AWS re:Invent

DAT209-L

Leadership session: AWS purpose-built databases

Shawn Bice
VP, Databases
Amazon Web Services

Tobias Ternström
Director, RDS & Aurora
Amazon Web Services

Joseph Idziorek
Principal Product Manager, DocumentDB
Amazon Web Services

App architectures & patterns have evolved over the years…

Mainframe

Client Server

Three tier

Microservices

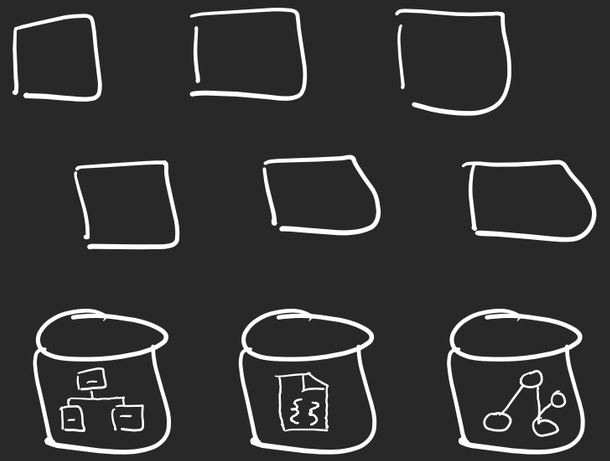

Common database categories

Top of mind for our customers

1. Move to Managed
2. Break free
3. New Modern Apps

Move to managed

Services

RDS
Aurora
ElastiCache
DocDB
Elasticsearch

Tools

SCT
DMS

Programs

MAP
ProServe
Partner

Break free from commercial

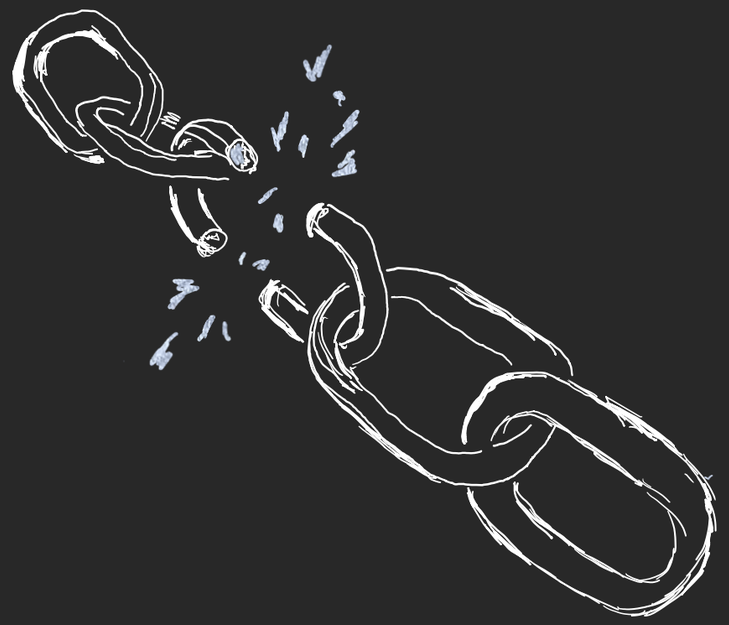

Services

Aurora

Amazon Redshift

Tools

SCT

DMS

Programs

MAP

DB Freedom

Pro Serve

Partner

New modern applications

New Requirements

Users:	1m+
Data Volume:	TB-PB-EB
Performance:	Milli-Micro sec.
Request Rate:	Millions+
Access:	Any device
Scale:	up-out-in
Economics:	Pay as you go
Developer Access:	Managed API

Modern (microservice) architectures

Customers

Key-value: Amazon DynamoDB

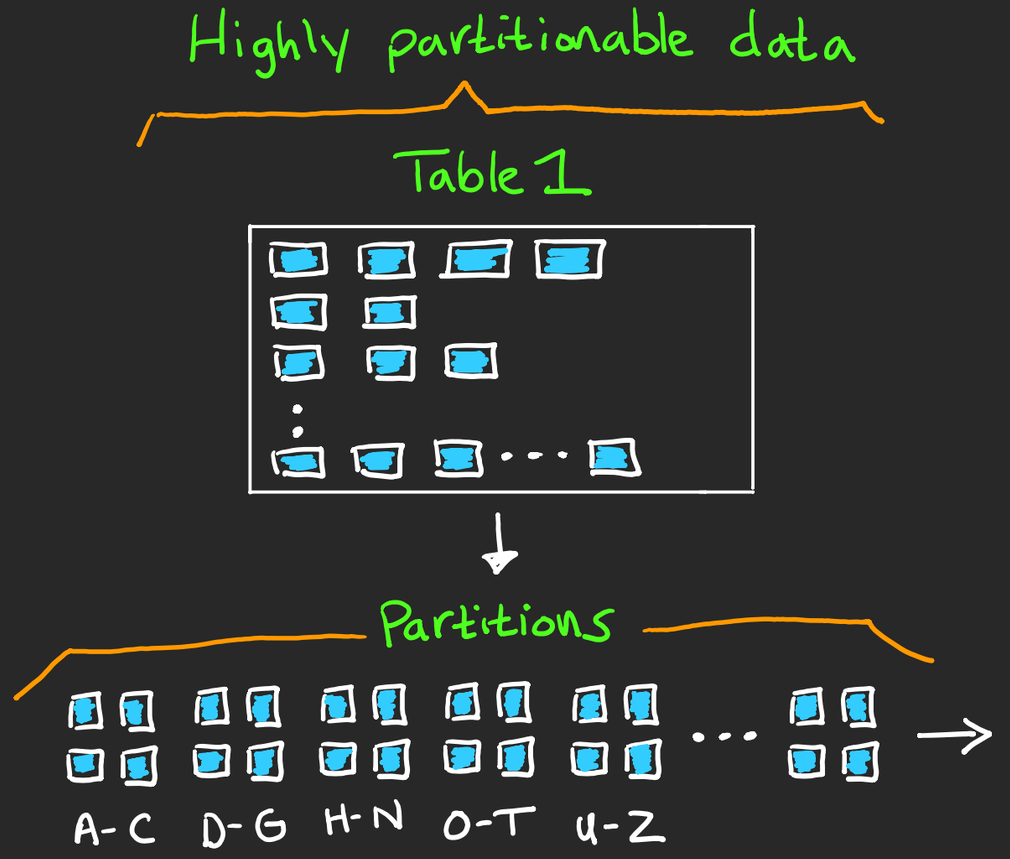

Relational: Amazon Aurora and Amazon RDS

ZipRecruiter®

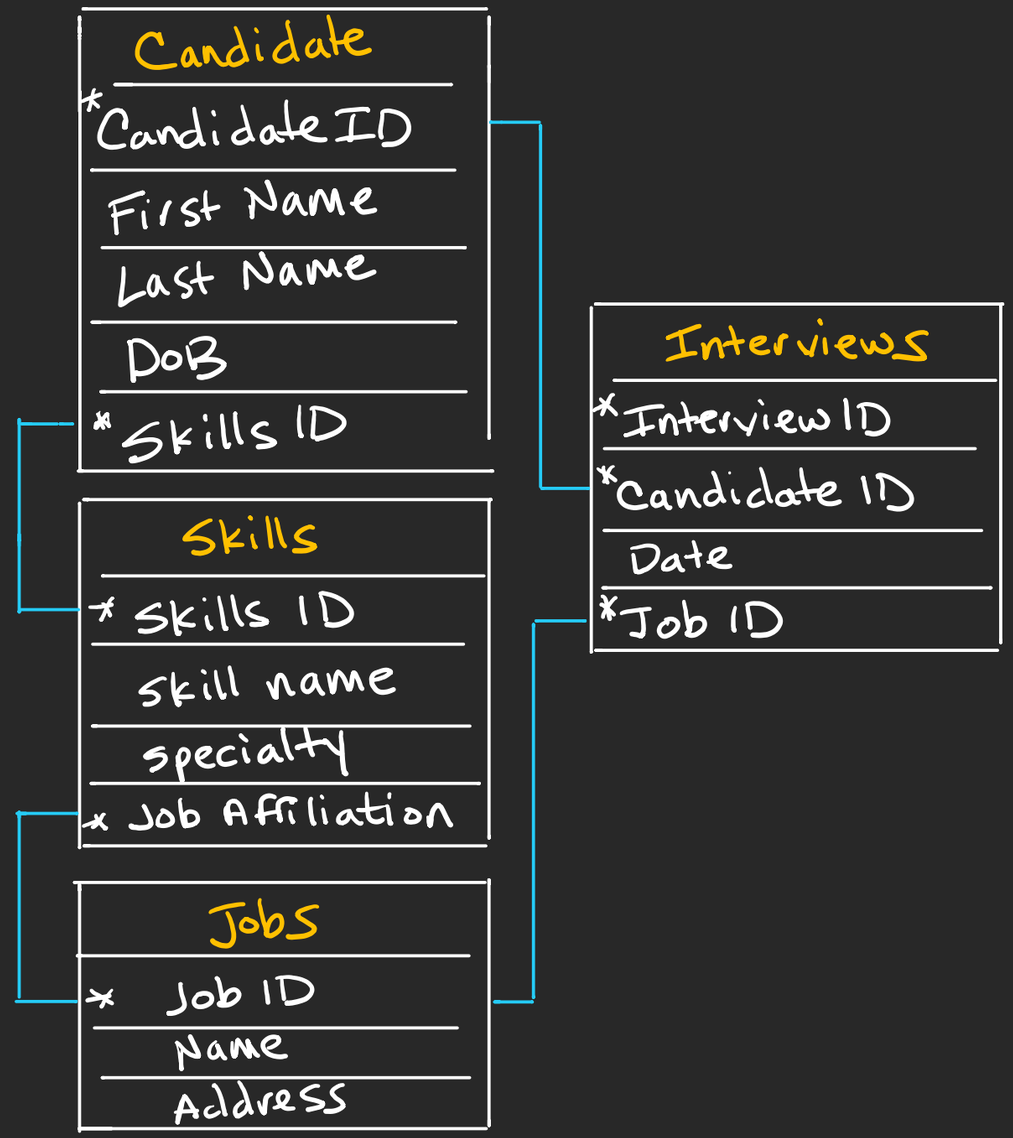

Document: Amazon DocumentDB


```
{
        "policyID": "0031852",
        "type": "car",
        "State": "Minnesota",
        "claims": [],
        "max coverage": 10000.00,
        "Address": {
                "Street": "22 Hawk Circle",
                "zip": "55701",
                "city": "Duluth",
                "state" : "MN"
                },
        "discount": ["good driver", "new
                customer",
        "student"],
        "createdAt": "2019-11-06T17:20:31Z"
}
```

In-memory: Amazon ElastiCache

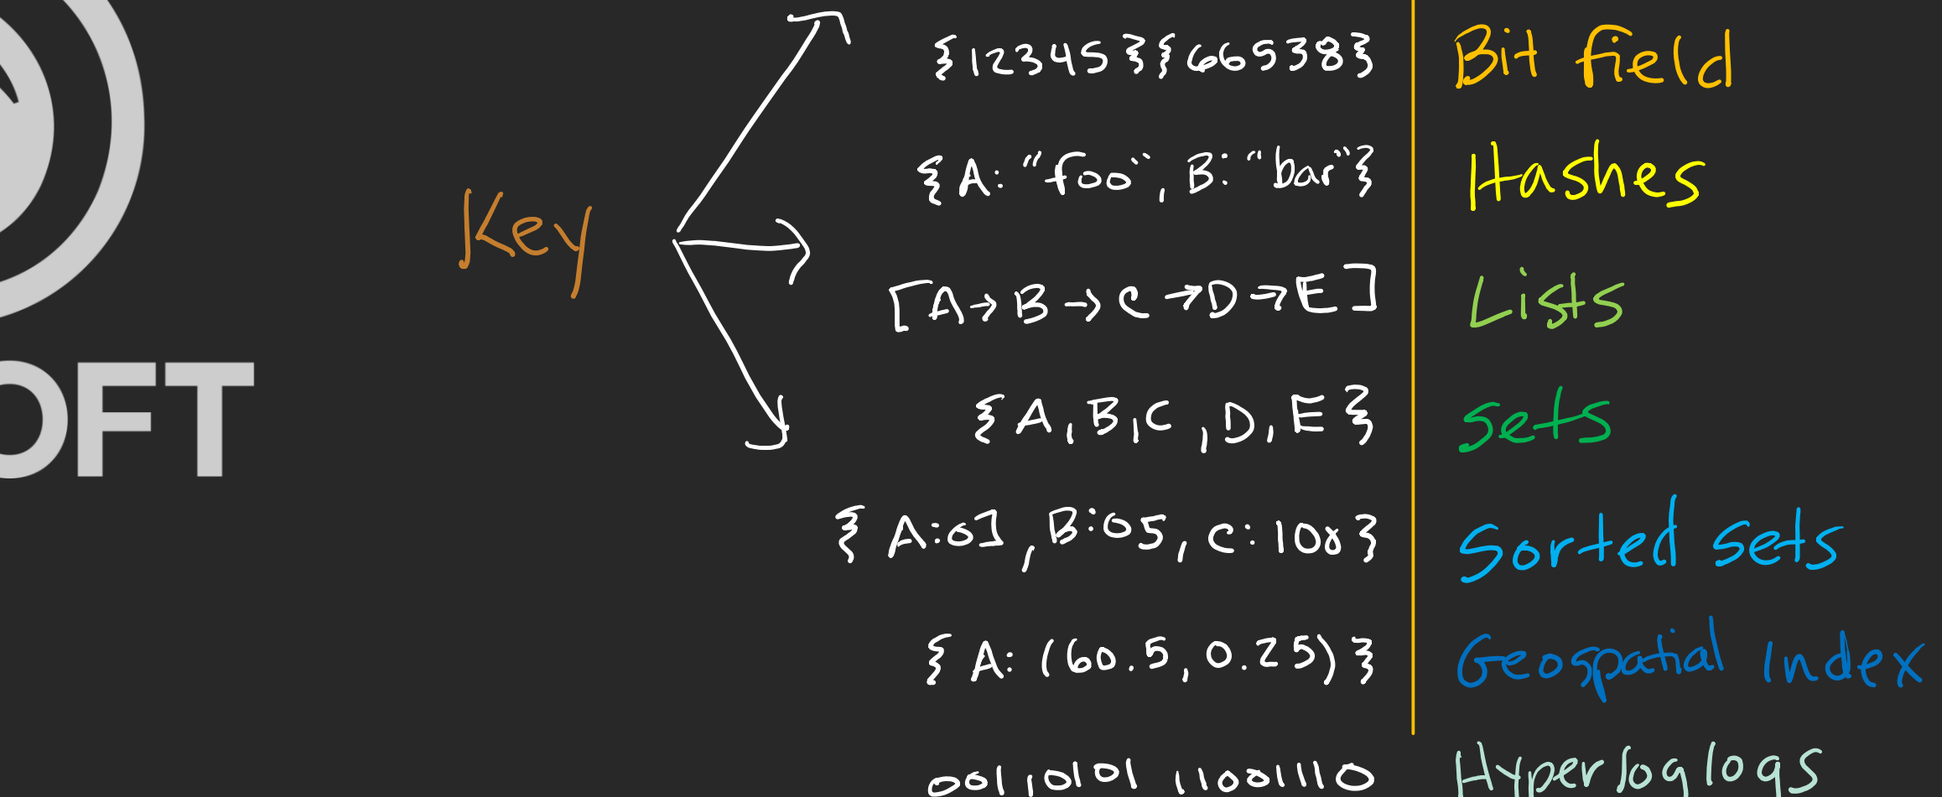

Graph: Amazon Neptune

Nike

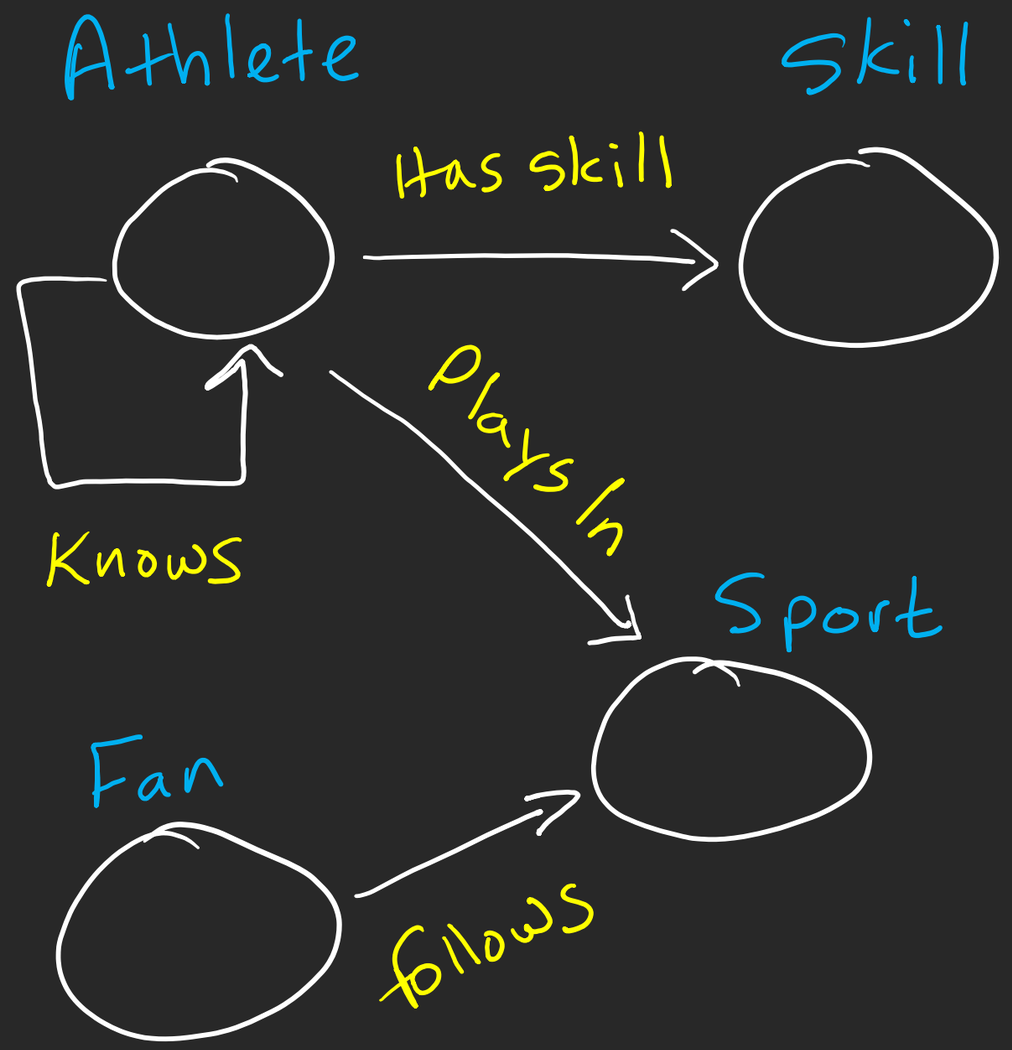

Ledger: Amazon QLDB

Klarna.

Smoooth payments

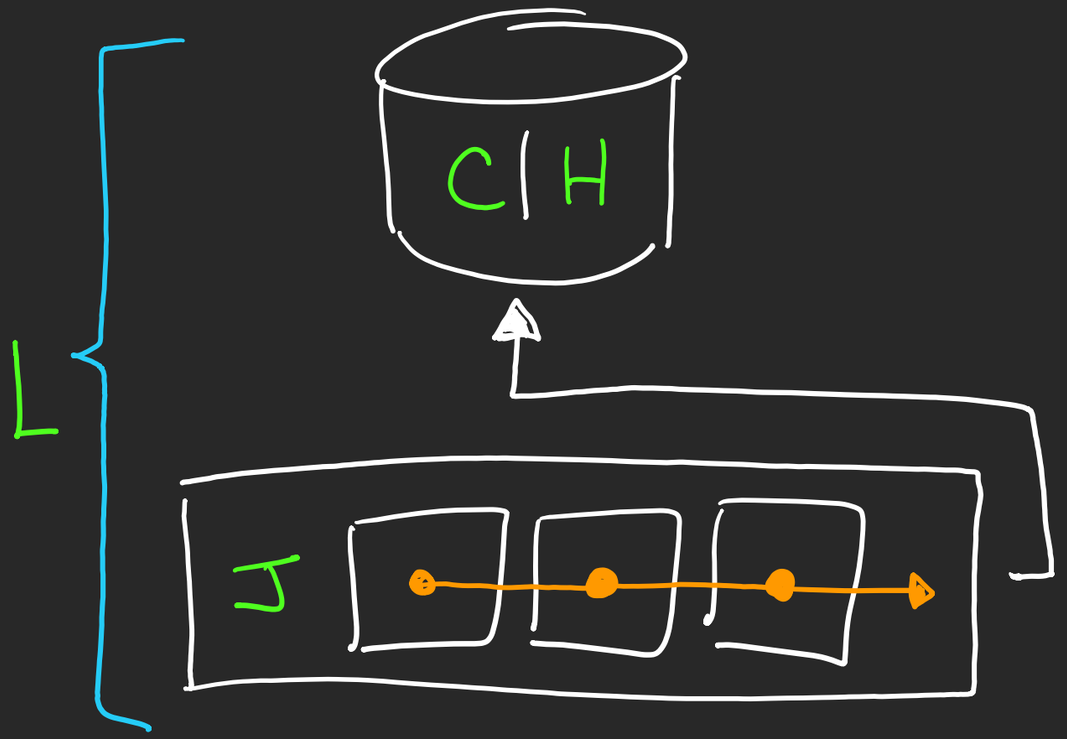

The right tool for the job

Woot

We started looking at how we could decouple our legacy Oracle database into smaller microservices, using the right tool for the right job.

The right tool for the job

Migrated their product data, images, and purchase orders from SQL Servers to DynamoDB, Amazon S3, AWS Lambda and Amazon ElastiCache. “We lowered costs by 20% and increased speed by 50% and migrated the whole system to the cloud in less than six months.”

Challenges querying data from multiple databases

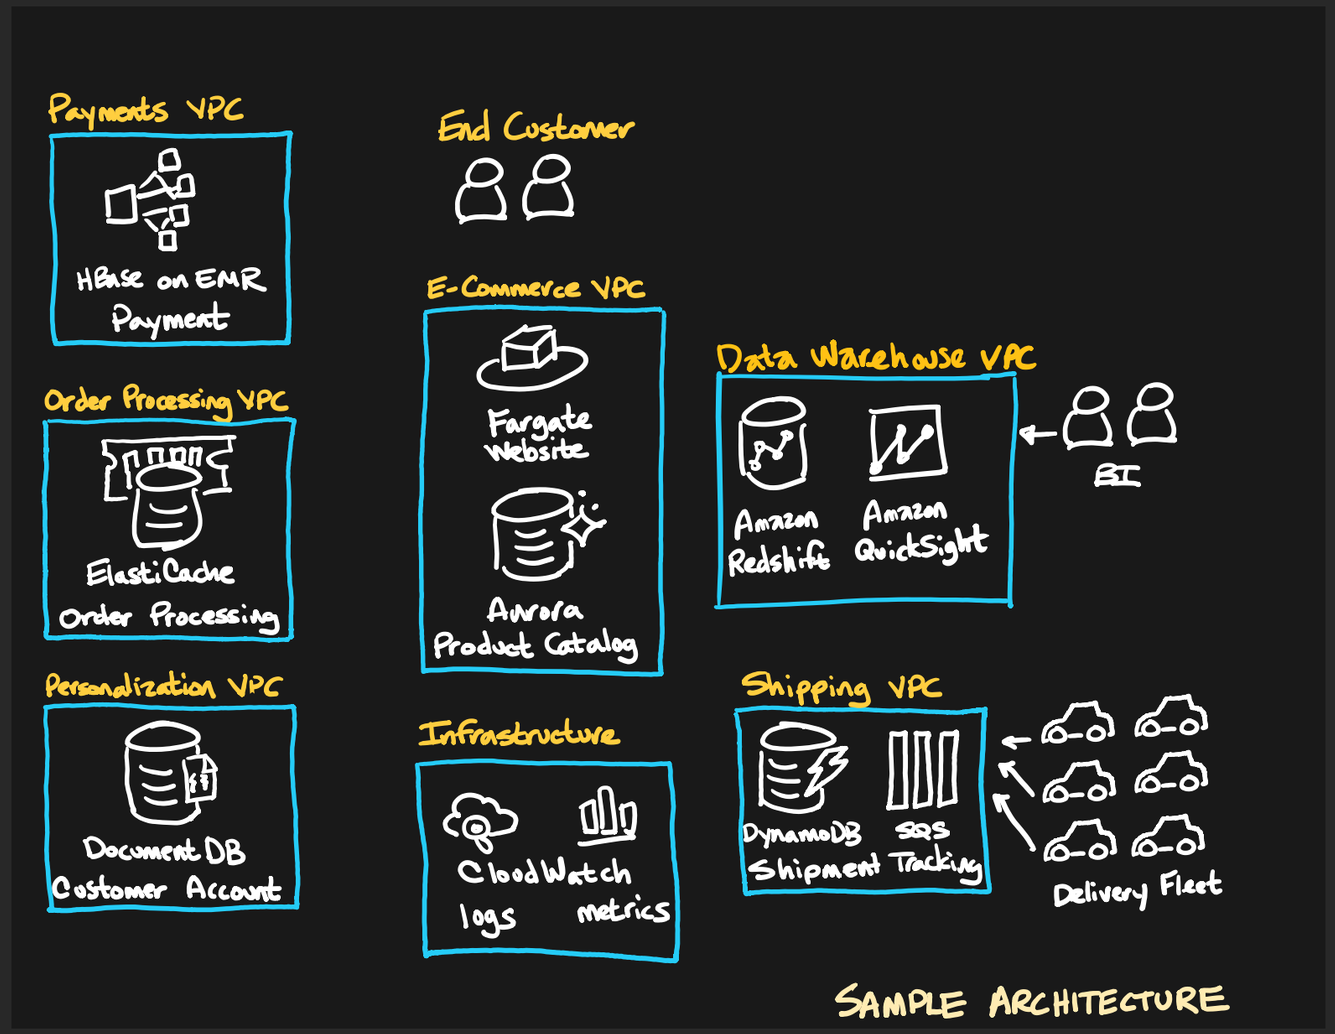

Imagine an e-commerce store with a microservices architecture

Accessing multiple systems can be challenging

Federated Query for Amazon Athena (Preview)

NEW!

Run SQL queries on data spanning multiple data stores

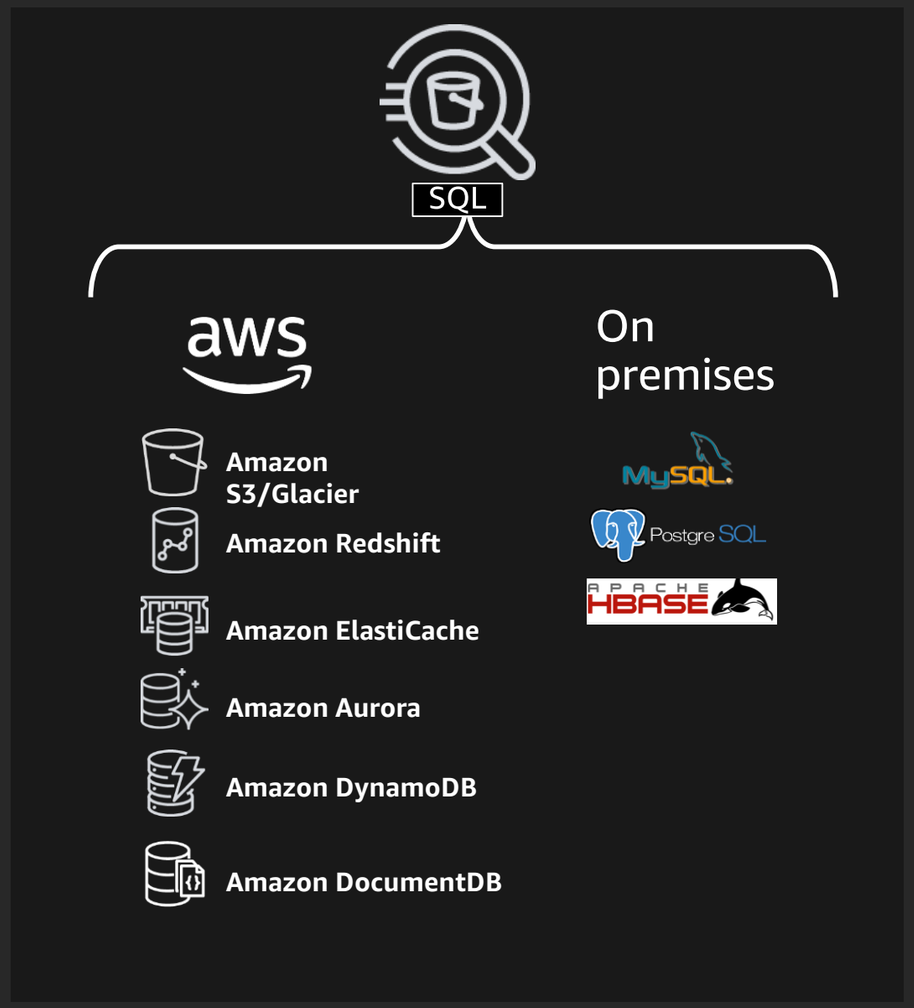

Run SQL queries on relational, non-relational, object, or custom data sources; in the cloud or on premises

Open-source connectors for common data sources

Build connectors to custom data sources

Run connectors in AWS Lambda: No servers to manage

Demo

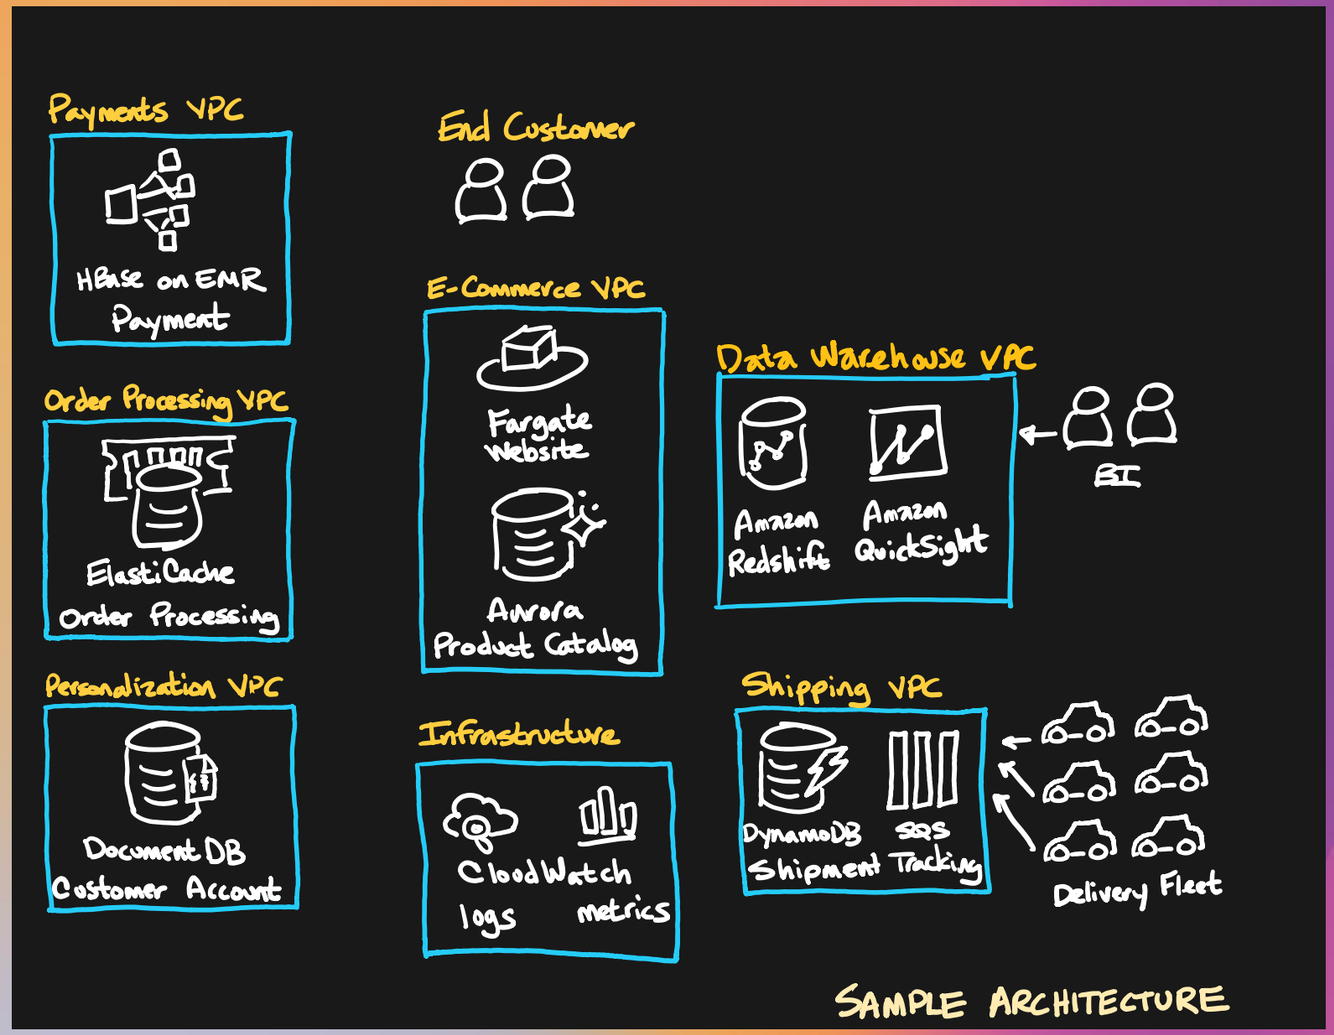

Challenges with integrating machine learning (ML) with your database

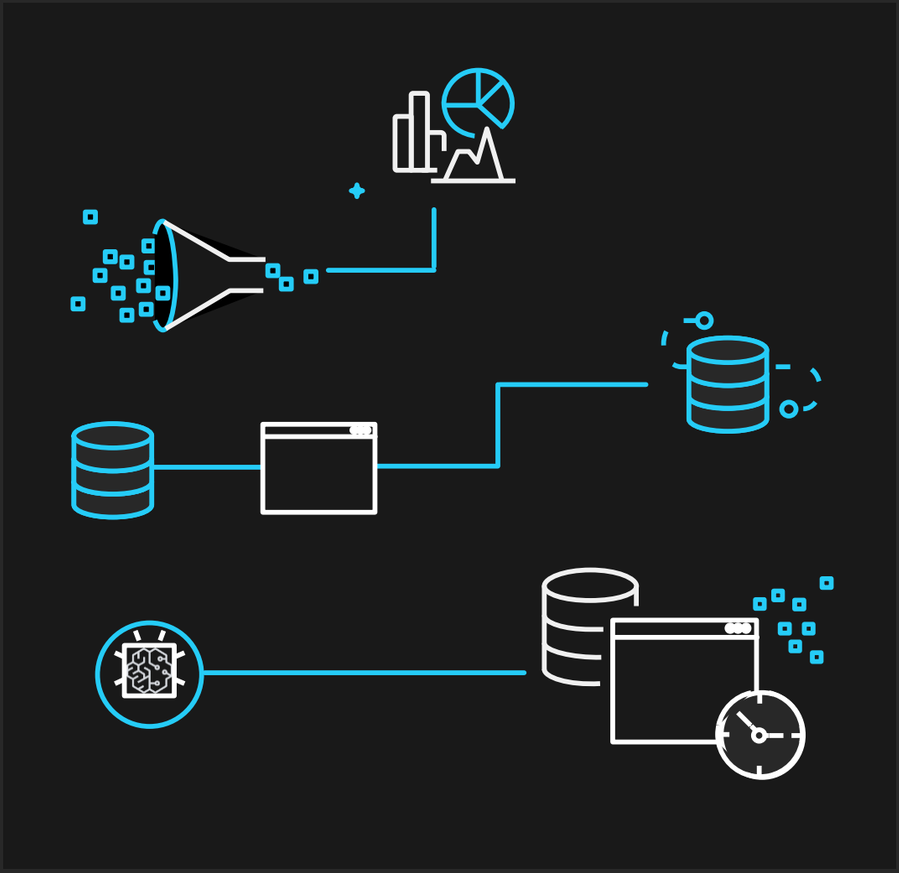

Select and train the model

Create application code to read data from the database

Query and format the data for the ML algorithm

Call an ML service to run the algorithm

Format the output

Amazon Aurora integration with ML

Simple, optimized, and secure Aurora, Amazon SageMaker, and Amazon Comprehend (in preview) integration

ML predictions on relational data

Integration with Amazon SageMaker & Amazon Comprehend

Familiar SQL language, no ML expertise

Low-latency, real-time

Security & governance

Demo

Challenges with serverless application development

Serverless applications open and close tens of thousands of connections within seconds

Leads to longer query response times

A database proxy server pools connections to an application, and communicates with the database through a smaller, more manageable number of connections

Proxy servers are difficult to deploy, patch, and manage

It’s challenging to manage large Cassandra clusters at scale

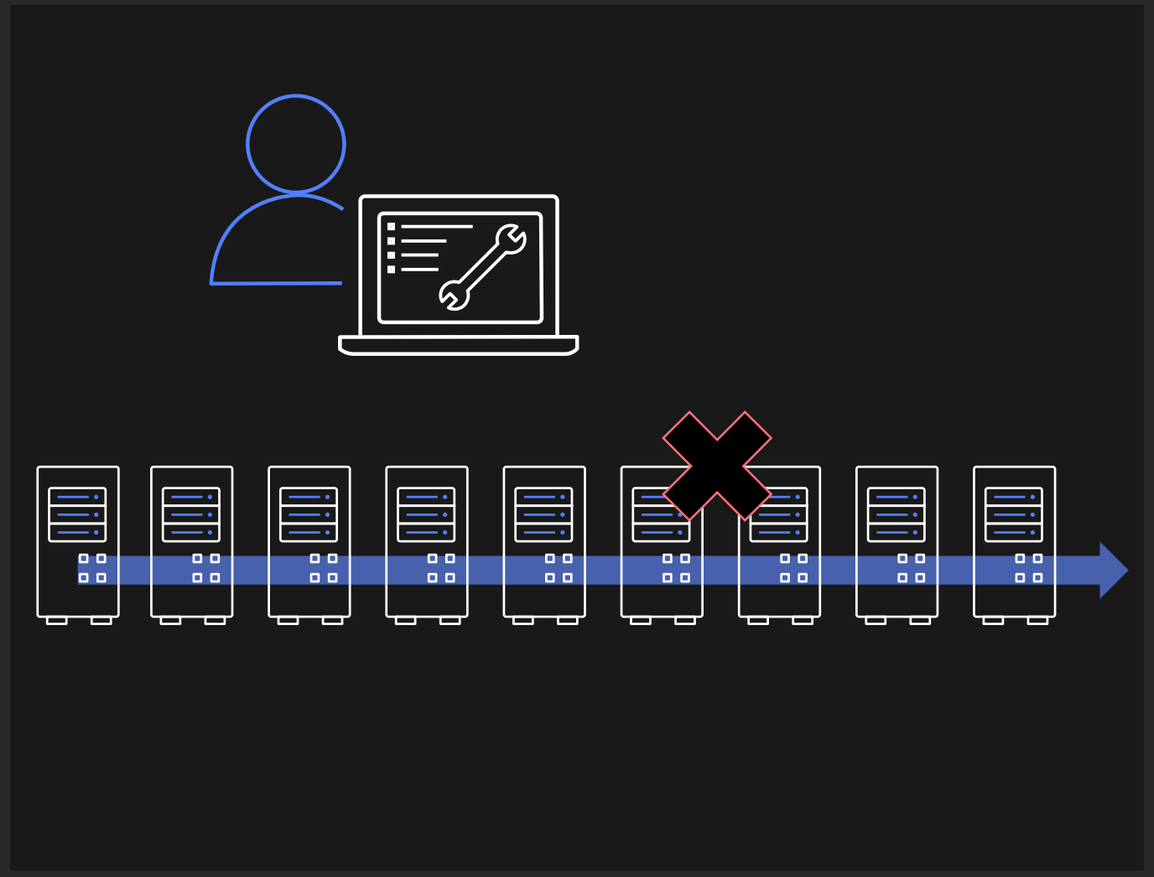

Specialized expertise to set up, configure, and maintain infrastructure and software

Scaling clusters is time-consuming, manual, and error-prone, so many overprovision capacity

Manual backups and error-prone restore processes to maintain integrity

Unreliable upgrades with clunky rollback and debugging capabilities

Demo

Related breakouts

DAT202-R – What's new in Amazon Aurora

DAT207-R – What's new in Amazon RDS

DAT220 – Real-world customer use cases with Amazon Neptune

DAT323-R – What's new with Amazon ElastiCache

DAT326 – Amazon DocumentDB deep dive

DAT380 – Amazon QLDB: An engineer's deep dive on why this is a game changer

DAT403-R – Amazon DynamoDB deep dive: Advanced design patterns

Learn databases with AWS Training and Certification

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

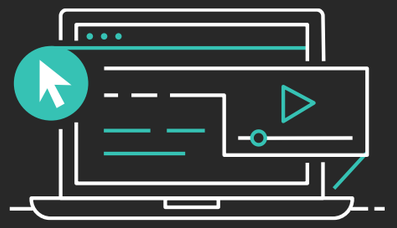

25+ free digital training courses cover topics and services related to databases, including:

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

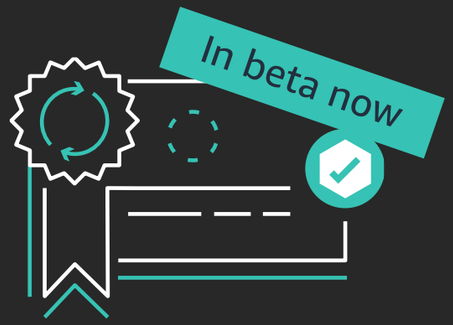

Validate expertise with the new **AWS Certified Database - Specialty** beta exam

Visit aws.training

aws training and certification

Thank you!

Shawn Bice
@shawnbice

Tobias Ternström
@TobiasSQL

Joseph Idziorek
@josephidziorek

AWS re:Invent

aws

Please complete the session survey in the mobile app.