

IT@INTEL

Transforming Intel's Supply Chain with Real-Time Analytics

Our integrated data platform provides Intel's supply chain experts with new, innovative ways to analyze data in real time and optimize processes. Our estimated five-year return on investment for our new supply chain platform is USD 208 million.

Craig Chvatal
Principal Engineer, Enterprise Solutions, Intel IT

Anil Varhadkar
Principal Engineer, SC Data Integration, Intel IT

Executive Overview

In today's fast-paced business world, making rapid, data-driven decisions is crucial to maintaining a competitive edge. These decisions are especially important for Intel's supply chain, which spans order taking, resource procurement, manufacturing, testing, and final delivery of products. Intel IT is transforming our legacy supply chain into a modern "glass pipeline" that improves our decision making capabilities and business agility. This glass pipeline is now possible with our new integrated data platform (IDP). Our IDP integrates the SAP HANA* in-memory database with Cloudera Distribution of Hadoop*, running on multiprocessor (four-socket) servers, which are currently based on the Intel® Xeon® processor E7 family.

Reports that used to take hours to generate can now be created in real time as a result of faster access to data. Supply chain planners and experts are more efficient and effective at optimizing business processes, driving operational excellence, and providing higher levels of customer satisfaction. The transformation of Intel's supply chain data is not yet complete. Our estimated five-year return on investment for our new supply chain platform is USD 208 million. The platform will deliver the following capabilities:

- Simplified supply chain and data pipelines
- Self-service analysis, which enables planners to make faster and better decisions
- Improved data quality by eliminating the reconciliation of data across systems
- Real-time analytics that identify, diagnose, and respond to issues

We plan to continue to add more advanced predictive analytics to enhance our "sense-and-respond" supply chain.

Contents

- 1 Executive Overview**
- 2 Business Challenge**
- 3 Solution**
 - Defining a Supply Chain Data Transformation Strategy and Data Model
 - Demonstrating Business Need and Value to Stakeholders
 - Creating a Modern Analytics Reference Architecture
- 6 Results**
 - Ongoing Business Value
 - Lessons Learned
- 8 Next Steps**
- 9 Conclusion**

Contributors

- Nghia Ngo**
Big Data Capability Architect
- Elaine Rainbolt**
Industry Engagement Manager
- Hima Reddy**
Supply Chain Architect, Intel IT
- Darin Watson**
Big Data Platform Architect
- Chandhu Yalla**
Big Data/AI Engineering Manager

Acronyms

- BI** business intelligence
- ECC** enterprise common core
- IDP** integrated data platform
- IoT** Internet of Things

Business Challenge

Intel's supply chain reflects the company's global operations—Intel does business in more than 100 countries, with over 450 supplier factories and 16,000 suppliers. In addition, Intel fulfills over 1 million orders a year from several factories and 30 warehouses. Intel processes over a terabyte of supply chain and manufacturing data every day. As Intel's business grows, evolves, and accelerates, supply chain efficiency and agility is critical to the company's continued success.

Reducing supply chain costs is always a focus for large enterprises. But it is equally important to continually innovate, while simultaneously maintaining supply chain stability. This concept—referred to as a bimodal supply chain¹—requires supply chain IT to serve as both a growth partner and operational caretaker. Intel's supply chain can help disrupt the competitive landscape, but at the same time it must preserve the strong, reliable core that keeps the business running.

Until just a few years ago, Intel's legacy supply chain architecture (see Figure 1) did not support business agility and innovation. The evolving business environment, with increasing cost pressure and new digital business models, made a traditional approach to supply chain management cumbersome. An aging data warehouse and batch-driven processes were characterized by the following disadvantages:

- Multiple data hops
- Data latencies of up to 12 hours
- Data fragmentation
- Data reconciliation and quality issues

¹ Gartner Says Bimodal Balance Is Required for the Future of the Supply Chain. Gartner. gartner.com/newsroom/id/3079617

Legacy Landscape: Data Hops Delay Decisions

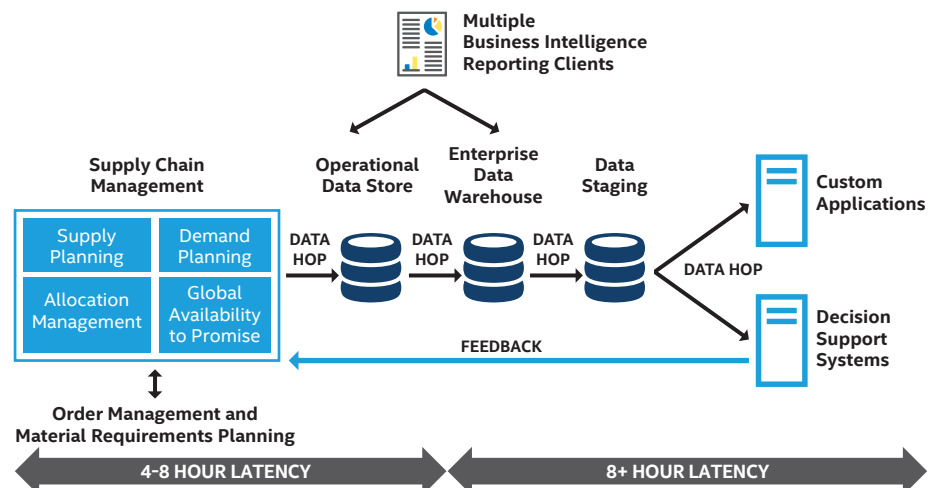


Figure 1. The limitations of our legacy supply chain architecture included latency, data fragmentation, lack of scalability, and excessive data hops.

The multiple data platforms, reporting tools, and business intelligence (BI) solutions—combined with time-consuming manual data stitching and analysis—delayed decision making and Intel's ability to respond to changing markets. It took 18 to 24 hours to derive business insights. In addition, advanced predictive analytics were impossible to implement.

We realized we needed to re-architect Intel's supply chain to provide new and improved decision making capabilities. We decided to deploy an in-memory database and advanced analytics with SAP HANA* technology.

Solution

Intel IT plays a leadership role in creating a supply chain that supports new business models and expands Intel's manufacturing capabilities. We are developing an integrated data platform (IDP) that integrates an SAP HANA system with our Cloudera Distribution of Hadoop* cluster.

SAP HANA offers real-time data management. That means our supply chain managers can analyze data within seconds of it being saved or created—no more having to wait hours for a report or to perform analysis. Such instantaneous access to data opens up endless possibilities. The IDP enables supply chain experts at Intel to ask “What can I do now that I could not do before?” and to explore the “art of the possible.”

SAP HANA supports data for all primary supply chain capabilities such as core manufacturing, order management, supply demand, supply planning, financial data, analytic systems, dynamic inventory tracking, warehousing, supplier security management, and capacity utilization functions. We use Hadoop to store and process many structured and unstructured types of data, such as input from Internet of Things (IoT) sensors and social media streams. The combination of SAP HANA's in-memory capabilities and Hadoop's big data capabilities creates a real-time “sense-and-respond” supply chain.

The IDP enables supply chain experts at Intel to ask “What can I do now that I could not do before?” and to explore the “art of the possible.”

Defining a Supply Chain Data Transformation Strategy

Our supply chain transformation strategy consists of three tenets:

- **End-to-end visibility.** We are transforming Intel's legacy batch and transactional supply chain system into a real-time platform. This “glass pipeline” with end-to-end data visibility supports real-time analytics, real-time data, and real-time processing.
- **Responsiveness.** Intel's supply chain managers need to explore supply chain data to identify risk areas, pinpoint root causes of issues, and perform what-if analysis to evaluate alternative solutions. We are implementing advanced BI algorithms to enable deep analysis and correlation of data from disparate sources, both structured and unstructured.
- **Simplification.** We are re-architecting our data foundation to coalesce all data elements and key supply chain metrics into a single version of truth. Previously, data was distributed across six data warehouses. Reducing the number of data hops and data latency improves data reliability and provides better supply chain visibility and operations.

This strategy provides the foundation for business growth (see Figure 2) through end-to-end supply chain visibility, faster decision making, and agile supply chain analyses. These capabilities are producing significant business value through maximized margins, faster product fulfillment, reduced operating costs, improved operational excellence, and reduced days of inventory.

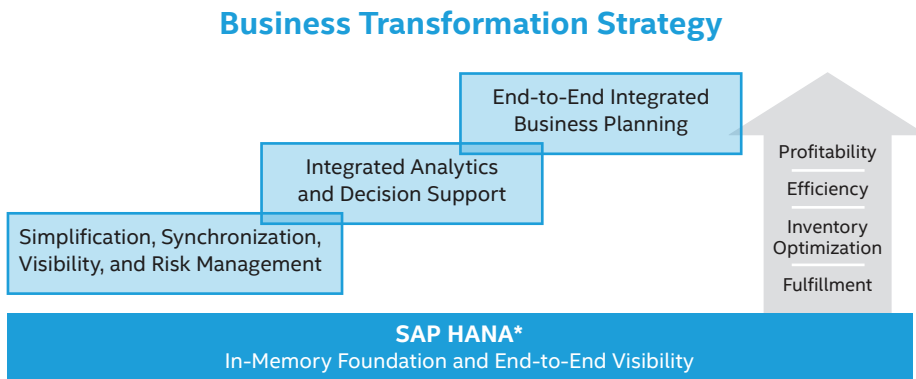


Figure 2. Our supply chain transformation strategy, based on an integrated data platform, lays the foundation for business growth and increased efficiency.

Demonstrating Business Need and Value to Stakeholders

Before starting any major IT initiative, we demonstrate the business need and value to upper management and business stakeholders. This approach helps drive adoption of new processes and technology across the business.

We communicated the negative business impacts of the legacy supply chain data architecture to management and discussed the analytics maturity curve. We needed to move beyond the ability to answer rudimentary questions such as "How much of X do we have?" and "Where are we at risk?" to performing root-cause analysis and what-if simulation to optimize the supply chain.

We also highlighted that initial investments were foundational and would provide more business value later. While the in-memory capabilities of SAP HANA and its hardware optimizations make existing processes faster, the full potential of the IDP lies in its ability to discover ways to improve the processes themselves. For example, we explored the following questions:

- Which supply chain processes are designed to accommodate performance or data delays?
- What tasks or processes are disabled or forbidden, and why?
- Where are workaround processes necessary?
- In which scenarios do we need to train employees to resolve supply chain performance challenges?
- Where would it be useful to drill down into transaction data?
- Where do we need to simplify the data model?

The business value drivers for the IDP are not limited to IT, supply chain managers, and the finance department. Intel's business units also benefit from the enhanced agility to meet customer demands, as shown in Figure 3.

Intel Ranks #6 in Gartner Supply Chain Top 25

In 2017, Intel was ranked #6 in Gartner's Supply Chain Top 25.¹ Two key aspects of the Supply Chain Top 25 ranking are the demonstration of demand-driven leadership and corporate social responsibility. Gartner noted three trends and innovative practices across the top global supply chains:

- Build digital supply chain capabilities, internally and with partners, in areas that best meet the strategic business objectives of the company and its customers.
- Apply agile development techniques and a modular supply chain service approach to more quickly respond to changing supply and market conditions.
- Invest in programs that ensure the ongoing health of the people, environmental, and partner ecosystems the company relies on for competitive advantage.

¹ The Gartner Supply Chain Top 25 for 2017. Gartner. gartner.com/doc/3728317

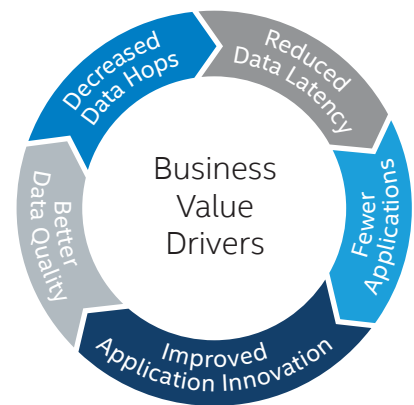


Figure 3. Business drivers for implementing an integrated data platform are compelling for Intel's business units.

Creating a Modern Analytics Reference Architecture

Figure 4 illustrates an IDP reference architecture, which includes the SAP HANA platform, Cloudera Distribution of Hadoop, and SAP Vora*. Self-service advanced analytics and predictive models are layered on top. The entire IDP is fueled by servers based on the Intel® Xeon® Scalable processor. The primary functions of the components are:

- **Intelligence tools and predictive models create proactive and predictive supply chain models.** Advanced statistical regressions and analytics algorithms predict supply chain conditions using R* and SAP HANA's built-in Predictive Analytics Library (PAL). Data visualization tools enable analysts to quickly understand complex data sets for quick root-cause analysis and faster response to changing conditions.
- **SAP HANA** is a highly scalable in-memory database that ingests and manages supply chain data in real time to support analytical decision making.
- **Cloudera Distribution of Hadoop** integrates and utilizes both internal and external data, including unstructured content from the supply chain, IoT sensor data, and market intelligence.
- **SAP System Landscape Transformation* (SLT*)**, SAP Data Services*, integrate and manage real-time data from across the enterprise into SAP HANA and Cloudera Distribution of Hadoop.
- **SAP Vora** integrates Hadoop data with the SAP HANA platform. This application incorporates a tremendous amount of new data types and information, and integrates it with “traditional” enterprise data in HANA.
- **Servers based on the Intel Xeon Scalable processor.** These servers provide the computing performance, security, and agility necessary for real-time supply chain management.
- **Intel® Optane™ memory** helps eliminate data center storage bottlenecks and allows for bigger data sets. It can accelerate applications and reduce transaction costs for latency-sensitive workloads.

Modern Analytics Reference Architecture

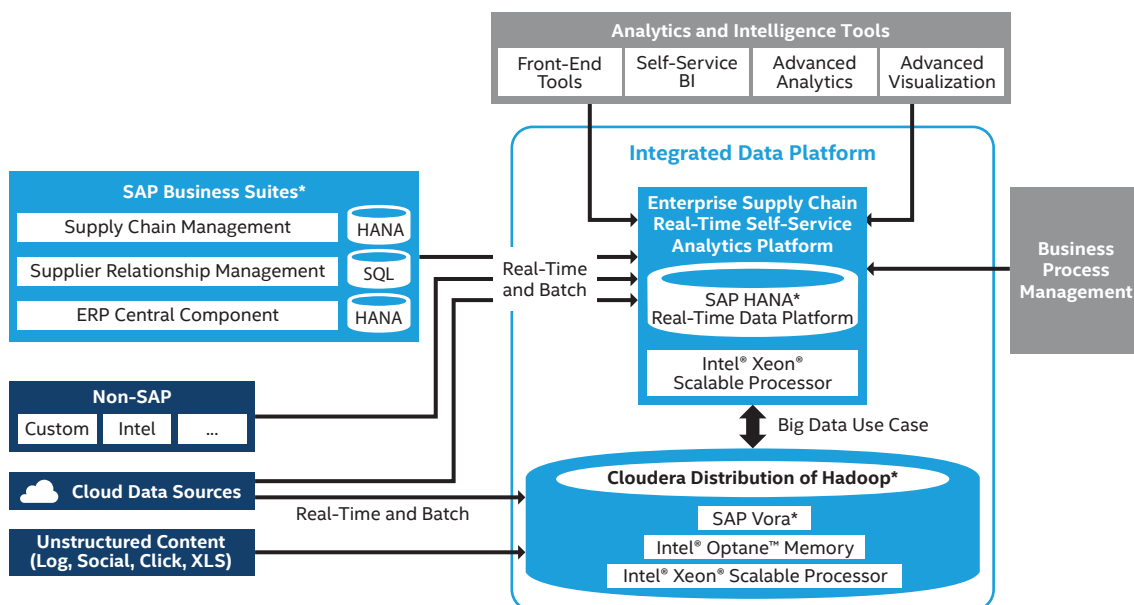


Figure 4. This modern analytics reference architecture connects SAP HANA* and Cloudera Distribution of Hadoop*, fueled by servers based on the Intel® Xeon® Scalable processor and Intel® Optane™ memory, for faster sense-and-respond capabilities.

An IDP can introduce potential risk from single point of exposure for critical supply chain data. An enterprise needs well-defined rules for storing, transporting, and accessing sensitive data. The IDP encrypts data while it is being transmitted through the network and when at rest in the SAP HANA database and Hadoop platform. In addition to standard authentication and authorization for data access, an IDP uses row-level security to help prevent users from accessing data they are not entitled to access.

Results

While we are only in the third year of our five-year SAP HANA supply chain data transformation, our IDP is already providing measurable business benefits (see Figure 5).

Ongoing Business Value

After our initial migration to an IDP, we were able to reduce the size of our enterprise resource planning database by 63 percent and the enterprise common core (ECC) database by 80 percent. Database transactions decreased by between 25 and 50 percent, and ECC processing time decreased by 50 percent. We also experienced a 75 percent efficiency gain in account reconciliation. Because our new platform is far less complex than the legacy architecture, we anticipate a 45 percent reduction of IT support staff, freeing technicians to work on higher-value tasks.

The full potential of our IDP is the ability to transform processes that take advantage of the new capabilities. Integrating analytics with business decision-making data improves planner efficiency and accuracy. Supply planners can access a worldwide view of supply chain hubs, prioritized by financial risk and business value, such as stock-up conditions or missed shipments. They can drill down on risk areas and root causes. Here are just a few of the transformational capabilities:

- Detailed real-time, historical, and future projections including inventory, supply, and demand
- Weeks of inventory-control charts that highlight outliers
- Summary of alerts with drill-down to product structure
- Demand signal evaluation and root-cause analysis
- Instant access to previous-period supply plans
- Ability to rebalance recommendation options to resolve shortages
- Email collaboration to review and ratify decisions

An intuitive interface (see [Figure 6](#)) makes these capabilities available to users through a self-service portal.

Putting the Integrated Data Platform (IDP) to Work

We identified two early use cases for our IDP:

- **Dynamic routing.** Real-time data about the gross or dimensional weight of boxes being shipped to customers enables quick and intelligent selection of shipping options such as shipping company and service method.
- **Parts management.** Intel's supply chain consists of parts that are necessary for factory equipment repairs. The IDP provides a real-time, highly accurate picture of what parts are in stock. It also enables factory operators to order parts at the optimal time, which reduces the costs of storing and then scrapping unused parts.

Opportunity: Speed Matters

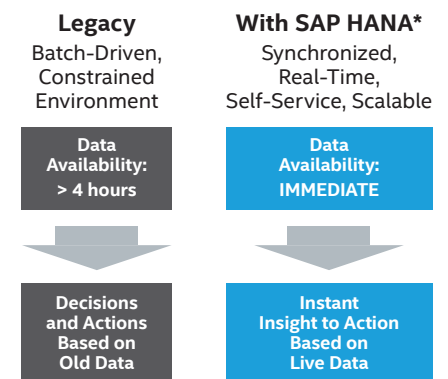


Figure 5. The integrated data platform we are building enables our supply chain experts to act on real-time data, which in turns leads to better and faster business decisions.

Example of a Self-Service Portal

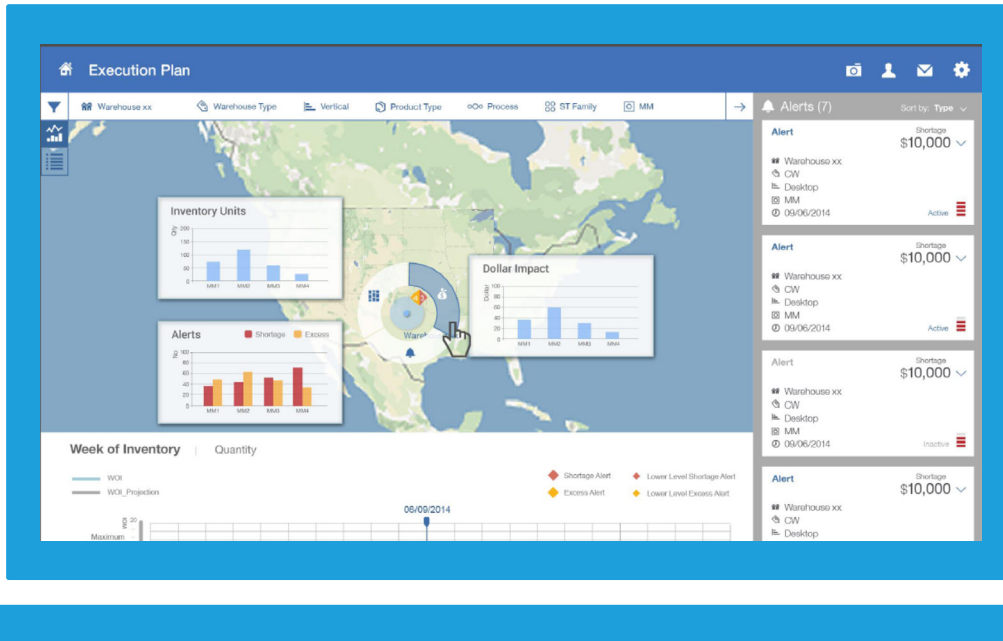


Figure 6. The easy-to-use, self-service interface enables Intel's planners to quickly drill down into data, make informed decisions, and optimize our supply chain.

Lessons Learned

Because we are still in the process of onboarding supply chain activities onto the IDP, we can learn as we go. Some of the important lessons we have learned so far include the following:

- To take full advantage of real-time data, business processes must change.
- Knowledge of business processes, data, and insight into data consumption is essential. Start with a good integrated data model.
- Start small and deliver business value incrementally.
- Although the necessary technical and database skills are similar to traditional database skills, IT staff must be open to new ways of doing things. Porting existing data queries and data models “as-is”—while they will work—will typically not result in 100x improvement. Like any new data system, if the model and queries are reworked around the strengths of the SAP HANA IDP, one can achieve amazing results.
- Real-time data flowing through the supply chain is a game changer, but it takes significant effort to optimize and align data flowing into the IDP at different frequencies and granularity.
- Data governance is key to success. We created structure, accountability, and prescriptive guidance for the team.

Porting existing data queries and data models “as-is” will typically not result in 100x improvement.

Customers Benefit from Intel and SAP Collaboration

Intel and SAP have been collaborating on the development of SAP HANA* for more than a decade. Intel began defining our system architecture with the four-socket Intel® Xeon® processor E7 family at the same time SAP started developing the SAP HANA platform. Engineering collaboration and innovation with SAP continues today with the Intel® Xeon® Scalable processors. The result is an extremely efficient analytics platform that reduces complexity and increases performance for mission-critical business processes.

Intel architects work closely with their SAP HANA counterparts to optimize SAP HANA source code to take full advantage Intel Xeon Scalable processor features such as Intel® Transactional Synchronization Extensions, Intel® Advanced Vector Extensions 512, and Intel® Run Sure Technology. The new Intel® Xeon® processor Scalable family demonstrates a significant jump in performance for SAP HANA database workloads in testing, with performance gains of up to 1.59 times faster compared to the previous-generation Intel Xeon processor family.¹ Additionally, the SAP HANA platform now is certified to support up to six times more memory with the Intel Xeon processor Scalable family than with the typical four-year-old systems that are in broad use in data centers today.²

When building out SAP HANA 2 configurations, customers can now select the Intel® Xeon® Platinum processor for industry-leading mission-critical performance and gain the largest SAP HANA 2 memory footprint supported out of the box. For example, expanded SAP HANA 2 memory configurations for eight-socket systems now support 6 TB of memory for online analytical processing and 12 TB of memory for online transaction processing. SAP HANA 2 Tailored Data Center Integration certification can provide further system configuration flexibility.³ Together, Intel and SAP are simplifying IT landscapes, redefining enterprise applications, and automating advanced analytics.

¹ Up to 1.59x higher online transaction processing (OLTP) performance versus Intel® Xeon® processor E7 v4 family performance estimate based on SAP HANA* internal S-OLTP workload (internal testing). Baseline configuration: one-node, four-socket Intel® Xeon® processor E7-8890 v4 with 1,024 GB total memory on SUSE Linux Enterprise Server* (SLES*) 12 SP1 versus estimates based on SAP internal testing on systems based on the one-node, four-socket Intel® Xeon® Scalable processor.

² The SAP HANA* database platform is certified to support up to 6x memory for online analytical processing configurations on the Intel® Xeon® processor Scalable family when compared to systems available in the market four years ago, which represent the typical data center installed base. For example, systems based on the four-socket Intel® Xeon® processor E7.

³ SAP HANA* deployment options: sap.com/products/hana/implementation/deployment

Next Steps

The digital transformation of Intel's supply chain is a multiyear journey. We plan to add more capabilities to streamline Intel's supply chain processes, as well as improve the data tiering and data virtualization between SAP HANA and Cloudera Distribution of Hadoop. This process will unlock our ability to deliver new analytical functionality while optimizing the IDP's total cost of ownership.

Other improvements we anticipate implementing include the following:

- SAP Business Suite 4 SAP HANA* (S/4HANA*), the next generation of SAP business software, will enable further consolidation of our supply chain landscape while providing advanced analytical capabilities.
- IoT-based analytics at the edge, in retail, and in the cloud and data center will create a truly connected supply chain.
- Machine learning algorithms will add supply chain intelligence throughout the enterprise.
- Servers based on the Intel Xeon Scalable processor for more compute performance, more memory, and other processor-level benefits.

- Cutting-edge nonvolatile memory to increase the performance and reliability of the IDP while decreasing total cost of ownership.
- Intel Optane memory, which helps eliminate data center storage bottlenecks and allows for bigger data sets.

Conclusion

The ongoing digital transformation of Intel's supply chain with our IDP is increasing Intel's efficiency, organizational agility, and competitive advantage in today's complex global market. Our IDP, built with SAP HANA's in-memory database technology and Cloudera Distribution of Hadoop, and fueled by multiprocessor (four-socket) servers, provides advanced real-time analytics. The end-to-end supply chain visibility enhances planning capabilities and enables faster, more agile business decisions.

Our estimated five-year return on investment for our new supply chain platform is USD 208 million. We are already seeing the benefits of better quality of data, better decisions, and an analytical foundation for a sense-and-respond supply chain. We will continue our transformation of Intel's supply chain data over the next two years. We will transition to servers based on the Intel Xeon Scalable processor, and add Intel Optane memory and persistent memory technology from Intel. These additions will increase performance and reliability of our IDP, supporting new business processes and insights.

For more information on Intel IT best practices, visit intel.com/IT.

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues, and we want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation:

- [Twitter](#)
- [#IntelIT](#)
- [LinkedIn](#)
- [IT Center Community](#)

Visit us today at intel.com/IT or contact your local Intel representative if you would like to learn more.

Related Content

If you liked this paper, you may also be interested in these related stories:

- [Optimizing Intel's Supply Chain with an In-Memory Data Platform paper](#)



Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes.

Software and workloads used in performance tests may have been optimized for performance only on Intel® microprocessors. Performance tests, such as SYSmark® and MobileMark®, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.

Configurations: System configurations and performance tests conducted are discussed in detail within the body of this paper. For more information go to intel.com/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. Check with your system manufacturer or retailer or learn more at intel.com.

THE INFORMATION PROVIDED IN THIS PAPER IS INTENDED TO BE GENERAL IN NATURE AND IS NOT SPECIFIC GUIDANCE. RECOMMENDATIONS (INCLUDING POTENTIAL COST SAVINGS) ARE BASED UPON INTEL'S EXPERIENCE AND ARE ESTIMATES ONLY. INTEL DOES NOT GUARANTEE OR WARRANT OTHERS WILL OBTAIN SIMILAR RESULTS.

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS AND SERVICES. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS AND SERVICES INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel, the Intel logo, Intel Optane, and Xeon are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others. Copyright Intel Corporation. All rights reserved.

Printed in USA

Please Recycle

0917/ERA/IC/PDF