



AWS
re:Invent

D A T 3 6 0 - R

Analytical use cases with AWS Database Migration Service

John Winford

Senior Manager
Amazon Web Services

Riyaz Shiraguppi

Senior SDE
Amazon Web Services

Abhinav Singh

Database Engineer
Amazon Web Services

Agenda

Analytics overview

Replication technology overview

Online streaming analytics

Business intelligence (BI) and reporting with replication

Demos, demos, demos

Related breakouts

DAT214 Automation framework to migrate relational databases

DAT342 Lessons learnt from migrating Oracle databases to Amazon Aurora

DAT343 Lessons learnt from migrating SQL Server databases to Amazon Aurora

DAT345 Assessing and categorizing your database migrations

DAT362 Dive deep into AWS SCT and AWS DMS

DAT363 Migrating your data warehouses to Amazon Redshift

DAT366 Running Oracle on Amazon RDS and migrating to Aurora PostgreSQL

DAT367 Running SQL Server on Amazon RDS and migrating to Aurora MySQL

DAT377 Getting started with AWS DMS and AWS SCT

AWS analytics overview

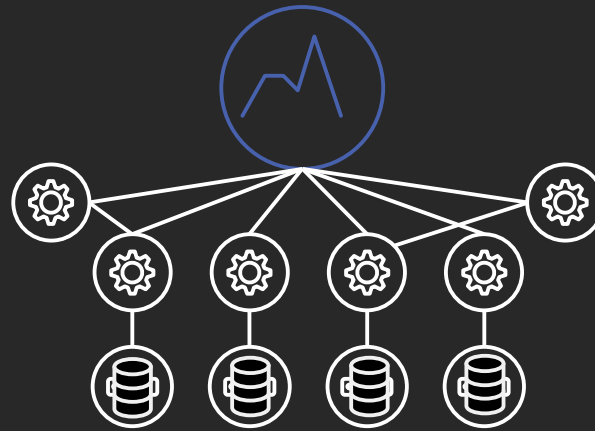
Trends that impact the way we think about data

Explosion of data



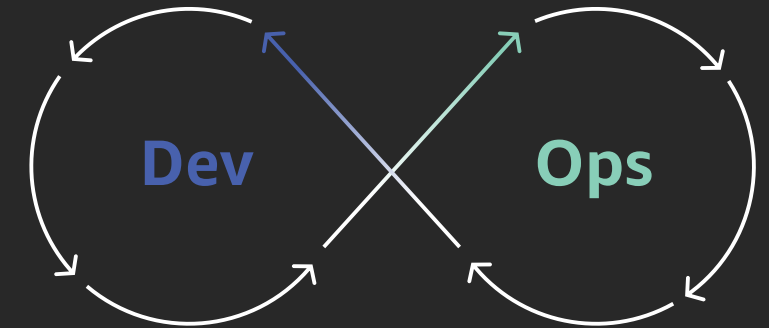
Data grows 10x every 5 years driven by network-connected smart devices

Micro-services change data and analytics requirements



Micro-services architecture decreases need for one-size fits all databases and increases need for real-time monitoring and analytics

Rapid rate of change driven by DevOps



Transition from IT to DevOps increases rate of change

Modern apps create new requirements



Ride hailing



Media streaming



Social media



Dating

Users: 1 million+

Data volume: TB, PB, EB

Locality: Global

Performance: Milliseconds, microseconds

Request Rate: Millions

Access: Web, Mobile, IoT, and devices

Scale: Up, down, out, and in

Economics: Pay for what you use

Developer access: Instant API access

Our portfolio

Broad and deep portfolio, purpose-built for builders

Business intelligence & machine learning



Amazon QuickSight



Amazon SageMaker

Analytics



Amazon Redshift
Data warehousing



Amazon EMR
Hadoop + Spark



Amazon Athena
Interactive analytics



Amazon Elasticsearch Service
Operational analytics



Amazon Kinesis Data Analytics
Real time

Databases



Amazon Aurora
MySQL, PostgreSQL



Amazon DynamoDB
Key value, document



Amazon QLDB
Ledger Database



Amazon RDS
MySQL, PostgreSQL, MariaDB, Oracle, SQL Server



Amazon ElastiCache
Redis, Memcached



Amazon RDS on VMware



Amazon Neptune
Graph



Amazon Timestream
Time Series



Amazon Managed Blockchain



AWS Blockchain Templates

Data lake



Amazon S3/Amazon S3 Glacier



AWS Lake Formation
Data lakes



AWS Glue
ETL & data catalog


Data movement

AWS Database Migration Service | AWS Snowball | AWS Snowmobile | Amazon Kinesis Data Firehose | Amazon Kinesis Data Streams

Amazon Redshift


Highly scalable cloud data warehouse

Virtually unlimited concurrency




Dynamically scales to support virtually unlimited number of concurrent users and growing data volumes

Extends your data lake



Analyze exabytes of data in the Amazon S3 data lake together with petabytes of data loaded into Amazon Redshift's high performance SSDs

10x performance



Get faster time-to-insight for all types of analytics workloads; powered by machine learning, columnar storage and MPP

1/10th the cost



Start at \$0.25 per hour, scale out as low as \$1,000 per terabyte per year

Amazon Kinesis


Real-time analytics and data streaming

Volume, velocity, and variety



Analyze and react in real time
to a variety of data events
affecting business

Elastic



Scales up and down as
demand changes, saving
money compared to
traditional options

Seamless integration



Leverage other core AWS
services such as Amazon
Elasticsearch Service,
Amazon S3, and AWS
Lambda

Easy to use



Launch a fully managed
analytics solution in
minutes

Amazon Elasticsearch Service

Fully managed, scalable, and secure

Easy to use



Fully managed; deploy production-ready clusters in minutes

Open



Direct access to Elasticsearch open-source APIs; supports Logstash and Kibana

Secure



Secure access with Amazon VPC to keep all traffic within AWS network

Available



Zone awareness replicates data between two AZs; automatically monitors & replaces failed nodes

Amazon Athena

Interactive query service to analyze data in Amazon S3 using SQL

Query instantly



Zero setup cost; just point to Amazon S3 and start querying

Pay per query



Pay only for queries run; save 30–90% on per-query costs through compression

Open



ANSI SQL interface, JDBC/ODBC drivers, multiple formats, compression types, and complex joins and data types

Easy



Serverless: zero infrastructure, zero administration; integrated with Amazon QuickSight

Amazon QuickSight

First BI service with pay-per-session pricing and machine learning (ML) insights for everyone in your organization



Serverless, cloud-powered BI service (no servers to manage)

Scale from 10s of users to 100s of thousands of users

Pay only for what you use


- Readers: \$0.30/30 min session with a \$5/user/month max
 - Authors: \$18/month/author
-

Integrates with Amazon S3, Athena, Amazon Redshift, Amazon RDS, Aurora, and Amazon EMR

Migration and Replication Basics

AWS Migration Tooling

Our goal: Allow customers the freedom to choose the best data platform for their needs **#DBFreedom**



AWS Schema Conversion Tool (SCT) converts your commercial database and data warehouse schemas to open-source engines or AWS-native services, such as Amazon Aurora and Amazon Redshift

AWS Database Migration Service (DMS) easily and securely migrates and/or replicate your databases *and* data warehouses to AWS

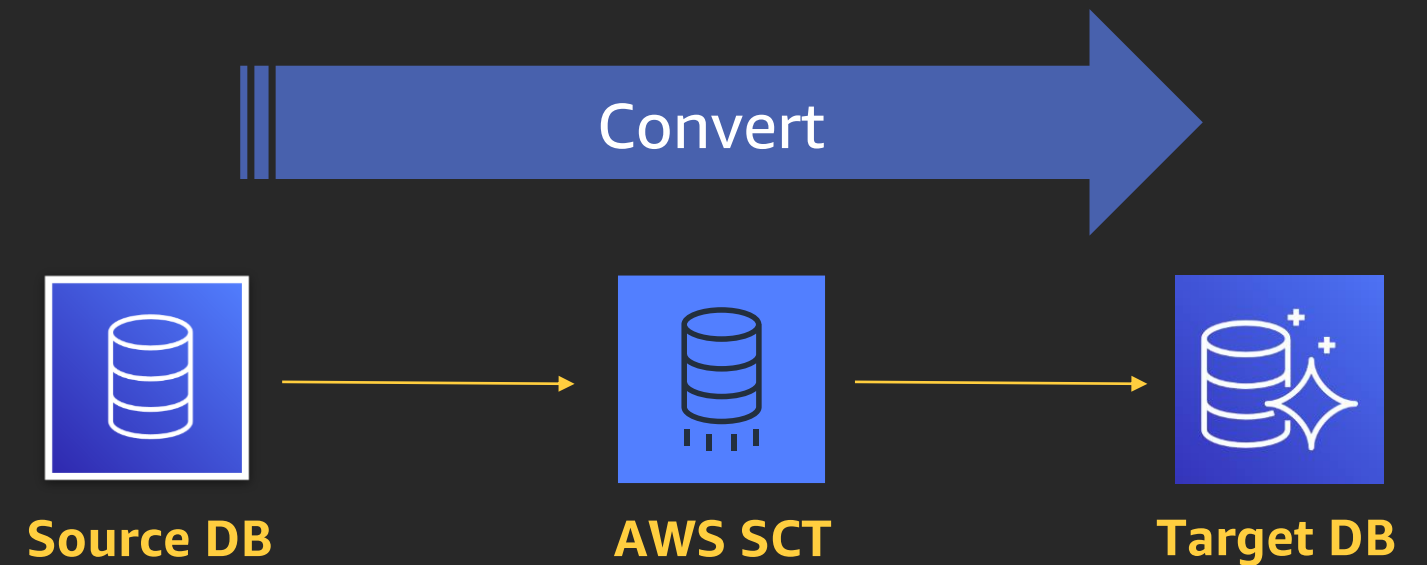


AWS Schema Conversion Tool

The AWS Schema Conversion Tool (AWS SCT) helps automate many database schema and code conversion tasks when migrating from source to target database engines

Features

- Create assessment reports for homogeneous/heterogeneous migrations
- Convert database schema
- Convert data warehouse schema
- Convert embedded application code
- Code browser that highlights places where manual edits are required
- Secure connections to your databases with SSL
- Service substitutions/ETL modernization to AWS Glue
- Migrate data to data warehouses using AWS SCT data extractors
- Optimize schemas in Redshift




AWS SCT product highlights



Assess



Plan



**Convert schema
and code**



Optimize



**Migrate
data warehouses**

Supported source and targets


Relational


NoSQL

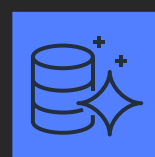
Data Lake


Data Warehouse*


Sources



Oracle



MySQL



Amazon Aurora


PostgreSQL


SQL Server


MariaDB


SAP ASE


Db2 LUW

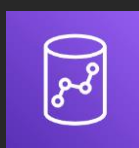

SQL Azure

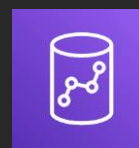

MongoDB

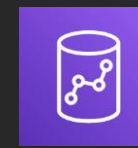

Cassandra

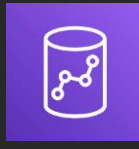

Amazon S3

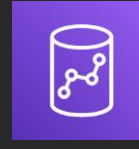

AWS Snowball

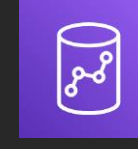

Oracle


SQL Server



Netezza



Greenplum

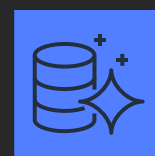

Teradata


Vertica


Targets



Oracle



MySQL

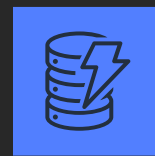

Amazon Aurora

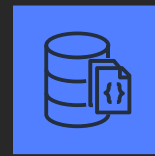

PostgreSQL

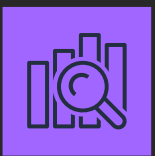

SQL Server



MariaDB

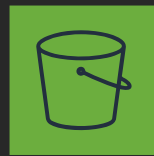

SAP ASE


Amazon DynamoDB


Amazon DocumentDB (with
MongoDB compatibility)


Amazon Elasticsearch
Service


Amazon Kinesis
Data Streams


Amazon S3


Amazon Redshift

* Supported via SCT data extractors


AWS DMS product highlights



Secure



Assess



Validate



**Snowball
integration**



Monitor



Stream data




Low cost



**Multiple
options**

The migration process



- Start a replication instance
- Connect to source and target databases
- Select tables, schemas, or databases



Application users

- Let AWS DMS load data, and keep them in sync
- Switch applications over to the target once in sync at your convenience

General use cases

Migrate



- **Migrate** business-critical applications
- **Migrate** data warehouses to Redshift
- **Upgrade** to a minor/major version
- **Consolidate** shards into Amazon Aurora
- **Archive** old data to Amazon S3
- **Migrate** from NoSQL to SQL, SQL to NoSQL or NoSQL to NoSQL

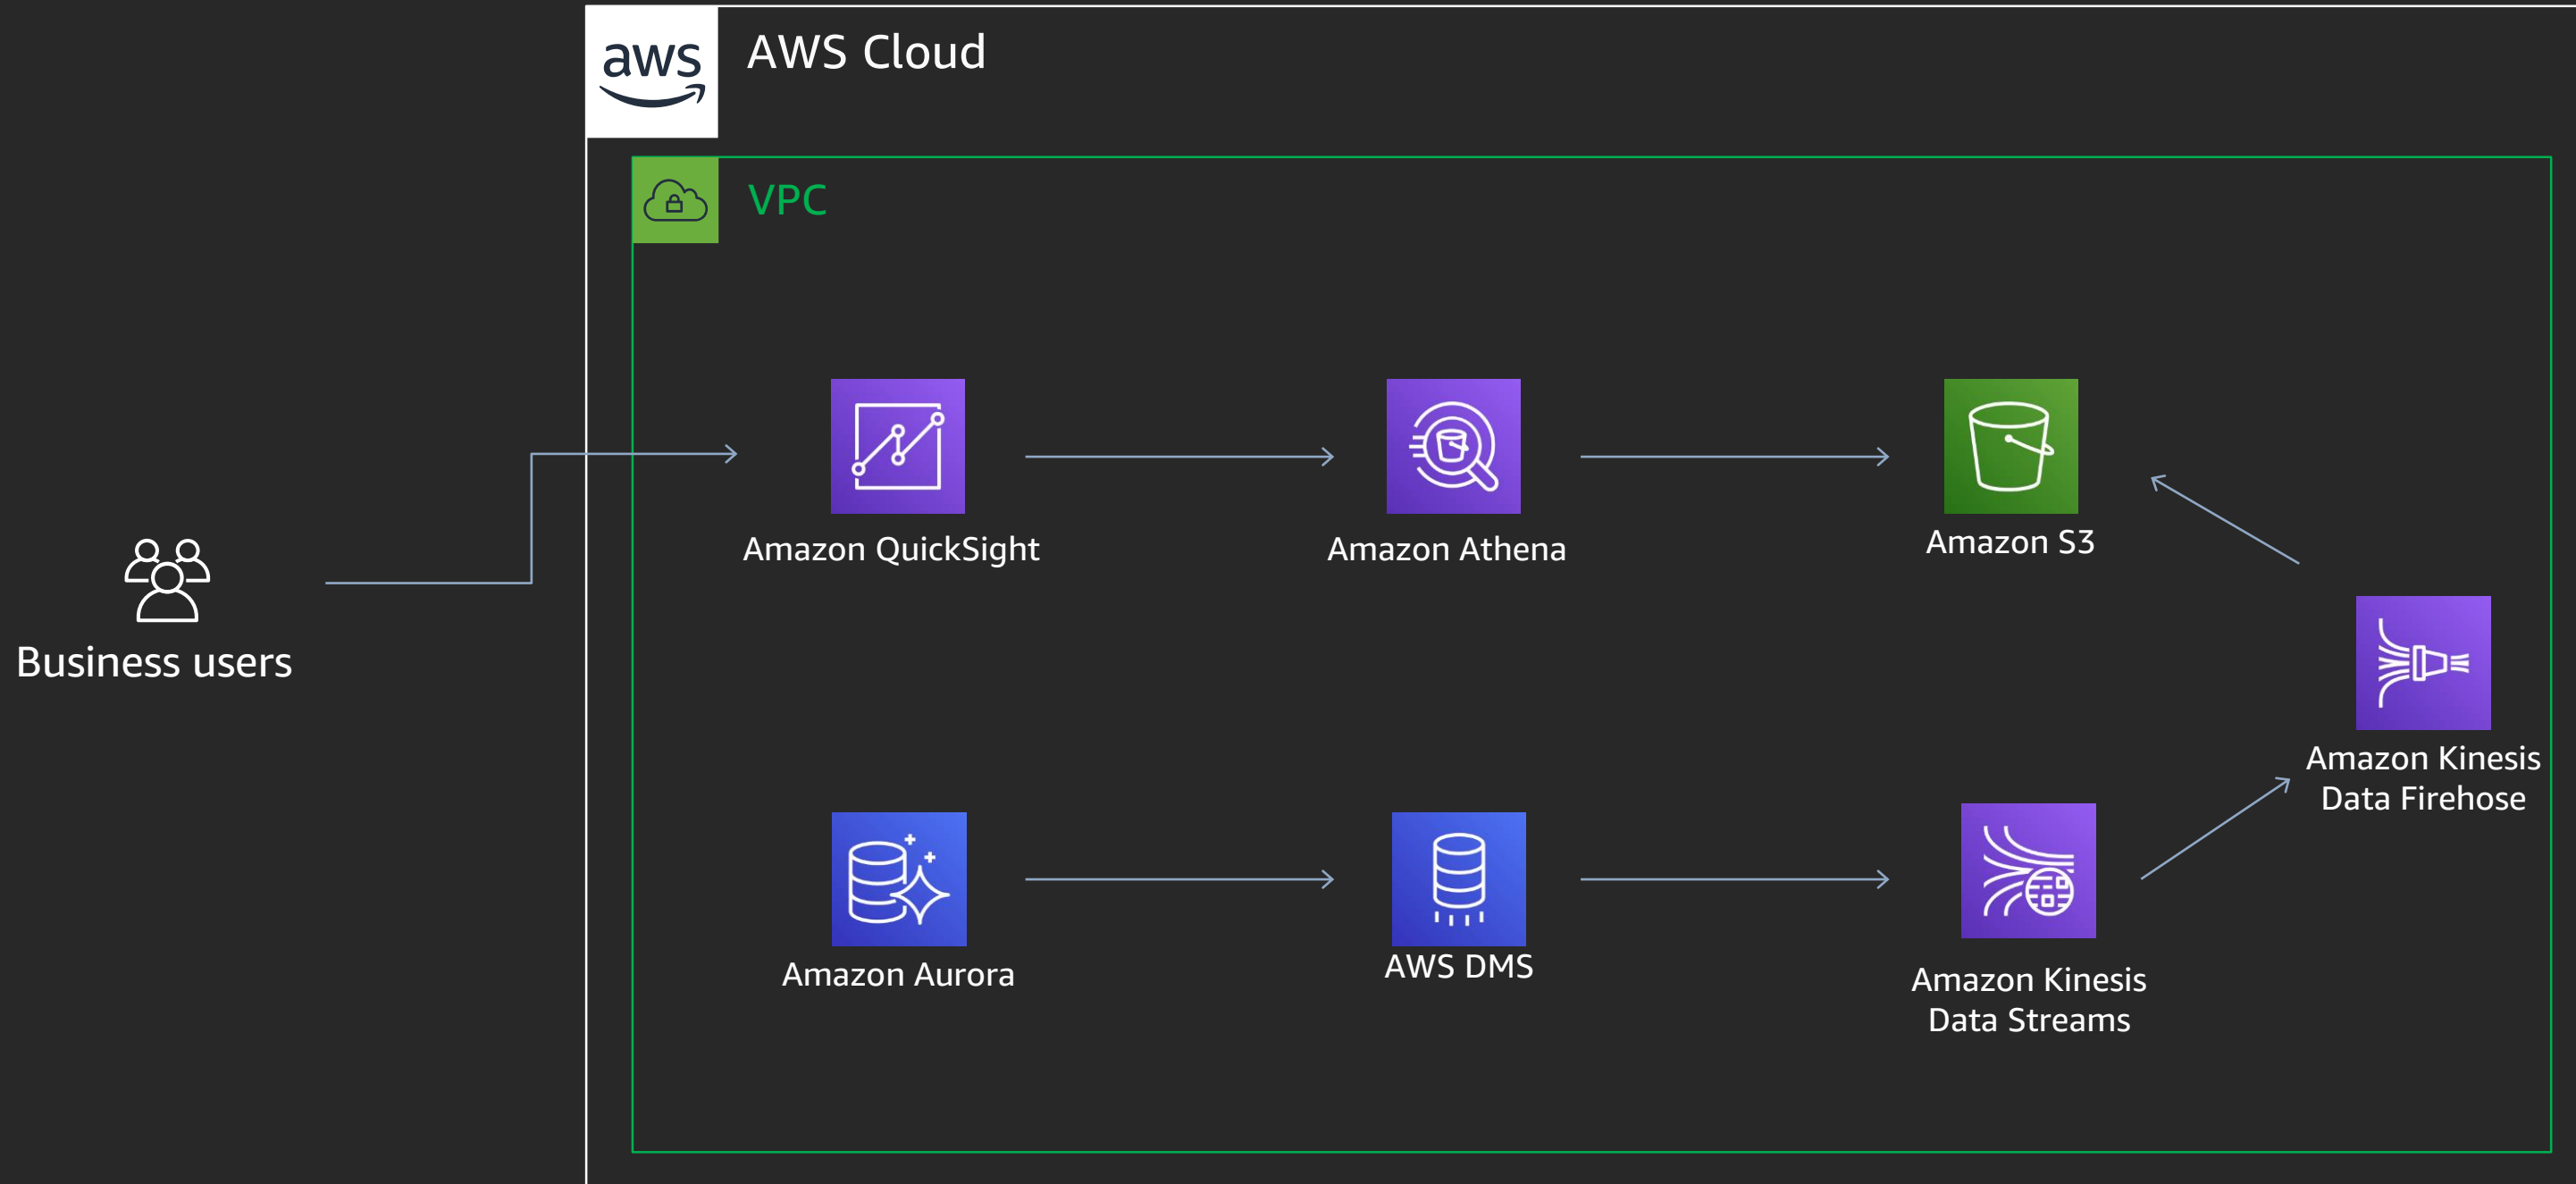
Replicate



- **Create** cross-region read replicas
- **Run** your analytics in the cloud
- **Hydrate** your data lakes
- **Replicate** to streaming platforms

DMS Analytics Use Cases

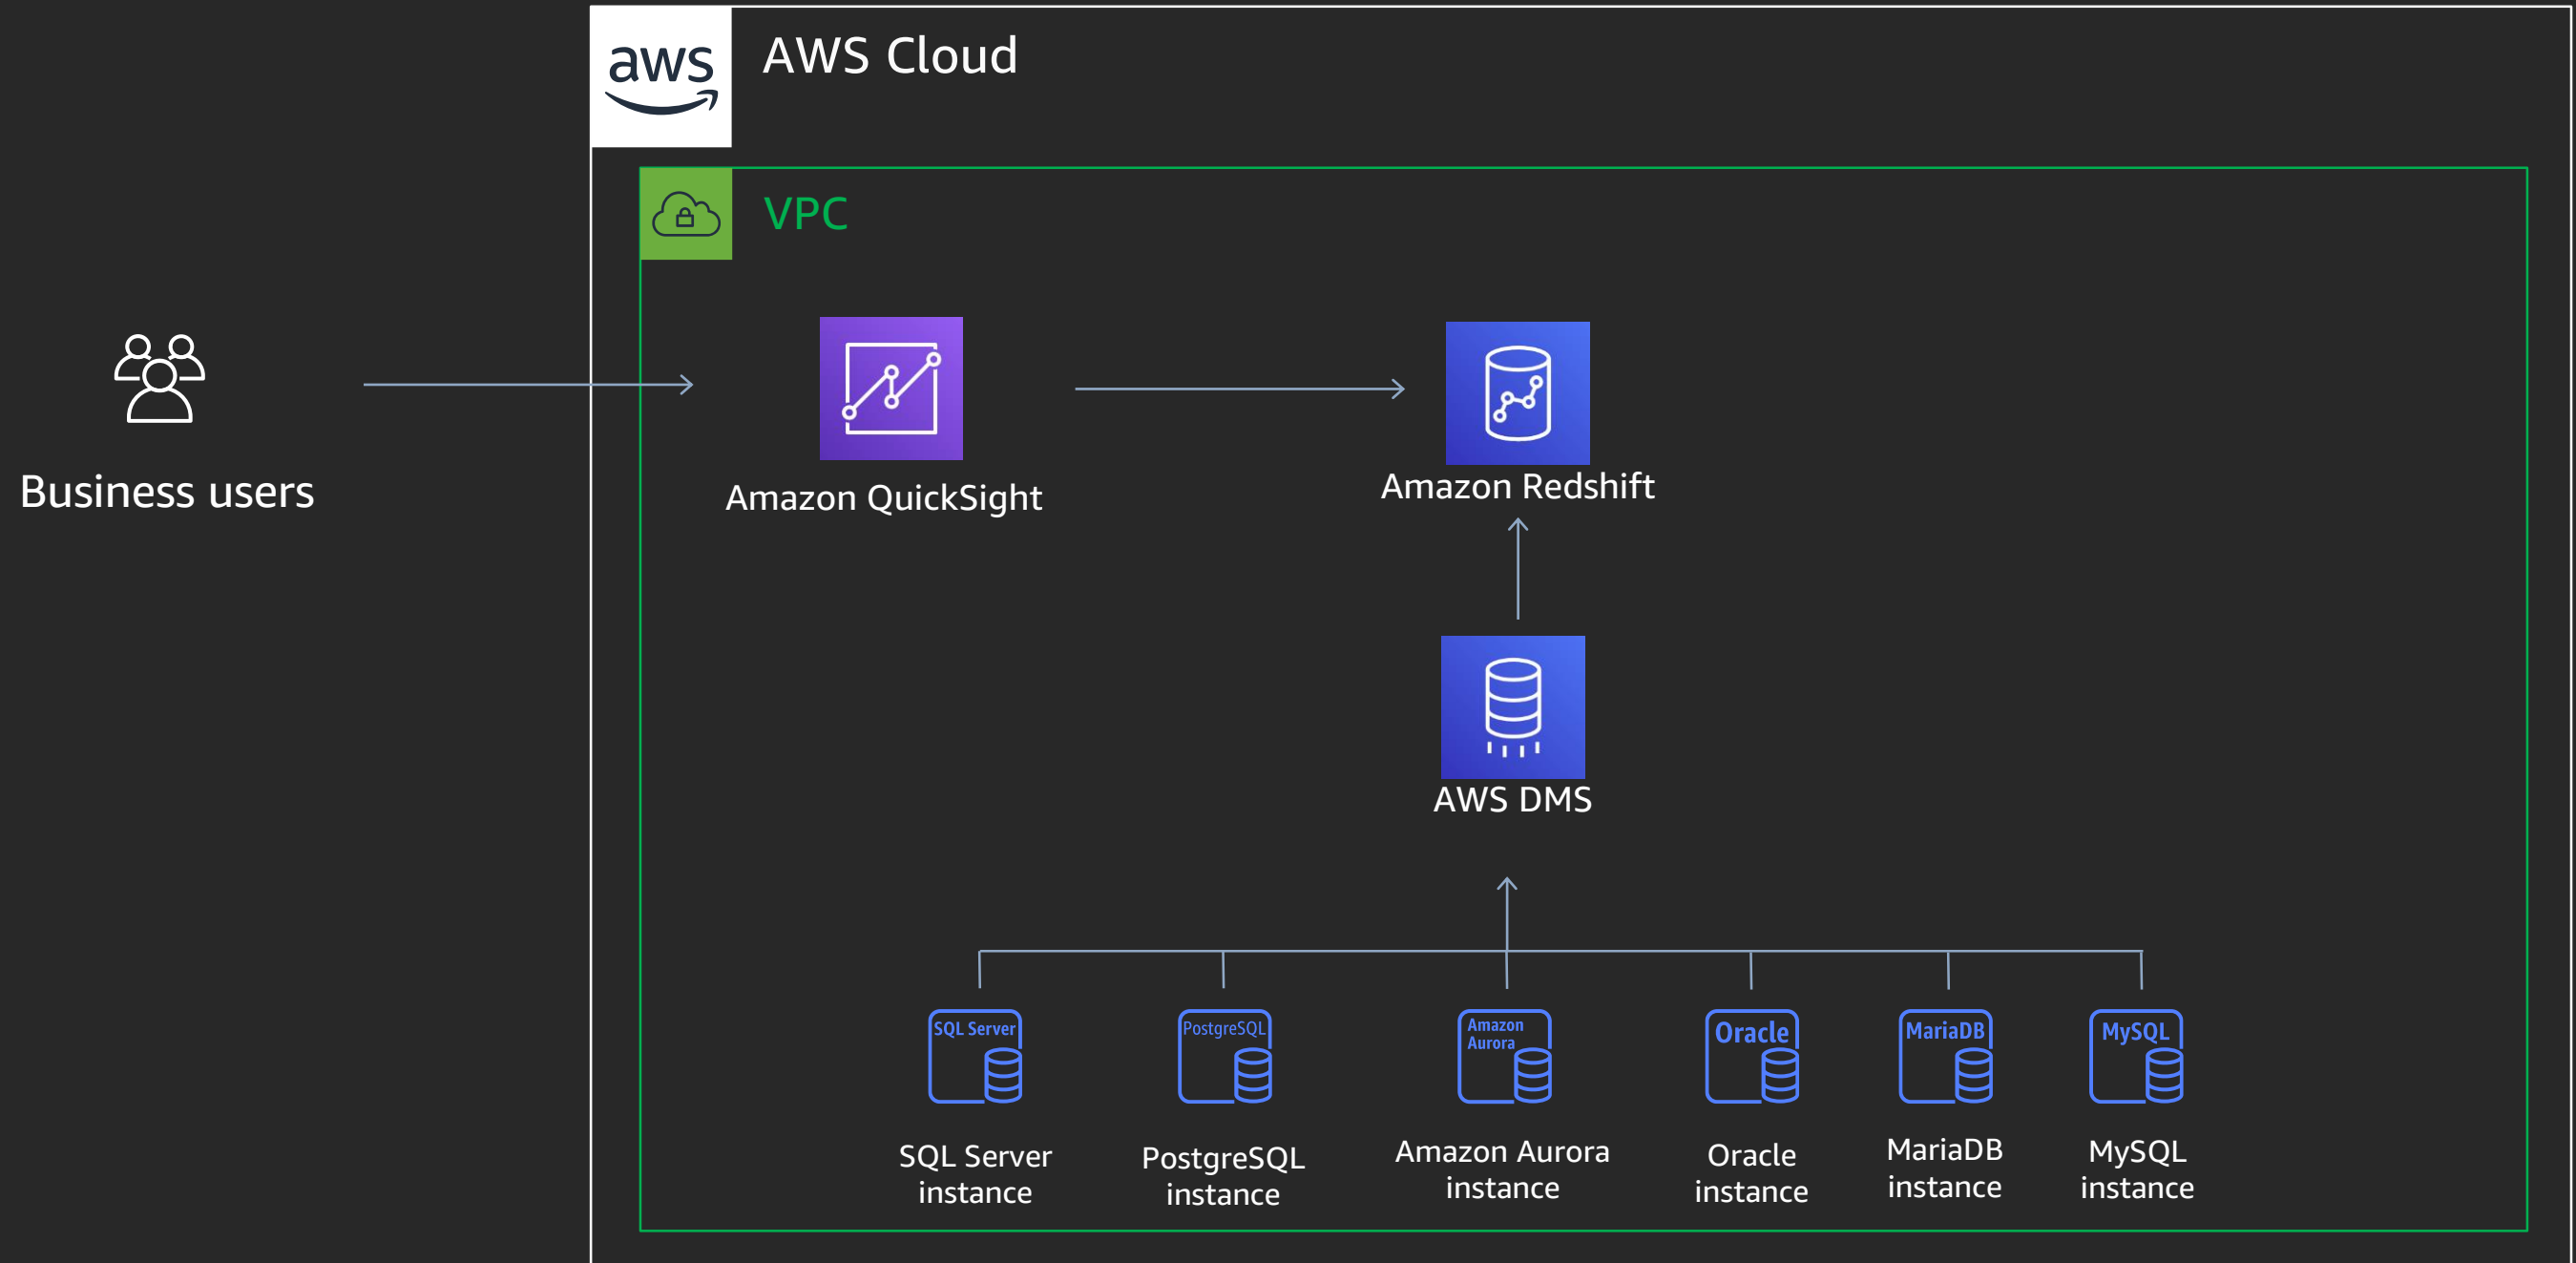
Analytics using AWS DMS - Example



Visualized in Amazon QuickSight



AWS Architecture - Demo




“Technology is cool, but you’ve got to use it as opposed to letting it use you.”

Prince

Demo


Learn databases with AWS Training and Certification

Resources created by the experts at AWS to help you build and validate database skills



25+ free digital training courses cover topics and services related to databases, including:

- Amazon Aurora
- Amazon Neptune
- Amazon DocumentDB
- Amazon DynamoDB
- Amazon ElastiCache
- Amazon Redshift
- Amazon RDS



Validate expertise with the new **AWS Certified Database - Specialty** beta exam

Visit aws.training

Thank you!



Please complete the session
survey in the mobile app.