

Customer Insight


LEADING THE HIGHLY COMPETITIVE MEMORY & STORAGE INDUSTRY WITH DATA ANALYTICS AND INSIGHT

////////////////////////////////////

Impact

- Shortened the time it took to identify misprocessed die from 7 days to under an hour
- Enabled IT to focus more on data analytics and less on data management
- Eliminated latency for Hadoop users performing data analytics

////////////////////////////////////



Time to identify misprocessed die

Micron is a world leader in innovative memory solutions that transform how the world uses information. For more than 40 years, the company has been instrumental in the world’s most significant technology advancements, delivering optimal memory and storage systems for a broad range of applications. Micron manufactures the industry’s broadest portfolio of flash memory and storage technologies. Many of these begin with a wafer of silicon which is cut into a die, a thin layer of semiconducting material, in one of Micron’s global fabrication (“fab”) sites and then assembled into flash memory products.

Greg Kincade, Senior Ecosystem Enablement Program Manager for the Micron Storage Solutions Center, explains the company’s capabilities further, “We build NAND, DRAM, and SSDs that are optimized for workloads and applications. What sets us apart is our ability to fold in insights and improvements from our developments to our product cycles and to our internal IT and manufacturing operations.”

“The highly competitive memory manufacturing industry is a \$162 billion business, growing at an annual rate of 25 percent,” adds Mike Cunliffe Senior Data Architect for Enterprise Data and Analytics, Manufacturing, at Micron. “Micron is one of the top memory manufacturers worldwide, and the only one of the top three based in the U.S.”

Kincade and Cunliffe both play a role in keeping Micron competitive. Kincade and his team test and validate hypotheses for flash storage improvements using data analytics, focused mostly on popular workloads and use cases. Cunliffe supports data scientists and engineers who are developing analytic solutions to provide Micron with a competitive advantage. He also supports the broader manufacturing process.

“Competitive advantage in this industry is measured in fractions of a percent, so we spend a lot of time working on yield, cycle time improvements, and reducing the time required to ramp up products and get them to market. It’s through the sum of all these fractional performance gains that we are able to retain and advance our competitiveness,” Cunliffe says.

Gaining Competitive Advantage through Data

At Micron, insight from data fuels this competitive advantage. For example, Kincade’s project management work at the Storage Solution Center uses data generated from a suite of tests to uncover how and where SSDs and memory may solve a specific data analytics, processing, or storage problem. The data is then used to document and demonstrate the proof points for the user community.

Data also fuels the highly complex process of semiconductor manufacturing, in which Micron has more than 40,000 patents. Micron leverages advanced analytics to identify patterns and relationships among inputs and discrete manufacturing process steps, which helps to optimize the factors that can have the greatest effect on yield. The company manages this manufacturing complexity with automation. “The only way to manage automation is with data, so data is really key to moving product in our manufacturing line. In fact, without data, our manufacturing process stops,” says Cunliffe.

“Hadoop and big data are our technology focus. We chose Cloudera because of its commitment to open source, its collaborative approach to development, and because it is a leader in the Hadoop space.”

Greg Kincade, Senior Ecosystem Enablement Program Manager for the Micron Storage Solutions Center

About Cloudera

At Cloudera, we believe that data can make what is impossible today, possible tomorrow. We empower people to transform complex data into clear and actionable insights. Cloudera delivers an enterprise data cloud for any data, anywhere, from the Edge to AI. Powered by the relentless innovation of the open source community, Cloudera advances digital transformation for the world's largest enterprises.

Learn more at cloudera.com

Enabling Analytics and Innovation

Micron has been deep into big data for decades. But, as Hadoop open source solutions for big data grew, the company needed a data platform that would allow it to focus more on analytics and less on data management. In addition, the company wanted to address the latency issues its Hadoop users were experiencing when using real-time data analytics. Noting that Cloudera could provide the modern data architecture the company was looking for, Micron chose Cloudera's Data Platform 3.0 to enable Micron to focus on analytics, to address latency, and to stay competitive.

Speeding Data Insight in Manufacturing

Micron's initial use of Cloudera's data platform was for its functional test data, a very large, complex dataset. This data checks all the boxes of what Cunliffe calls “the three Vs.” The volume of this data is in excess of a petabyte per year, the variety with over 10,000 discrete measurements per chip and in some cases over a million measurements per chip, and velocity of data capture, as the company collects more than a trillion measurements per day over millions of integrated circuits or die.

Micron enjoyed many successes with this initial functional test, but one in particular stands out to Cunliffe—a semiconductor manufacturing excursion (the term for a deviation in product or process) where some die got out of spec and needed to be identified and contained, and the excursion cause fixed. The QA group, which was using the legacy systems and trolling through the production data, believed they could identify the die in approximately seven days. Knowing that was not fast enough, they turned to Cloudera and queried the data using Hive.

“We were able to identify all the die that were affected in under an hour,” Cunliffe says. “It was a huge success for Micron and our customers – it was also a huge success for Cloudera.” Since that initial functional test dataset, the team has added more than 100 other datasets into the environment, as well as more than 100 data scientists and data engineers to develop analytical solutions. Those solutions have brought in more than 1000 users that are currently using the environment. Micron uses Apache NiFi for data ingestion pipelines, to provide a secure global of manufacturing processes.

“On any given day, there are over 500,000 jobs that run through our environment and the analytics that they produce are critical to Micron's business success,” says Cunliffe. “Most of the analytics that we run are around manufacturing to drive yield, cycle time, and new product ramps.”

Improving Performance for Hive Acceleration

Micron's first investigation of the Cloudera's data platform at the Storage Solutions Center was to test Hadoop performance with all hard disk drive (HDD) storage vs hybrid HDD and SSD. The first test scenario combined a single SSD per node with Cloudera's data platform and a Hive acceleration environment. That test showed a two and a half times performance improvement over a 10K RPM HDD cluster—speeding up results of big data analysis with Hadoop.

“Our challenge was to use our Micron IP to see if we could eliminate the latency for our Hadoop users, and we were able to do that with our SSDs and Cloudera's data platform,” says Kincade. “The platform includes key innovations, ultimately leading to faster insights for our business. We're just now beginning our journey with Cloudera, testing how well it performs and scales with our SSDs and our memory.”

Driving Success with Cloudera's Support

As its data environment has grown and evolved, Micron has appreciated the support Cloudera has provided. “Cloudera's support in engineering has been key to our success with the platform and has really helped propel us to provide excellent solutions for our customers,” says Cunliffe.