Guidance for Financial Regulatory Reporting on AWS

This reference architecture shows how to build a data analytics platform on AWS for improved regulatory reporting.

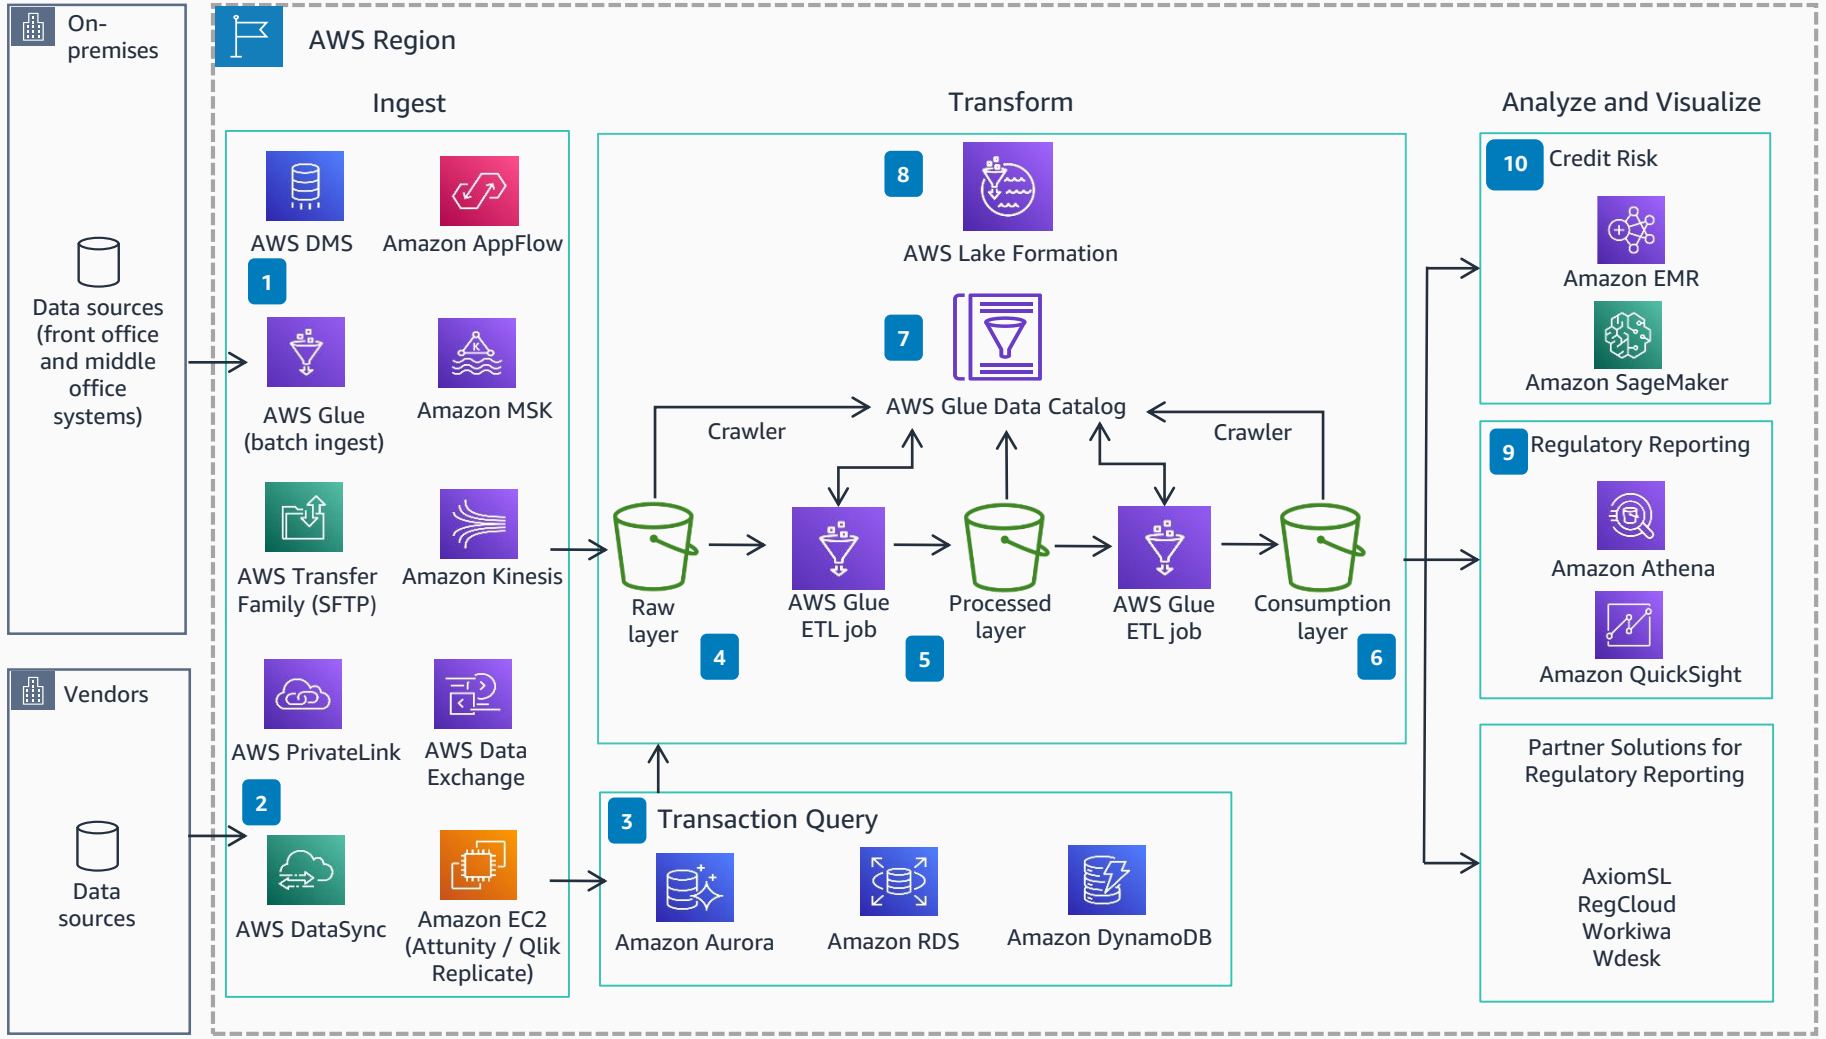

1. Ingest on-premises data using services such as **AWS Database Migration Service (AWS DMS)** for databases, **AWS Glue** for batch data, **or Amazon Managed Streaming for Apache Kafka (Amazon MSK)** for streaming data.
2. Access vendor data by connecting to the vendor account directly using **AWS PrivateLink** or by using services such as **AWS Data Exchange** and **AWS DataSync**.
3. Replicate transaction data from data sources on AWS (such as **Amazon Aurora**, **Amazon Relational Database Service [Amazon RDS]**, or **Amazon DynamoDB**) to a **Amazon Simple Storage Service (Amazon S3)** bucket using a service such as **AWS DataSync**.
4. All the data is stored as-is in the raw layer without undergoing any changes.
5. Data undergoes basic transformations in the processed layer, such as normalizing the date to a certain format or cleaning up empty rows.
6. The consumption layer contains the final "cleansed" copy of the data to be used across a number of different use cases, including regulatory reporting.
7. **AWS Glue Data Catalog** provides a view of the metadata of all the data across the different **S3** buckets.
8. **AWS Lake Formation** centrally manages access to the available datasets and applies fine-grained permissions for all users accessing the data.
9. Use **Amazon QuickSight**, a data visualization and business intelligence service, for reporting. Use **Amazon Athena** for interactive analytics.
10. Use services such as **Amazon EMR** and **Amazon SageMaker** for credit risk calculations and forecasting.

Reviewed for technical accuracy January 9, 2023

AWS Reference Architecture