



SIEMENS

Omneo Big Data Analytics for Product Performance Intelligence Helps Manufacturers Save Millions

Overview

Today's consumers have high expectations for the products we use on a daily basis, particularly when it comes to electronic devices. We want new products to come out faster, at lower prices, and with more capabilities than previous versions. But we also demand increased reliability with these products. Siemens PLM Software, a veteran in the enterprise manufacturing and supply chain space, saw this trend and identified an opportunity. "Why is it," Bill Boswell, Senior Director Marketing and Business Strategy, Cloud Services, Siemens PLM Software, asked, "that in spite of all of the investment in design, quality assurance, supplier and manufacturing controls and monitoring, device manufacturers are still having seemingly simple issues? Why do cell phone batteries still overheat and melt down?"

"Supply chain quality has been the holy grail of manufacturing that nobody has attempted to solve yet," added Dave Mitchell, CTO, Cloud Services, Siemens PLM Software.

Siemens PLM Software is addressing this with its supply chain cloud solution, Omneo™ services, which runs on an [enterprise data hub \(EDH\)](#) from [Cloudera](#). The Omneo EDH provides electronic device manufacturers with a holistic, comprehensive, and interactive solution that helps them resolve supply chain issues before they impact the customer experience, or worse yet, the bottom line. These capabilities are helping Siemens PLM Software win business with global brands, including two consumer electronics companies that are known for their leadership in supply chain operations and product quality. Based on conservative numbers from clients, customers have reported that the Siemens PLM Software supply chain cloud solution, Omneo, has driven total annual operational savings between USD \$15 and \$25 million.

The Challenge

Electronic device manufacturers are responsible for delivering millions of products, each comprised of hundreds of components. "You have hundreds of millions of components being sourced globally, put together, and pushed through distribution channels to customers. So much can go wrong. Even if the components are designed perfectly, unexpected compatibility issues may surface downstream, when they are used in conjunction with other parts; something that is nearly impossible to predict in advance. Human beings can't possibly monitor and conceive of all the possibility for error. You need to monitor product quality across the entire value chain," explained Bill Boswell.

Dave Mitchell added, "The manufacturing sector has grown not just vertically, but also horizontally. Single site manufacturing, which in and of itself generates tons of data, has become very complicated because of the growth of the global supply chain. So, when we talk about big data in terms of velocity, volume, veracity, variety, they are excessively exacerbated in a global supply chain solution."

Key Highlights

Industry

- Manufacturing
- Supply Chain

Locations

Charlotte, NC, USA

Business Application Supported

- Enterprise data hub empowering 360-degree view of product performance across the supply chain

Impact

- Annual savings among client base conservatively estimated between USD \$15-25 million based on ability to identify and address supply chain issues in near real time
- Interactive analytics enabled, vs. waiting days for data
- Flexibility to support various workloads and users

Technologies in Use

- Hadoop Platform: Cloudera Enterprise
- Hadoop Components: Cloudera Manager, Cloudera Search, HBase, Hue, HDFS, MapReduce, Oozie, Pig
- Servers: Dell PowerEdge C8000

Big Data Scale

- 48 production CDH nodes and growing
- Millions of data points ingested daily from dozens of data sources

That's what Siemens PLM Software set out to do: build a global monitoring solution using information and data driven systems. In most supply chain processes, the data exists and is being captured to some extent, but is very disparate—spread across service centers, suppliers, manufacturing sites, repair centers, and retail outlets. "It's all in different formats, different locations, different types," explained Bill Boswell, "and there's no way to look at it cohesively. It's distributed, and on top of that, the data quality is not consistent. Fields are missing. Information is not always the same. It's not always where you want it, when you want it. By the time people get to it and bring it together, it's often too late for the kind of decisions that we need to make.

On its mission to address the problem, Siemens PLM Software quickly realized that its traditional relational database management system (RDBMS) wouldn't suffice. It used a highly normalized structure, and the data volumes from manufacturing clients were already starting to exceed the system's limitations. Data older than six months had to be archived; if a client wanted to do root cause analysis on something that happened a year ago, it couldn't. As Kathleen deValk, chief architect of Omneo explained, "It was not designed to grow to that scale, and it was not very cost effective to maintain. The licensing costs got extreme, and we were starting to miss our performance targets."



The team evaluated several big data technologies with a few key criteria in mind:

- Scalability to grow with customers' needs over time
- Flexibility to accommodate the needs of all customers and their diverse data sets in a multi-tenant environment
- Low total cost of ownership (TCO) to offer an efficient big data solution that clients would be able to afford

An IBM Netezza data warehouse appliance was evaluated, and delivered strong performance but was relatively expensive and the scale-out path wasn't attractive. "Having to decommission one server to buy a bigger, better server when we needed to grow wasn't really the way we wanted to go," explained deValk. Columnar stores, such as Infobright, were also evaluated, which demonstrated strong analytical capabilities but weren't flexible enough to accommodate clients' diverse needs universally. "We also looked at NoSQL technologies like Cassandra and MongoDB, but we really were very attracted to [\[Apache\] Hadoop](#) in general. It was an up and coming technology, and a lot of people were adopting it. We really liked the promise of linear scalability in a cost effective way—we don't have to have the biggest server in the world. We don't need this giant big data appliance. We can just put a couple nodes together, and grow it over time as the needs of our customers evolve."

Siemens PLM Software's Omneo team also saw value in the open source aspect of Hadoop, promising flexibility and rapid innovation at low cost. Without any upfront investment, they could download the software and try it out. Nobody on the team had experience with Hadoop, but [Cloudera Express](#) presented a packaged solution to help them get started, bundling Hadoop and related open source projects with [Cloudera Manager](#) for end-to-end management of the cluster. "That was the primary reason we went with Cloudera at first—just the ease of use. I deployed a virtual machine right on my laptop and had it up and running in a day. It was really easy to download, really easy to get up and running," said deValk.

After a few months of testing out Cloudera Express with promising results, Siemens PLM Software did an assessment to compare Cloudera's TCO to that of IBM Netezza and the legacy RDBMS. Cloudera came out on top, at about 75% lower cost per terabyte (TB) than IBM Netezza and 90% lower cost per TB than the incumbent. "It looked like Cloudera was the way to go," said deValk. "It was going to be cost-effective, it was going to hit our requirements and performance targets, and it could handle the large volumes of data."



TCO Evaluation of Big Data Technologies

The Omneo team also spent time comparing a Cloudera Enterprise subscription with other commercial Hadoop offerings. The company decided to move forward with Cloudera over other Hadoop vendors for four key reasons:

- Long-term company strategy and viability:** “We want to partner with a company that is innovative and that will be around for the long haul,” explained Bill Boswell. “We need a vendor that understands it’s not just about Hadoop—it’s about adding value to the Hadoop ecosystem. And Cloudera does that very well.” Dave Mitchell added, “Of all of the vendors we evaluated, Cloudera was the only one that gave a complete, end-to-end solution that aligned with our big data strategy, and shared our vision for creating a single, unified platform that can serve simple search, analytics, or data mining.”
- Ease of use and maturity of Cloudera Manager:** The team knew and liked their CDH cluster; it was working smoothly and they appreciated the functionality of Cloudera Manager. “I can log in and see everything that’s going on in my cluster all at the same time,” explained deValk. “We have an alerting system that monitors if a node is in bad health, and will send me an email or can send a text message to my systems administrator. We always know the health of our cluster and its nodes. We really can stay in touch with what’s happening on the system, and we can deploy and manage things really easily.”
- Enterprise-grade support:** Siemens PLM Software sought a partner that could support a variety of demands from its multinational customers with speed, agility, and deep domain expertise.
- Dedication to open source:** The Omneo team had an affinity toward an open source solution that would ensure rapid platform innovation and flexibility over the long term.

Solution

Siemens PLM Software has deployed a multi-tenant enterprise data hub from Cloudera as the platform behind its new Omneo solution, running on [Dell](#) hardware that leverages the jointly developed [Dell Cloudera reference architecture](#). Its clients—consumer electronics brand owners—have indirect access to the EDH via Omneo’s web portal offering a highly configurable user interface. Omneo ingests machine-generated and existing system data from the manufacturing process, including clients’ own factory data in addition to data

“The enterprise data hub delivers this raw storage format which is fault tolerant, redundant, and reliable. The different technologies that make the Cloudera offering so appealing is that the MapReduce, Search, and analytic storage frameworks make it possible for us to build our industry-specific solutions on top of it.”

Kathleen deValk, Chief Architect, Siemens PLM Software

“We really liked the promise of linear scalability in a cost effective way... We can just put a couple nodes together, and grow it over time as the needs of our customers evolve”

Kathleen deValk, Chief Architect, Siemens PLM Software

“We want to partner with a company that is innovative and will be around for the long haul. We need a vendor that understands it's not just about Hadoop—it's about adding value to the Hadoop ecosystem. And Cloudera does that very well.”

Bill Boswell, Senior Director Marketing and Business Strategy, Cloud Services, Siemens PLM Software

from their suppliers, field services, Internet of Things (IoT), and after-market repair and re-manufacturing data. All of that data in its raw form is loaded into Cloudera, and from there it is transformed, contextualized, stored, and analyzed using various tools to support different workloads.

- **MapReduce** is used to transform and manipulate data into any structure needed.
- **HBase** provides an easy way to access specific records in real time.
- **Cloudera Search**, which natively integrates Apache Solr with CDH, enables Omneo to rapidly index all of its raw data in a way that makes sense for customers. “One of our customers has about 9 billion records in the search engine, and we can search all that in seconds,” noted deValk.

Siemens PLM Software relies on Cloudera Manager to keep its EDH running smoothly, and values its ability to manage and monitor not only MapReduce and HDFS, but also HBase and Solr, through a single, simple interface.

Impact: Better Visibility, Better Products

The Omneo big-data analytics for product performance intelligence solution brings together data from end to end in the supply chain—from suppliers, manufacturing, equipment, field service, Internet of Things (IoT) and repair/re-manufacturing operations—delivering a 360-degree view of the supply chain process in seconds.

“Our customers are drastically reducing the time they spend finding and resolving product issues, and using that time to ensure higher customer satisfaction.” – Dave Mitchell, “It’s instantaneous gratification to non-technical resources to help them get their jobs done better.”

The Omneo solution allows manufacturers to access their data in different ways.

The enterprise data hub delivers this raw storage format which is fault tolerant, redundant, and reliable. The different technologies that make the Cloudera offering so appealing is that the MapReduce, Search, and analytic storage frameworks make it possible for us to build our industry-specific solutions on top of it.

Over time, as we build up our platform, we expect to leverage most of these technologies. You could consider Omneo the supply chain enterprise data hub.”

Bill Boswell explained, “If something happens at any supplier that drives a sudden increase in quality issues, we can figure out where the issue is stemming from and why. We can answer those questions for the supply chain and manufacturing community in minutes or worst case, hours, versus investigations that would normally take weeks and months.”

Impact: \$15 to \$25 Million in Annual Operational Savings

The increased business visibility provided by the Omneo solution means that problems can be identified and addressed as soon as they start occurring, and manufacturers can spend more time resolving challenges and less time trying to pinpoint them, which directly impacts the bottom line. Siemens PLM Software clients report total annual savings between USD \$15 and \$25 million each conservatively, due to insights from its EDH.

In one case, the Omneo solution helped a global brand recover millions of dollars. This company had an issue managing component consumption at one of their contract manufacturers, who'd been instructed to use certain allocations of alternative suppliers for particular parts (for example, sound cards). Compliance with the set percentage allocation between suppliers is very important to the manufacturer because it gets rebates based on economies of scale with given suppliers. Without the Omneo software, manufacturers couldn't get a real-time

cloudera.com

1-888-789-1488 or 1-650-362-0488
Cloudera, Inc. 1001 Page Mill Road, Palo Alto, CA 94304, USA

© 2015 Cloudera, Inc. All rights reserved. Cloudera and the Cloudera logo are trademarks or registered trademarks of Cloudera Inc. in the USA and other countries. All other trademarks are the property of their respective companies. Information is subject to change without notice

cloudera-casestudy-omneo-103



view into the parts being used in current manufacturing cycles—making it nearly impossible to ensure compliance. The quarter would end and, often times, rebates would be missed. Now, with Omneo, the global brand can monitor operations on a daily basis, to ensure suppliers and contract manufacturers are complying and their rebates are earned. This amounted to millions of dollars in cost recovery for the brand.

Another customer of Siemens PLM Software is focused on providing a superior customer experience with its products. Utilizing Omneo, the company has a global view into every geography across the 40 countries it sells into, and it can see—within hours—when there is a service-related issue. The manufacturer can jump on the problem, remedy it, and proactively address customers' concerns without waiting for the customers to contact the manufacturer or merchant with a complaint. This results in a better customer experience and reduced support costs.

See www.omneo.com for more information on the Omneo solution

About Cloudera

Cloudera delivers the modern platform for data management and analytics. The world's leading organizations trust Cloudera to help solve their most challenging business problems with Cloudera Enterprise, the fastest, easiest, and most secure data platform built on Apache Hadoop.