept system for storing and analyzing large-scale, high-velocity data sets. This solution brief describes the design philosophy and benefits of a multi-tier architectural approach, the system's hardware and software components, and the results of performance testing.

The Challenges of Big-Data Analytics

Today's enterprises can capture massive amounts of real-time data from thousands or even millions of sources, ranging from social-media sites to manufacturing-plant sensors. This broad range of collection sources can generate a variety of data types, including photographs, videos, and text. Traditional disk-based relational database management systems (RDBMSs), which require highly rigid schemas, might not be the best solution for collecting and analyzing such unstructured data. Complex ad-hoc queries over large database tables are often slow, which can frustrate users and impair decision making processes. In addition, poor application performance can impact users attempting daily tasks such as running reports or gathering statistics, which can ultimately affect data security. As enterprises push more data into traditional databases, frustrated users might turn to data-extraction methods to pull data from protected, primary databases into smaller, local databases, which can compromise security and privacy policies.

Another challenge that enterprises face is the problem of storing both new and old data within the same database. The frequency with which data is accessed tends to decrease over time. A plant manager might want to view real-time plant production statistics, while an auditor might need to view historical production statistics on a quarterly basis. In both of these cases, newer data is constantly combined with older data, and the database engine must query an increasing volume of data within the database. This can lead to diminishing performance that can limit critical real-time analysis.

The goal for Intel and SAP was to build an analytics system that would address the performance challenges of ingesting, storing, and organizing large, unstructured datasets while still maintaining query performance.

Design Philosophy: Dynamic Data Tiers

To meet the challenges of balancing performance versus velocity, volume, and variety, Intel and SAP based their system design upon the concept of dynamic data tiering for multi-temperature data. This design philosophy places frequently accessed data in fast, random access memory (RAM), while less frequently accessed data moves to slower, less expensive storage.

Multi-temperature data fall into three categories:

- "Hot" data is frequently accessed by an organization; for example, this could include data for real-time forecasting and reporting, or ad-hoc queries. Hot data often must reside in fast RAM as opposed to disk-based or even solid-state drive (SSD)–based storage to provide real-time results.
- "Warm" data is not accessed often, but it is used frequently enough that the data still needs to reside on disk- or SSD-based storage—such as the Intel® SSD DC P3700 Series—accompanied by a fast analytics engine. For example, a call center might need to run monthly call-statistic reports that don't require real-time results. Data often resides in this state until it is archived.
- "Cold" data is accessed sporadically and is considered archival. For instance, a manufacturing plant might not often need data captured from more than a few years ago, but it might want to keep such data available to show year-over-year performance gains from process improvements.

Project Goals

Project goals included ingesting hundreds of millions of rows per hour, supporting ad-hoc queries across millions of rows of data, providing exploratory and predictive-analysis tools, and dynamically managing data movement across multiple system tiers.

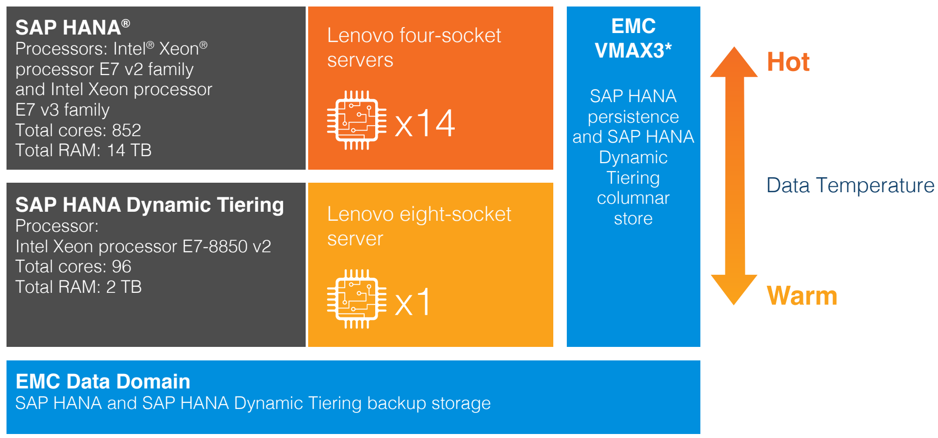

Figure 1. Architectural overview of the Intel and SAP analytics solution; data flows through hot and warm states based on pre-defined expiration criteria

For smaller data sets, all three data temperatures can reside on a single system or a small cluster of systems. But increased growth and usage might require a more complex, tiered approach using multiple levels of compute, memory, and storage.

While Intel and SAP focused on the multi-temperature design philosophy as the foundation for their system's architecture, only the "hot" and "warm" tiers were implemented and tested.

Architectural Overview

Intel and SAP designed and built a multi-tiered data-analysis and storage system that can handle thousands of users across the globe and hundreds of terabytes of data. As part of this proof-of-concept, specific performance goals included:

- Ingesting hundreds of millions of rows of data per hour from thousands of data-collection devices, and allowing querying of new data within 15 minutes
- Supporting ad-hoc queries and reporting across hundreds of millions of rows of data for thousands of global users without noticeable performance loss
- Providing integrated exploratory and predictive-analysis tools to drive better understanding of past activities and future trends
- Dynamically managing data movement by temperature across multiple storage and compute system tiers

The system consisted of two data-management and analysis tiers defined by using the data-temperature model, including:

1. A "hot" data tier built on SAP HANA® and servers powered by the Intel® Xeon® processor E7 v2 family and Intel Xeon processor E7 v3 family.[1] This tier uses EMC VMAX3* for persistent storage.
2. A "warm" data tier built on SAP HANA Dynamic Tiering and a server powered by the Intel Xeon processor E7 v2 family. This tier uses the same EMC VMAX3 storage systems for longer-term data storage.

Hot-Data Management: SAP HANA®

To meet the goal of ingesting hundreds of millions of rows of data per hour, Intel and SAP implemented SAP HANA on a cluster of servers powered by the Intel Xeon processor E7 v2 family and Intel Xeon processor E7 v3 family. SAP HANA is an in-memory database platform that differs from traditional, disk-based relational databases. All data is held in system memory, which provides access speeds that are orders of magnitude faster than disk-based databases. The in-memory technology of SAP HANA enables real-time analysis of large, high-velocity data sets, while the SAP HANA cluster configuration helps achieve the goals of ingesting and analyzing large amounts of data in real-time.

Warm-Data Management: SAP HANA Dynamic Tiering

While the hot data tier works well for real-time analysis of high-velocity data, the warm data tier needs to perform predictive analysis across relatively large data sets stored on less expensive storage. Intel and SAP implemented the dynamic tiering option of SAP HANA running on a server powered by the Intel Xeon processor E7-8850 v2 product family to meet these goals.

With the dynamic tiering option of SAP HANA, data can move out of system memory and onto less-expensive disk or SSD storage once it reaches a pre-defined expiration threshold. SAP HANA Dynamic Tiering makes use of the data-lifecycle-manager feature of SAP HANA Data Warehouse Foundation to move expired data to near-line storage, which is managed by SAP® IQ.

SAP HANA Dynamic Tiering enables large-scale, advanced analytics across structured, semi-structured, and unstructured data sets. It provides fast bulk loading capabilities and more cost-effective near-line storage for warm data.

Table 1. Hot-data-tier server configuration

Hot-Data-Tier Server Configuration	
Servers	14 four-socket Lenovo servers with the Intel® Xeon® processor E7 family[2]
Total Cores	852
Total Memory	14 TB

Table 2. Hot-data-tier network configuration

Hot-Data-Tier Network Configuration	
Internal Network	10 gigabit Ethernet (GbE) connection using a Lenovo RackSwitch G8272*
Client Network	10 GbE connection using a Lenovo RackSwitch G8272
Management Network	1 GbE connection using a Lenovo RackSwitch G8052
Storage Network	8 gigabit Fibre Channel connection linking the EMC VMAX3* family storage system with the SAP HANA® server cluster

Table 3. Warm-data-tier server and storage configuration

Warm-Data-Tier Server and Storage Configuration	
Server	One eight-socket Lenovo server[3]
CPU	Intel® Xeon® processor E7-8850 v2
Total Cores	96
Total Memory	2 TB
Storage	EMC VMAX3* family storage system

Table 4. Warm-data-tier network configuration

Warm-Data-Tier Network Configuration	
Client Network	Four 10-GbE connections aggregated into a logical 40-gigabit connection using a Lenovo RackSwitch G8272*
Management Network	1 GbE connection using a Lenovo RackSwitch G8052
Storage Network	Two 8-gigabit Fibre Channel connections linking the EMC VMAX3* family storage system with the SAP HANA® Data Tiering server

As a result, SAP HANA Dynamic Tiering provides a powerful and cost-effective intermediate storage and analytics platform for warm-data analysis.

Proof-of-Concept Performance Results

In many aspects, the proof-of-concept system exceeded the original design goals.[4]

- The system was able to achieve load speeds of 24 billion rows of data per hour, far exceeding the original goal of hundreds of millions of rows of data per hour.
- Loading data into the warm data tier used only 264 GB of memory and generated a peak utilization of 14 percent on the Intel Xeon processors, which let the system respond quickly to ad-hoc queries while the data load was taking place. This result aligned with the design goal of providing ad-hoc query services to users without noticeable performance loss while data was loading.
- In the hot data tier,1.8 billion rows of data were made ready for analysis in 88 seconds. In the warm data tier, aggregating and analyzing 1.8 billion rows using the SAP HANA Predictive Analysis Library (PAL) required less than two minutes. The extract, transform, and load (ETL) action for 2 billion rows took 4 minutes and 48 seconds, and it only required 88 seconds to make ready for analysis. The entire process finished in 6 minutes, which exceeded the original design goal of 15 minutes.
- A total of 800 billion rows of data were loaded, equal to 51 terabytes of uncompressed data. The EMC VMAX3 storage family can scale to a capacity of 4 petabytes and can achieve performance throughput of 55,643.78 megabytes per second.[5]
- The SAP HANA Data Lifecycle Manager moved 429 million rows from the hot data tier to the warm data tier in 2.5 hours, which aligned with the data-mobility design goal.

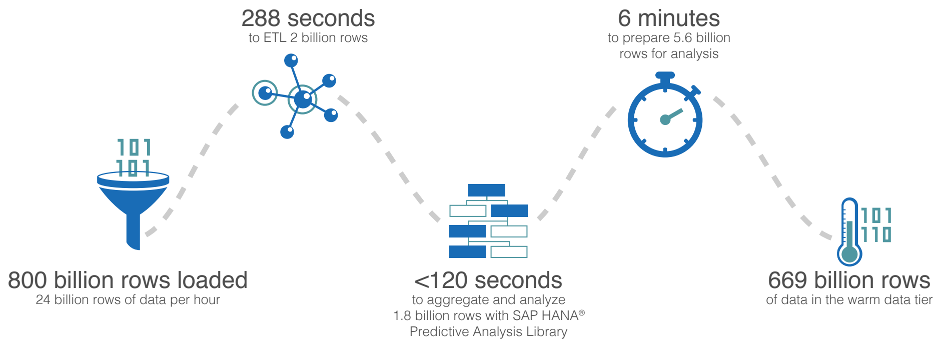

Figure 2. Intel and SAP help accelerate the process for data gathering and analysis while optimizing long-term data storage

The Intel and SAP combination of hardware and software provided a powerful foundation for ingesting and analyzing large, high-velocity data sets. The combination of SAP HANA running on a cluster of servers powered by Intel Xeon processors provided the means to ingest a significantly higher amount of data than expected. In addition, the proof-of-concept demonstrated that SAP HANA Dynamic Tiering running on Intel Xeon processors is capable of aggregating and analyzing large data sets with little to no effect on user-based ad-hoc queries.

Summary

Sophisticated data collection and analysis can help organizations gain a competitive advantage in today's data-driven business world. But the ever increasing velocity, volume, and variety of data and the sophistication of analysis tools can tax even the most powerful relational databases. To overcome the limitations of traditional data analysis and storage methods, Intel and SAP worked together to design and build an analytics system that can provide tiered data storage and analytics on large data sets, which can be queried quickly and efficiently. This system exceeded the original proof-of-concept goals of data load speed, availability of the system during data loads, and the speed in which data is available for analysis after being loaded.

[1] Though some of the servers in the Intel and SAP environment were powered by the Intel® Xeon® processor E7 v2 family, Intel and SAP recommend servers powered by the Intel Xeon processor E7 v3 family for new configurations because Intel® Transactional Synchronization Extensions (Intel® TSX) can improve performance by up to 6x. Up to 6x performance improvement for transactional workloads with new Intel TSX claim based on SAP® internal OLTP insert and select tests measuring transactions per minute (TPM) on SUSE* Linux* Enterprise Server 11 SP3.
Configurations:

a. Baseline 1.0: 4S Intel Xeon processor E7-4890 v2, 512 GB memory, SUSE Linux Enterprise Server 11 SP3, SAP HANA® 1 SP8 scoring 14,327 TPM.
b. Up to 1.8x more TPM: 4S Intel Xeon processor E7-4890 v2, 512 GB memory, SUSE Linux Enterprise Server 11 SP3, SAP HANA 1 SP9 scoring 26,139 TPM.
c. Up to 2.7x more TPM: 4S Intel Xeon processor E7-8890 v3, 512 GB memory, SUSE Linux Enterprise Server 11 SP3, SAP HANA 1 SP9—Intel TSX disabled scoring 39,330 TPM.
d. Up to 6x more TPM: 4S Intel Xeon processor E7-8890 v3, 512 GB memory, SUSE Linux Enterprise Server 11 SP3, SAP HANA 1 SP9—Intel TSX enabled scoring 89,619 TPM.

For more complete information visit http://www.intel.com/performance/datacenter.

[2] The hot data tier contained 14 servers with the following CPUs, with each server containing 1 TB of RAM and four sockets: two servers containing the Intel® Xeon® processor E7-8890 v3, four servers containing the Intel Xeon processor E7-8890 v2, five servers containing the Intel Xeon processor E7-8880 v2, two servers containing the Intel Xeon processor E7-4890 v2, and one server containing the Intel Xeon processor E7-4860 v2. Total number of cores in the hot data tier was 852, and total RAM was 14 TB.

[3] The warm data tier contained a single eight-socket server containing the Intel® Xeon® processor E7-8850 v2. The total number of cores was 96, and the total RAM was 2 TB.

[4] Total data loaded was 800 billion rows, which was approximately 51 TB of uncompressed data. The EMC VMAX3* system achieved 800 MB per second (approximately 2 percent of the available capacity) while writing 24 billion rows per hour to the warm data tier. The data load consisted of 100 data files being loaded in parallel using 100 of the 192 cores on the warm data tier, with the remaining cores being used for ad-hoc queries. Each data file was 5 GB in size and contained approximately 77 million rows of data. Note that performance was not measured using industry-standard benchmarks.

[5] For complete EMC VMAX3* benchmark information, visit http://www.storageperformance.org/benchmark_results_files/SPC-2/EMC/B00073_EMC_VMAX-400K/b00073_EMC_VMAX-400K_SPC-2_executive-summary.pdf.

0316/TA/PRW/PDF 333337-001US