



Real-time platform for Big Data

A \$20B aerospace company needed an enterprise-wide, real-time data platform. A system that could secure and control access to highly sensitive commercial and military data. With a worldwide presence and 70,000 employees in 300 locations, the company had hundreds of data sources with hundreds of thousands of tables and data types, and a need to ingest up to a petabyte of information each day.

In partnership with Amazon Web Services (AWS), productOps enabled aggregation of these large and complex data sources to increase customer satisfaction, improve operational efficiency and accelerate innovation.

The Challenge

Previously, the company had required months of lead time to obtain access to data or to undertake any kind of reporting or Business Intelligence analysis. With an enterprise-wide, real-time platform that could be accessed efficiently and securely this would be cut significantly. In addition, in this highly regulated industry, the security and classification of data was an important consideration in moving to a new platform. It was essential the end-to-end solution would be General Data Protection Regulation (GDPR) and International Traffic in Arms Regulations (ITAR) compliant.

productOps and AWS

Data governance and data discovery were key challenges as the platform would touch all aspects of the business. The company needed actionable and accessible data in real-time so staff could pull the information they needed in a uniform way to support data science, business intelligence and aerospace application performance.

The wisdom of a cloud-based solution was self-evident, but Microsoft Azure was the company's preferred choice. After a detailed comparative analysis was presented and productOps performed a stunning sales coup the company understood they needed to leverage productOps experience and AWS.

In short time, considering the scale of data involved, productOps built a robust, modern stack - a highly scalable, cloud-hosted data platform. It included elements such as a data ingestion tool, data lake and a data warehouse to support the management and storage of thousands of different entities throughout the organization.

KEY AWS SERVICES USED

Glue

Data preparation is time consuming. Glue allowed us to catalog data for analytics.

IAM

Security was paramount in this project. Identity and Access Management services were crucial in controlling individual and group access to AWS resources.

Kinesis

The real-time streaming platform that allowed Collins Aerospace to level up their data access.

S3

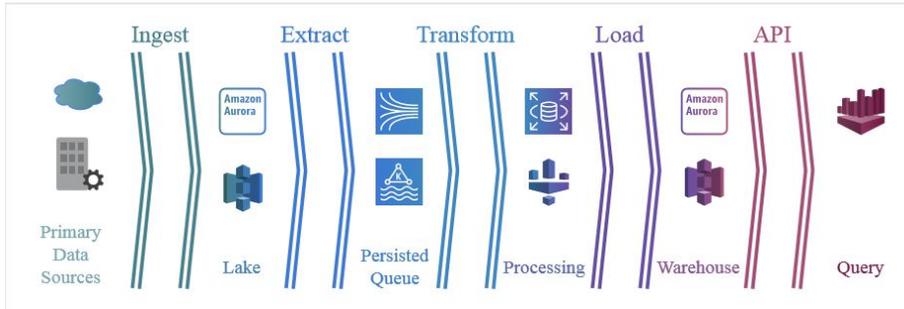
A fully managed, petabyte-scale data lake capable of handling the huge scale of this project.

Postgres Aurora

The cloud-first SQL service that makes it easy to analyze data using standard SQL with no infrastructure to setup or manage.

productOps and AWS Bring Needs-Based Change

The first steps involved a data compliance and scrubbing operation that would clean the data up so it could be ingested into the new centralized system. Next a platform had to be built that could handle the massive number of transactions (estimated to be 200,000 per second) required by the system. APIs, visualizations and web applications were built on top of the data lake and data warehouse to be able to efficiently access the data.



Security of the data in the cloud was a top priority and platform functionality had to comply with a myriad of regulations governing the security and access aspects of the business. By using AWS they were assured a credible and accredited service that covered their very specific needs.

Even though productOps is an AWS Advanced Consulting Partner, key to the rapid implementation of the project was support from Amazon. An AWS Solutions Architect was assigned to the project from the start and present at the meetings where crucial road blocks could be resolved without delay.

“Where typically we would wait 1 or 2 weeks for an answer relating to the platform with services we’ve used previously, the AWS Solutions Architect assigned to this project would be sitting right there ready with an answer that allowed us to move the project forward immediately.” Paul Iverson- Lead System Architect of productOps.

Capable of ingesting and storing the siloed data from across the organization, the new platform also supported real-time streaming of changes throughout the enterprise. This enabled easy monitoring of events throughout the organization. Turn around time of updates that previously took upwards of one hour to propagate throughout the existing infrastructure could now be accomplished in seconds.

Today, the company has a secure, efficient platform to pull insights from. Since the inception of the new data platform they are able to ingest streaming and bulk data at a scale of hundreds of terabytes a day. They now have an easy to maintain, scalable and highly secure cloud-based data platform that allows them to continue to be a leader in innovation in the aerospace industry.

USE CASE

Story of the Part

Problem:

No single source of record for aircraft part data. Component usage data tracked manually (often by .csv file)

Compelling Benefit:

End to end visibility of a part from manufacturing, through repair and support, allowing customers to track the lifecycle of parts through their fleet.

Key Challenges:

Aggregation of part usage from multiple sources (previously a manual process.)

Parts changing operators making it more difficult to track

Solution:

A clean, integrated system that pulls from the part information master

“The value of integrating part usage data is incredibly enlightening. We can now get data back to our design community and predictive data to our Maintenance, Repair and Operations staff.”