Modernizing Semiconductor Manufacturing Analytics Platform using AWS Data Analytics Services

In the high-tech semiconductor industry, data is paramount to achieving operational excellence, product quality, and competitive advantage. The sheer volume and complexity of data generated at various stages of the semiconductor manufacturing process necessitate advanced analytics to derive actionable insights.

This blog explores prevalent manufacturing analytics use cases in the semiconductor industry, underscoring their role in optimizing processes, enhancing yield, improving quality, and fostering innovation. Additionally, it provides insights into how to establish a robust manufacturing analytics data lake utilizing AWS services, enabling organizations to harness the power of data-driven decision-making and gain a competitive edge in their respective industries.

Manufacturing Analytics Use Cases

Yield optimization, a crucial process in semiconductor manufacturing, involves the comprehensive analysis of data from various stages of the manufacturing process to maximize the number of functional chips produced from a wafer. This data-driven approach encompasses sort, inline, and defect data, providing valuable insights into yield performance. Sort data, derived from electrical and functional tests on individual dies, offers detailed information on the yield of each die. Inline data, collected real-time during wafer processing, enables monitoring of process conditions and helps identify potential issues early on. Defect data, obtained through inspection stages, identifies the root causes of yield loss, enabling manufacturers to take targeted actions to improve process control and reduce defects.

Analytics techniques such as root cause analysis, statistical process control (SPC), and predictive modeling play a pivotal role in yield optimization. These techniques help manufacturers identify patterns and trends in the data, enabling them to make data-driven decisions to improve process control and reduce defects. By leveraging advanced analytics, semiconductor manufacturers can enhance product quality, reduce production costs, and improve their competitive position in the market.

Quality control and defect reduction ensure that semiconductor products meet stringent quality standards. By analyzing data from various inspection and testing stages, manufacturers can identify defects early and implement corrective actions to enhance product quality. Comprehensive data is essential for effective quality control, including inspection data, sort data, and wet process data. Inspection data from optical and electron microscopy inspections provide detailed information on surface and structural defects. Sort data from final electrical and functional testing of dies reveals performance issues and defects. Wet process data from cleaning, etching, and chemical deposition processes help monitor and control process consistency. Analytics techniques such as defect mapping, correlation analysis, and root cause analysis help identify defects early, pinpoint root causes, and implement targeted corrective actions

Predictive maintenance leverages data analytics to foresee equipment failures before they occur, allowing manufacturers to perform maintenance proactively rather than reactively. This approach minimizes unplanned downtime, extends the lifespan of machinery, and optimizes maintenance schedules. Predictive maintenance relies on various data sources such as sensor data, historical maintenance records, and operational logs. Sensor data includes temperature, pressure, vibration, and current readings from manufacturing equipment, providing insights into equipment health. Historical maintenance records and operational logs offer data on past maintenance activities, equipment failures, usage, and performance over time, essential for developing predictive models. Analytics techniques used in predictive maintenance include machine learning models, anomaly detection, and trend analysis.

Predictive maintenance, yield optimization, and quality control and defect reduction are critical use cases in the semiconductor industry that rely on advanced data analytics. These use cases not only enhance operational efficiency and reduce costs but also drive continuous improvement and innovation, positioning semiconductor manufacturers for sustained growth and competitive advantage in a rapidly evolving market.

Let us delve into the case study of a semiconductor customer who successfully migrated from an on-premises data warehouse to Amazon Redshift Manufacturing Analytics data lake, achieving peta byte scale data ingestion with sub second query performance.

Customer background :

The customer operates a network of fabs spanning different geographical locations, generating approximately ~7 petabytes (PB) of data daily. This data encompasses a diverse range of sources, including fabrication data such as lot tracking, inventory, cycle time fault detection, as well as business data such as sales, finance, and order fulfillment. The sheer volume of this data, amounting to ~10 terabytes (TB) per day, posed a significant challenge for the customer's legacy data lake, which struggled to ingest and process such a vast amount of information in a timely manner. The consumer applications, including Power BI, Storyboards, dashboards, and downstream queries, serve a diverse set of over 1,000 business users. These tools are designed to accommodate a substantial number of concurrent users, exceeding 300 during peak business hours, while maintaining sub-second response times. The growing demand for low-latency data is primarily driven by applications that leverage Machine Learning and Generative AI, which require real-time access to large datasets for processing and analysis.

Recognizing the need for a scalable and efficient solution, the customer sought to migrate their Manufacturing Analytics data lake to Amazon Redshift.

The move to Amazon Redshift was driven by the customer's strategic objective to leverage the extensive range and seamless integration of native services offered by the Modern Data Architecture on AWS. This move involved the modernization of four data warehouses, encompassing over 20+ trillion rows of data, into the AWS environment, with a daily ingestion rate of over 50+ billion rows combined.

Solution Overview

Amazon Redshift is designed for high-performance data warehousing, which provides fast query processing and scalable storage to handle large volumes of data efficiently. Its columnar storage format minimizes I/O and improves query performance by reading only the relevant data needed for each query, resulting in faster data retrieval. Lastly, you can integrate Amazon Redshift with data lakes like Amazon Simple Storage Service (Amazon S3), combining structured and semi-structured data for comprehensive analytics.

The following diagram shows a high-level illustration of the architecture.

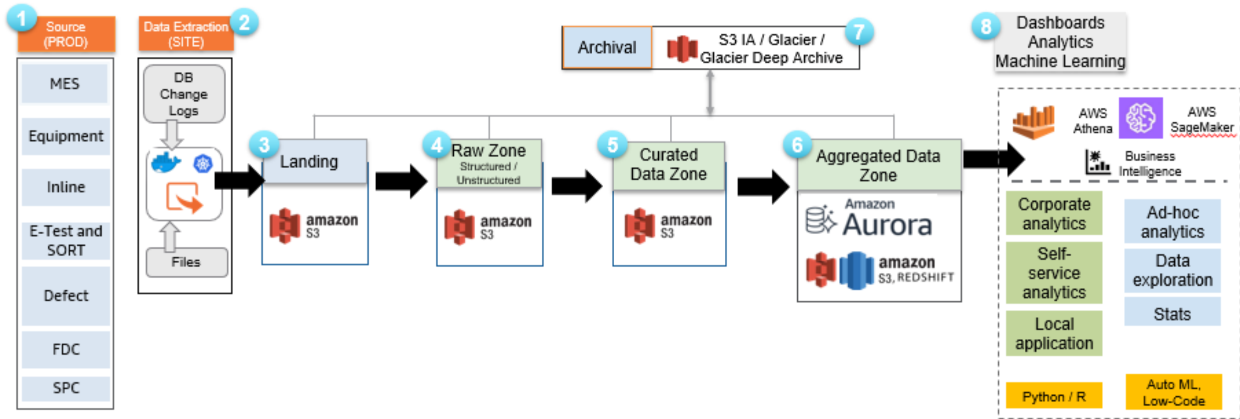

The customer captured raw unstructured data from sources like sensors, MES and DB logs using SteamSets in raw unstructured storage buckets, such as Amazon S3. This data is then enriched and processed using tools like Amazon EMR, Apache Spark and AWS Lambda. These tools can process and analyze large volumes of raw data in parallel, enabling faster and more efficient data processing. Once the raw data is transformed, it is stored in curated data lakes. A curated data lake follows a specific schema or data model, which is defined by a central authority.

Once the data is in the data lake, it is ready for further analysis and processing, Amazon Athena which is a is a serverless, cost-effective, and fully managed query service that makes it easy to analyze data stored in Amazon S3. It supports a wide range of SQL-based queries, allowing users to extract, transform, and analyze data from the data lake using a simple and intuitive interface.

The final stage of data ingestion involves aggregating the data from the data lake into a final copy before loading it into a hot database tier. This hot database tier is typically a relational database management such as Amazon Aurora PostgreSQL-Compatible Edition and /or Amazon Redshift. The purpose of this final copy is to provide a high-performance and scalable storage solution for real-time analytics and reporting.

Amazon Redshift allowed the customer to meet their scalability and performance requirements. Amazon Redshift features such as Workload Management (WLM), massively parallel processing (MPP) architecture, concurrency scaling, federated query feature and parameter groups helped address the requirements:

- WLM provided the ability for query prioritization and managing resources effectively by allowing to adjust the Workload Management Rules for high volume ingestion data warehouse to use manual workload management to throttle down the ingestion rate, while allowing separate queue for consumption workload. For other low volume ingestion data warehouse use Automatic Workload management, with queue priorities.
- Amazon Redshift federated query feature was a key differentiator for the customer's decision to adapt to Amazon Redshift. It allowed them to extend their data warehouse to run operational analytics on their transactional data store on Amazon Aurora PostgreSQL. There was a need to improve federated query performance while joining Redshift local tables to Aurora postgres tables,

as some of the key consumptions queries were taking up to 5 minutes (vs 5 secs) when run from Redshift. Working with Redshift specialist and service team the Team achieved SLAs by : 1/ Pushing down the entire Aurora query ; 2/ Pushing down entire Aurora sub query, including if Aurora query is in create table as. The current release of Redshift includes these optimizations.

- The MPP architecture model provided horizontal scalability
- Concurrency scaling added additional cluster capacity to handle unpredictable and spiky workloads
- Parameter groups defined configuration parameters that control database behavior

Together, these capabilities allowed them to meet their scalability and performance requirements in a managed fashion.

Results

The customer’s decision to use Amazon Redshift for their solution was further reinforced by the platform’s ability to handle both structured and semi-structured data seamlessly. Amazon Redshift allows the customer to efficiently analyze and derive valuable insights from their diverse range of datasets, including equities and institutional data, all while using standard SQL commands that teams are already comfortable with.

Through rigorous testing, Amazon Redshift demonstrated remarkable performance, meeting the customer's stringent SLAs and delivering exceptional sub second query response times with impressive latency. With the AWS migration, the customer achieved a 3X improvement in query performance compared to on-prem enabling faster data analytics contributing to yield optimization. This data lake also serves as a rich repository for Gena/ML use cases. The Cloud ML predictive maintenance solution helped extend maintenance from 50,000 wafers to > 100,000 wafers. The Virtual Metrology AI initiative is driving capex reduction and tool throughput increase. This platform also helped in lead time reduction for NPI on RF tech ramp from 21 days to 9 days (57% reduction) enabling speed to market.

Conclusion

In this post, we covered how a large semiconductor customer improved performance and scalability by migrating to Amazon Redshift. This enabled the customer to grow and onboard new workloads into Amazon Redshift for their business-critical applications.

Appendix:

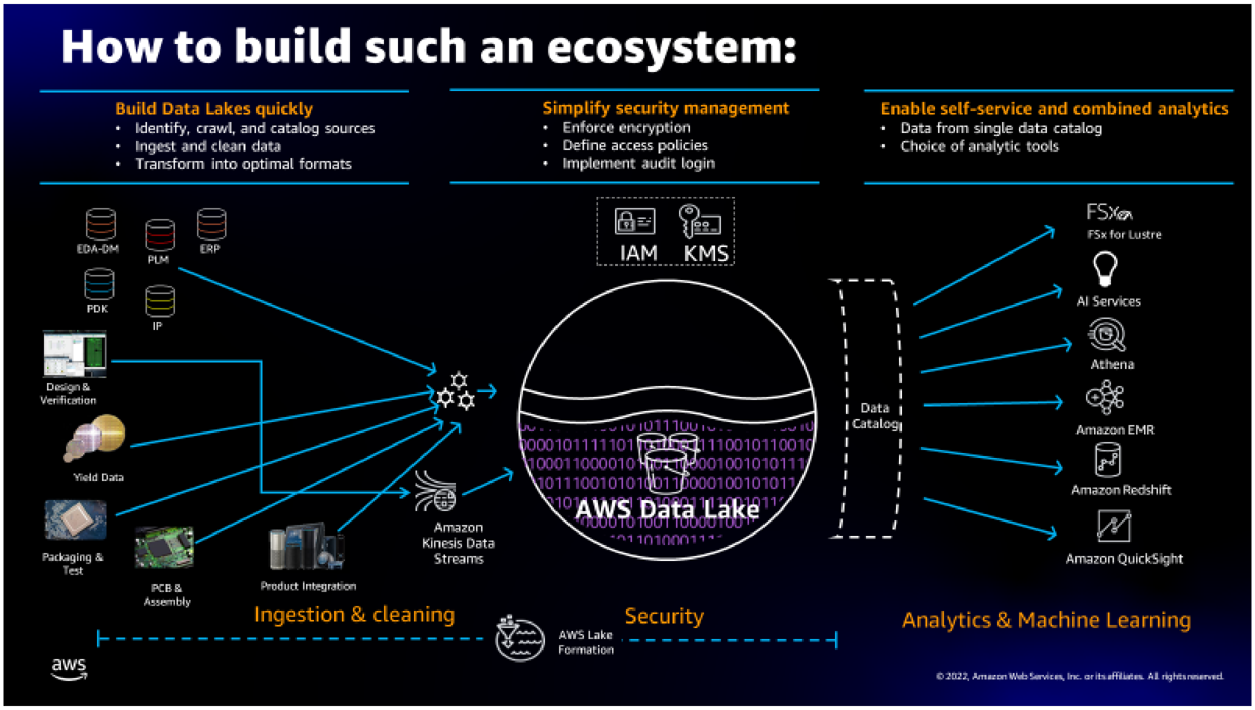
How to build such an ecosystem:
Build Data Lakes quickly
• Identify, crawl, and catalog sources
• Ingest and clean data
• Transform into optimal formats
Simplify security management
• Enforce encryption
• Define access policies
• Implement audit login
Enable self-service and combined analytics
• Data from single data catalog
• Choice of analytic tools
EDA-DM
PLM
ERP
PDK
IP
Design & Verification
Yield Data
Packaging & Test
PCB & Assembly
Product Integration
Amazon Kinesis Data Streams
IAM
KMS
AWS Data Lake
Data Catalog
FSx for Lustre
AI Services
Athena
Amazon EMR
Amazon Redshift
Amazon QuickSight
Ingestion & cleaning
Security
Analytics & Machine Learning
AWS Lake Formation
aws
© 2022, Amazon Web Services, Inc. or its affiliates. All rights reserved.