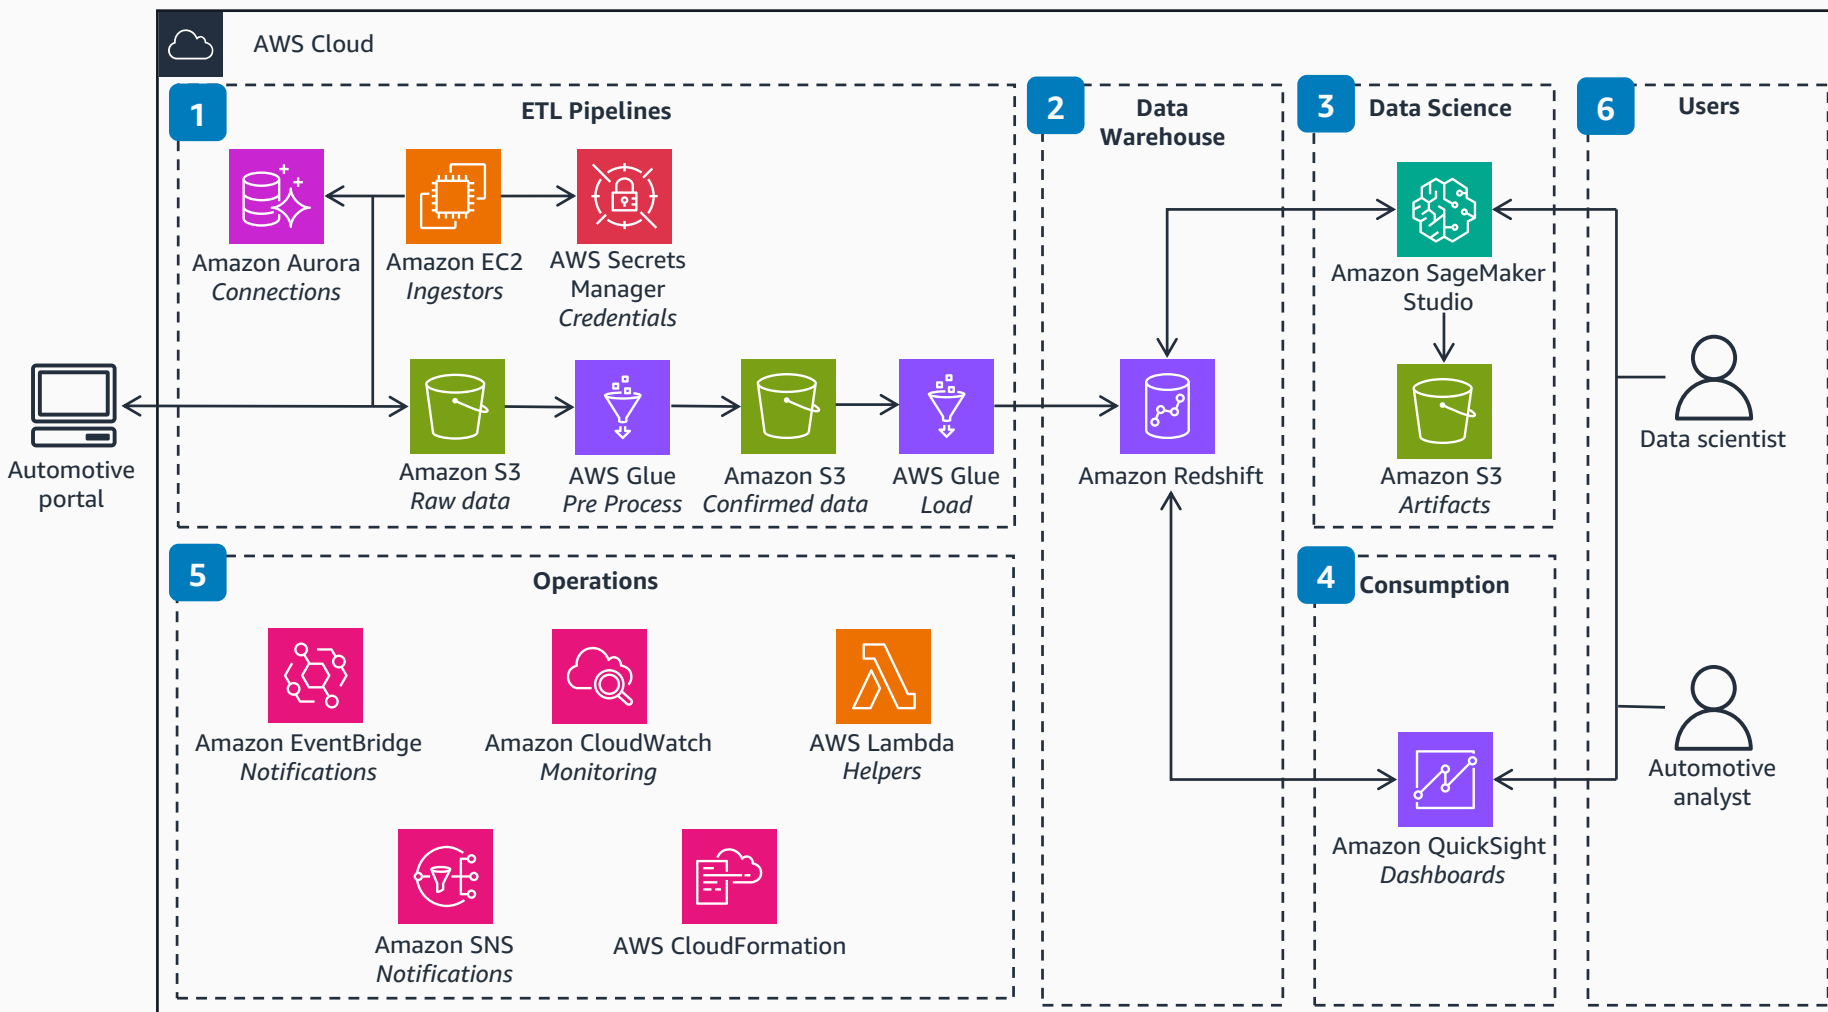


Guidance for Automotive Warranty Analytics on AWS

This architecture provides a repeatable way to ingest, transform, and consume inferences from an ML model to predict major product defects and recalls. It then enable corrective actions.



- 1 Scripts running on **Amazon Elastic Compute Cloud (Amazon EC2)** periodically ingest automotive warranty claims from automotive portals and store them in **Amazon Simple Storage Service (Amazon S3)** buckets. This data then goes through scalable extract, transform, load (ETL) pipelines implemented in **AWS Glue**.
- 2 **Amazon Redshift**, a centralized data warehouse, then hosts this data, which includes enriched data and tables for specific analytical dashboards. **Amazon Redshift** can scale to meet the needs of an entire enterprise or organizational unit.
- 3 **Amazon SageMaker Studio** provides data scientists and analysts with a comprehensive tool chain for data exploration, model training, and machine learning operations (MLOps) pipelines, all in one place.
- 4 The broader community of analysts and users obtain specific actionable recommendations through dashboards deployed on **Amazon QuickSight**.
- 5 **Amazon EventBridge** initiates periodic ingestion and ML pipelines. **Amazon CloudWatch** provides logging and monitoring capabilities.
- 6 Data scientists and automotive analysts iteratively develop and review analyses developed in **SageMaker Studio** and **QuickSight** dashboards.

