
AWS
re:Invent

GPS303

Application modernization with AWS purpose-built databases

Ravi Madabhushanam

Sr. Solutions Architect, Databases

AWS

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Agenda

What is a modern application?

Why consider purpose-built databases?

AWS purpose-built databases

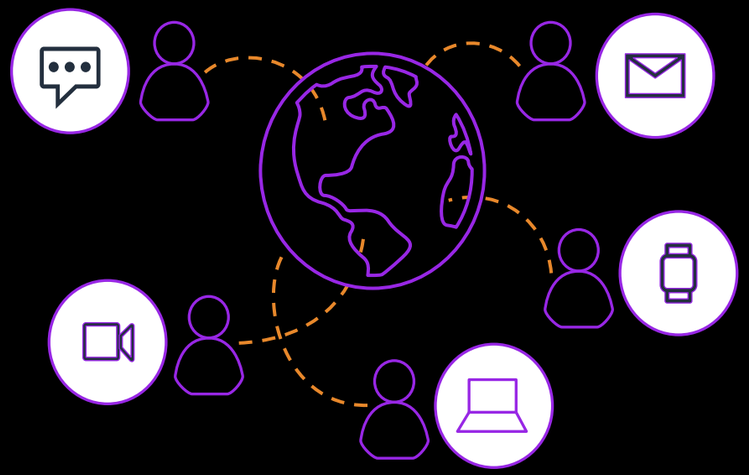
How to leverage AWS purpose-built databases



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

Modern application requirements

MORE PERFORMANCE, SCALE, AND AVAILABILITY

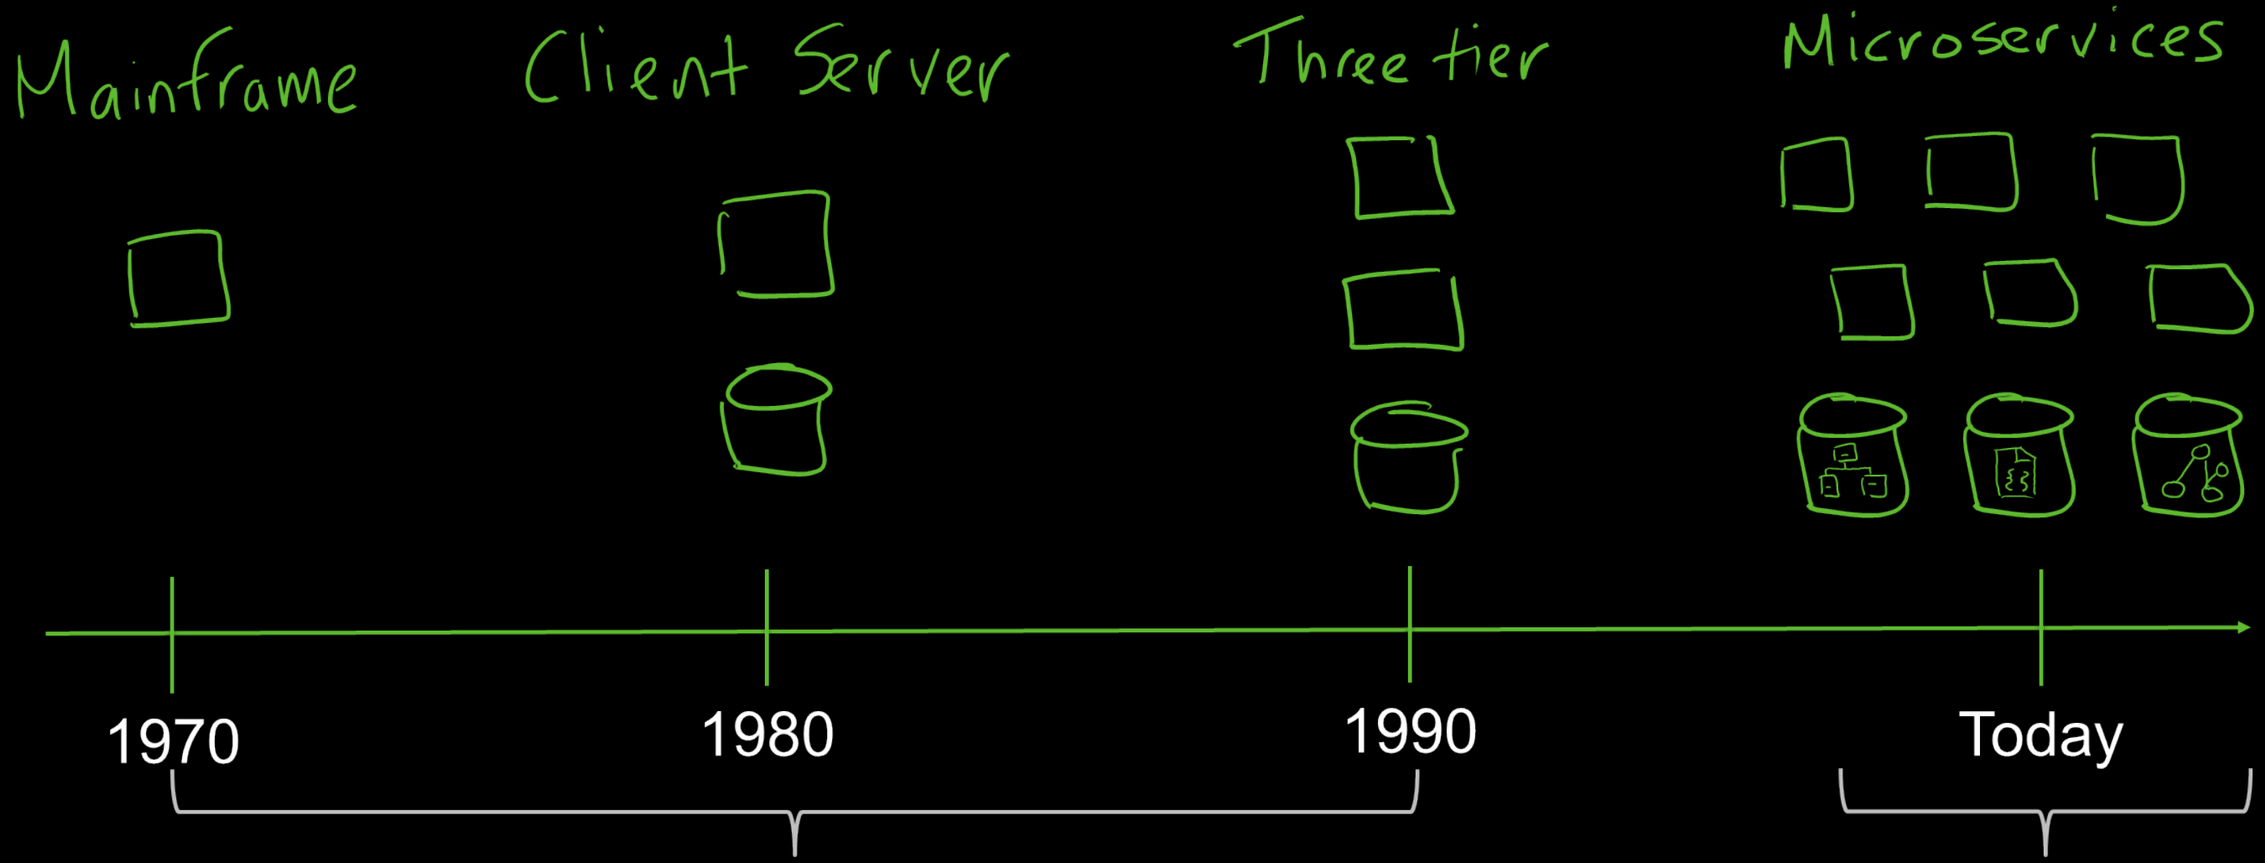




E-commerce | Media streaming | Social media | Online gaming | Shared economy

Users	1M+
Data volume	Terabytes to petabytes
Locality	Global
Performance	Microsecond latency
Request rate	Millions per second
Access	Mobile, IoT, devices
Scale	Virtually unlimited
Payment model	Pay as you go
Developer access	Instant API access
Development	Apps and storage are decoupled

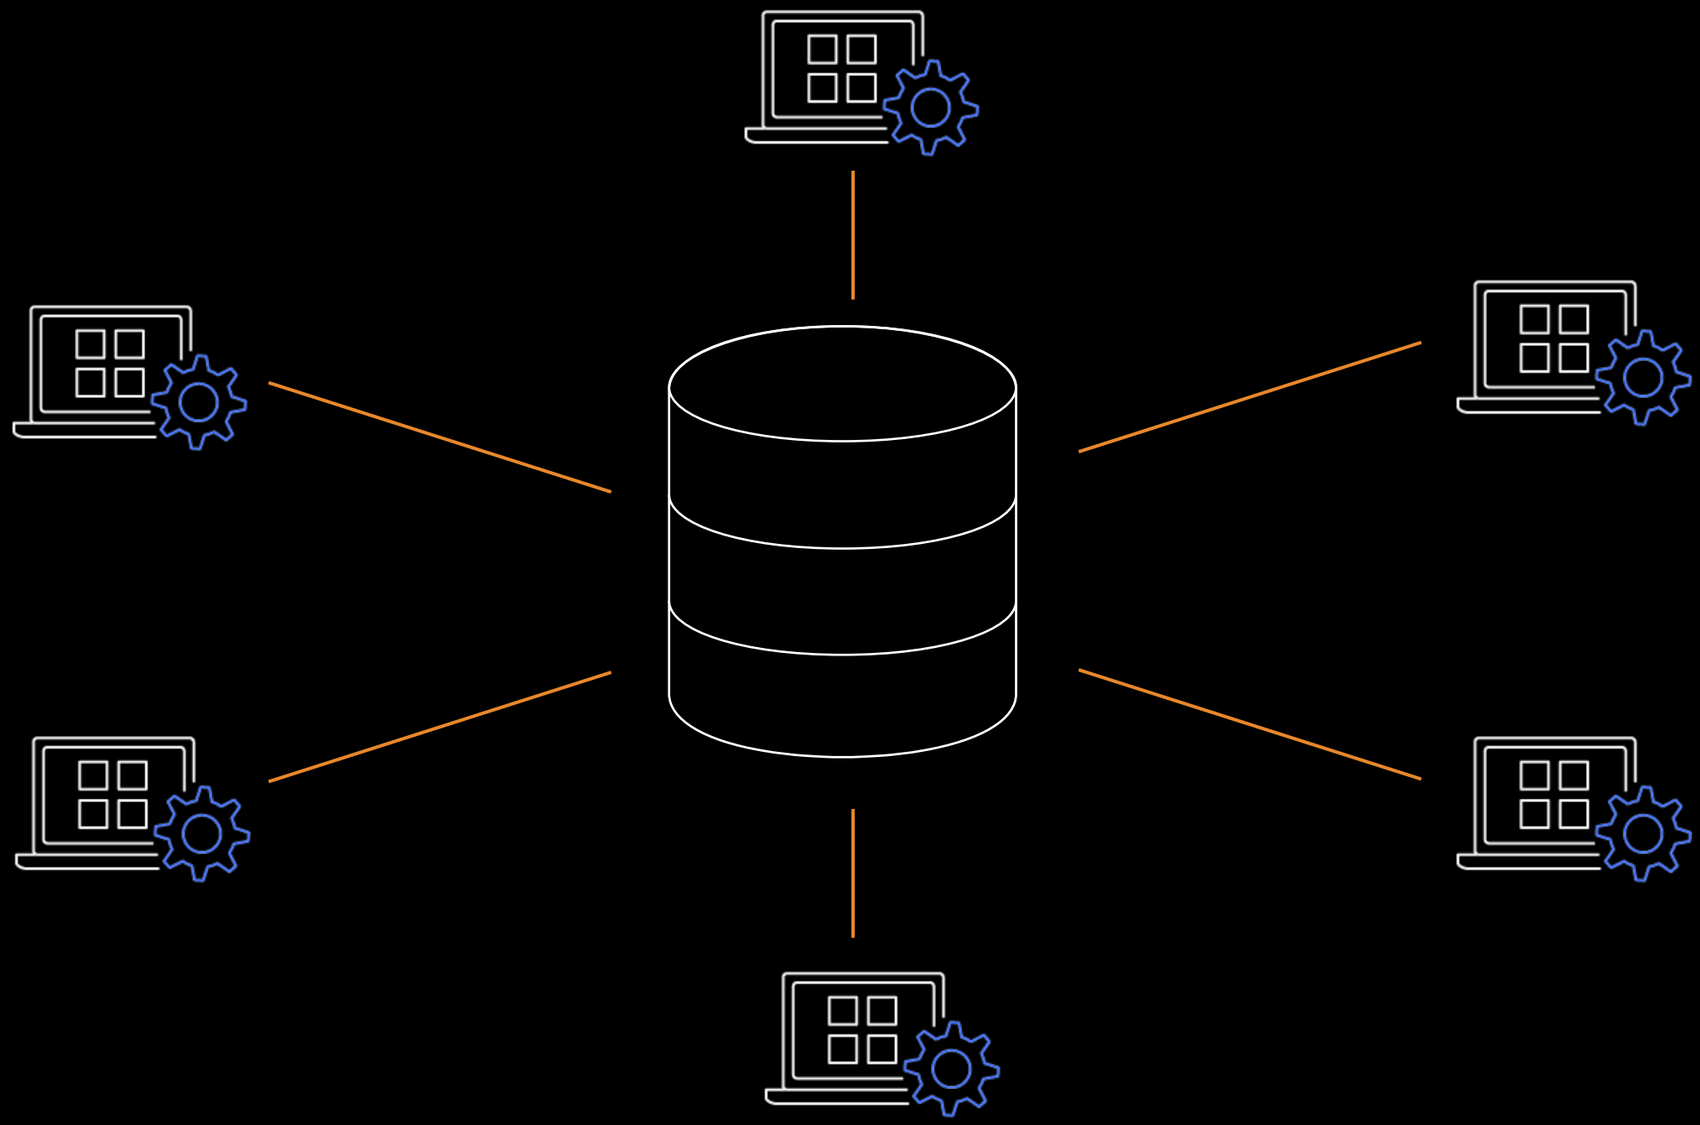
App architectures and patterns have evolved



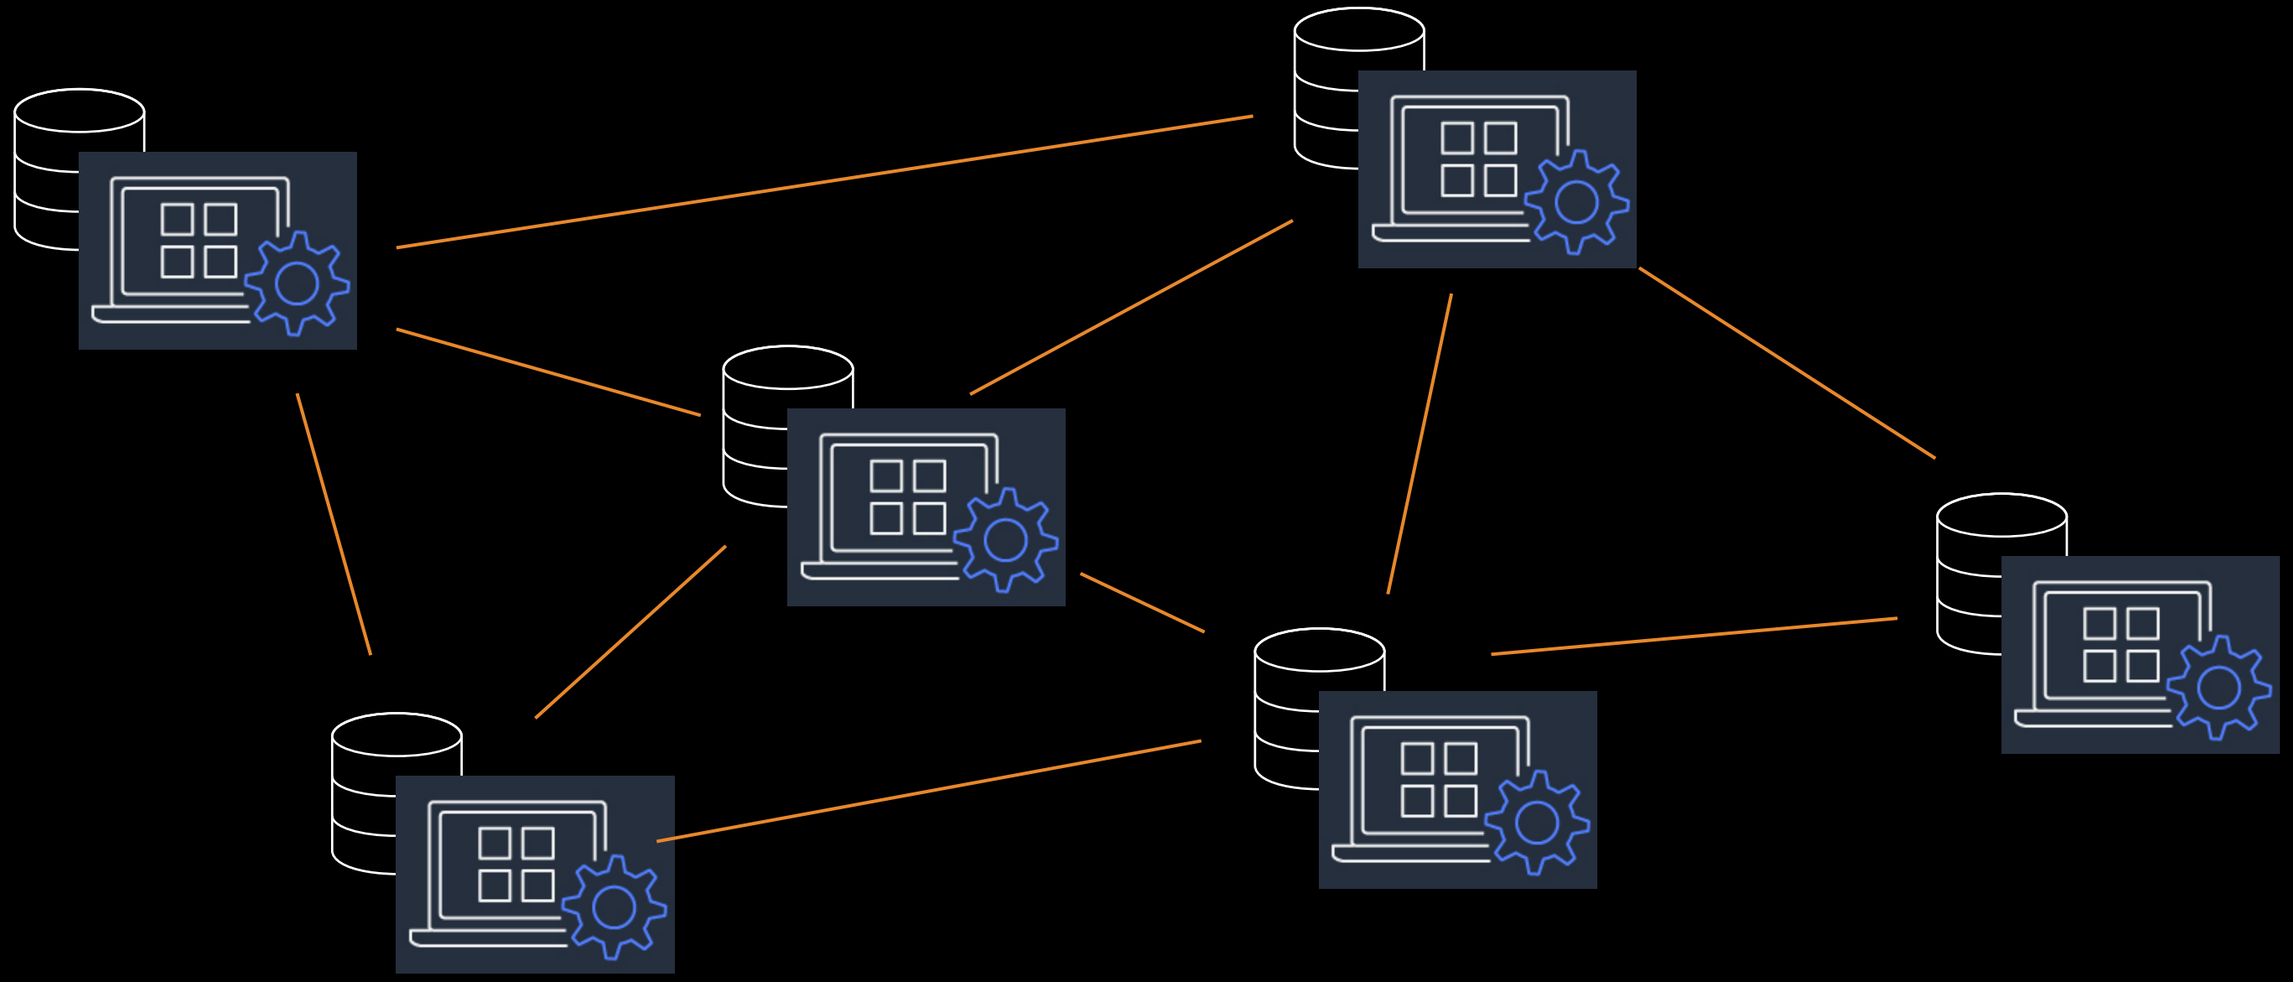
All data was modeled as relational, regardless of the shape or function of the data

Multiple databases and data models are used within the same application

Decoupling your data



Decoupling data into each service improves resilience

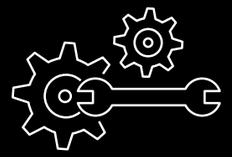


Purpose-built databases

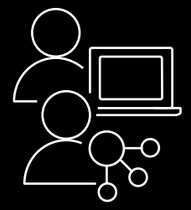
What are the data needs of each modular service?



Massive data volume?



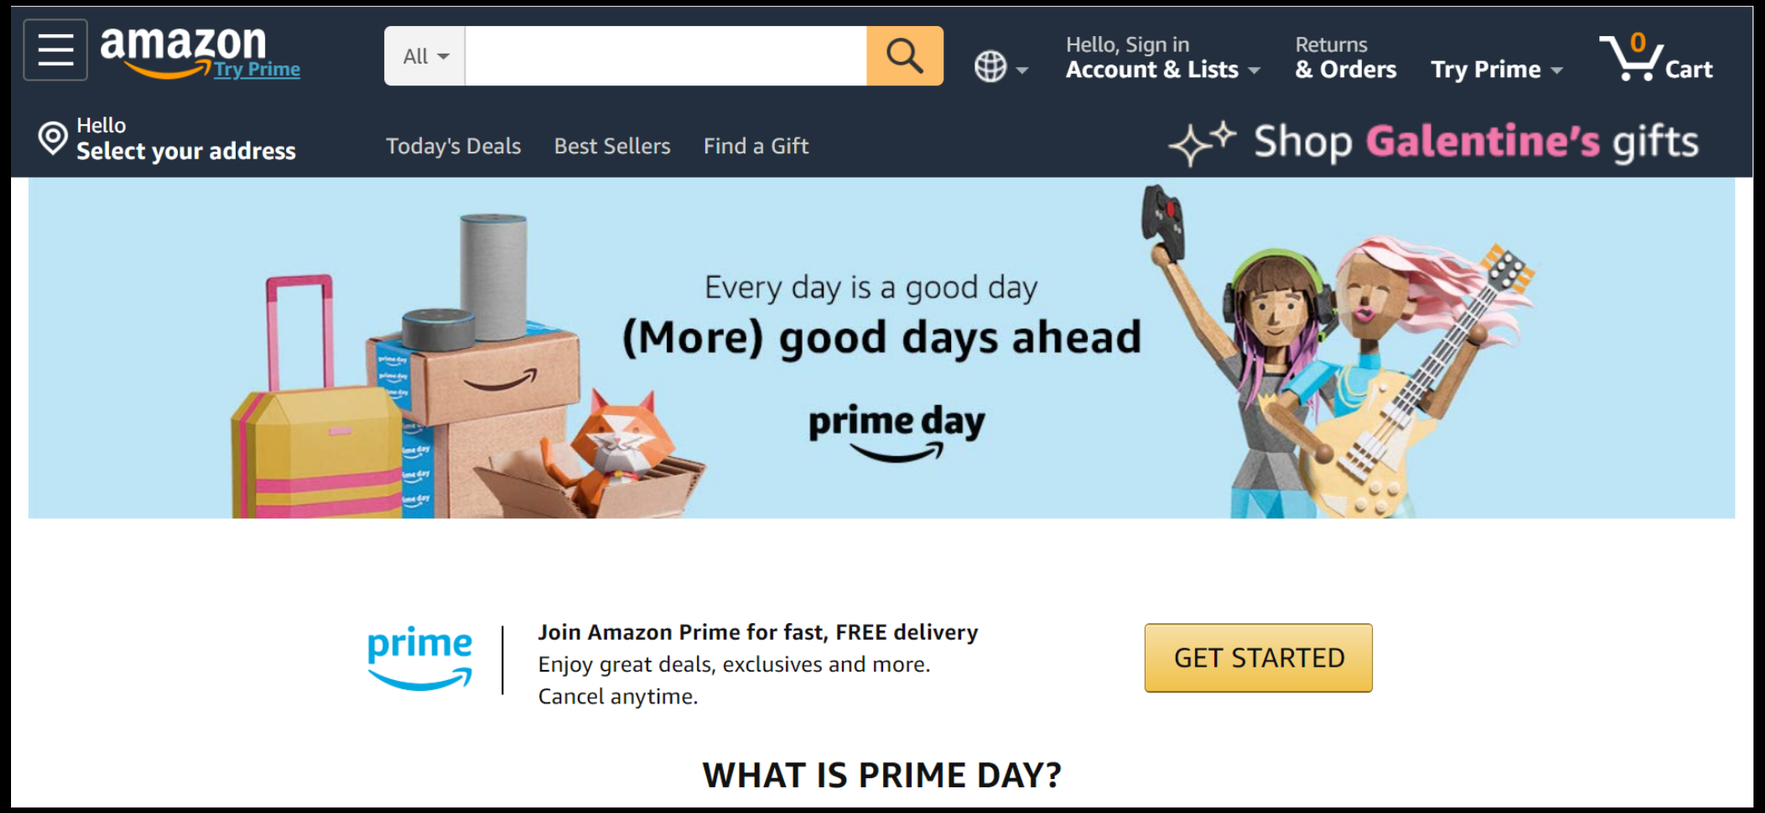
High-speed rendering?



Data warehousing?

One size does not fit all

Internet-scale e-commerce



The world's largest e-commerce business, Amazon.com, runs on **purpose-built databases** because of their **scale**, **performance**, and **maintenance benefits**

AWS delivers the right database for the right job

	Relational	Key value	Document	In memory	Graph	Time series	Ledger	Wide column
	Referential integrity, ACID transactions, schema-on-write	High-throughput, low-latency reads and writes; endless scale	Store documents and quickly access querying on any attribute	Query by key with microsecond latency	Quickly and easily create and navigate relationships between data	Collect, store, and process data sequenced by time	Complete, immutable, and verifiable history of all changes to application data	Scalable, highly available, and managed Apache Cassandra–compatible service
AWS service(s)	RDS Aurora Redshift	DynamoDB	DocumentDB	ElastiCache	Neptune	Timestream	QLDB	Keyspaces Managed Cassandra
Common use cases	Lift-and-shift, ERP, CRM, finance	Real-time bidding, shopping cart, social, product catalog, customer preferences	Content management, personalization, mobile	Leaderboards, real-time analytics, caching	Fraud detection, social networking, recommendation engine	IoT applications, event tracking	Systems of record, supply chain, health care, registrations, financial	Build low-latency applications, leverage open source, migrate Cassandra to the cloud

Database strategy for modern applications



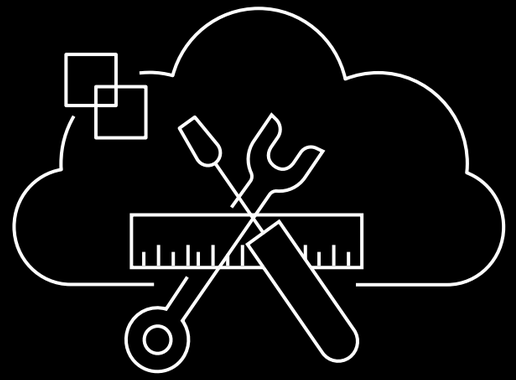
AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

2 fundamental areas of focus



Lift-and-shift existing apps to the cloud



Quickly build new apps in the cloud

Lift-and-shift existing apps to the cloud



Lift-and-shift existing apps to the cloud



Quickly build new apps in the cloud

Amazon Relational Database Service (Amazon RDS)

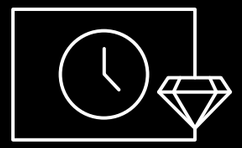
MANAGED RELATIONAL DATABASE SERVICE WITH A CHOICE OF 6 POPULAR DATABASE ENGINES

Amazon Aurora | MySQL | PostgreSQL | MariaDB | Microsoft SQL Server | ORACLE



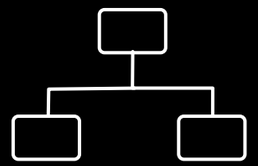
Easy to administer

No need for infrastructure provisioning, or installing and maintaining DB software



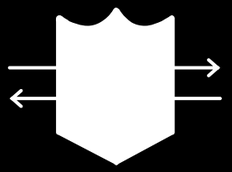
Availability and durability

Automatic multi-AZ data replication; automated backup, snapshots, and failover



Highly scalable

Scale database compute and storage with a few clicks, with no app downtime



Fast and secure

SSD storage and guaranteed provisioned I/O; data encryption at rest and in transit

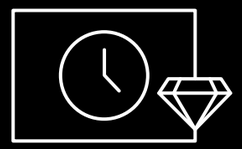
Amazon Aurora

MYSQL- AND POSTGRESQL-COMPATIBLE RELATIONAL DATABASE BUILT FOR THE CLOUD



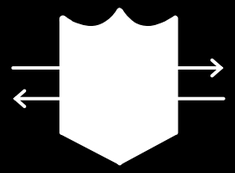
Performance and scalability

5x throughput of standard MySQL and 3x of standard PostgreSQL; scale out up to 15 read replicas



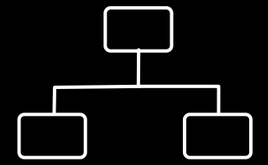
Availability and durability

Fault-tolerant, self-healing storage; 6 copies of data across 3 AZs; continuous backup to Amazon Simple Storage Service (Amazon S3)



Highly secure

Network isolation and encryption at rest / in transit



Fully managed

Managed by Amazon Relational Database Service (Amazon RDS) – On your part, no server provisioning, software patching, setup, configuration, or backup

AWS Database Migration Service (AWS DMS)

MIGRATING DATABASES TO AWS

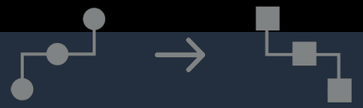
100,000+
databases migrated



Migrate between on premises and AWS



Migrate between databases



Automated schema conversion



Data replication for
zero-downtime migration

Target applications for move to managed / replatform



Is self-managing databases on premises, in Amazon Elastic Compute Cloud (Amazon EC2), and / or in another public cloud



Wants to reduce DB admin burden and reallocate DBA resources to app-centric work



Does not want to rearchitect applications

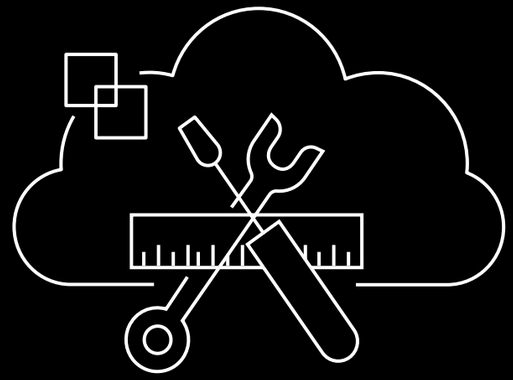


Needs a simple path to a managed service in the cloud for DB workloads

Quickly build new apps in the cloud

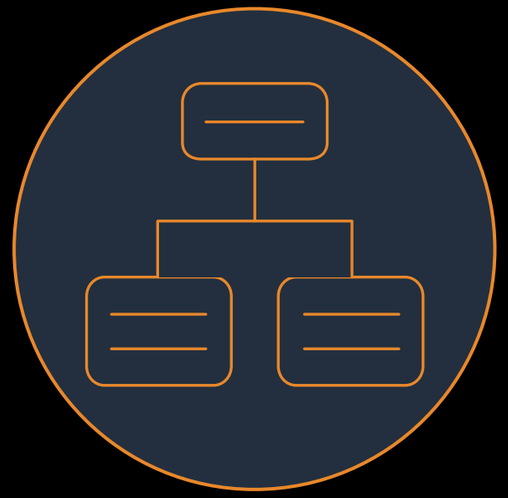


Lift-and-shift existing apps to the cloud

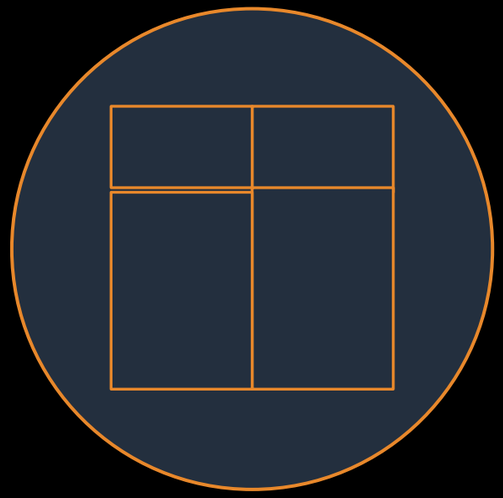


Quickly build new apps in the cloud

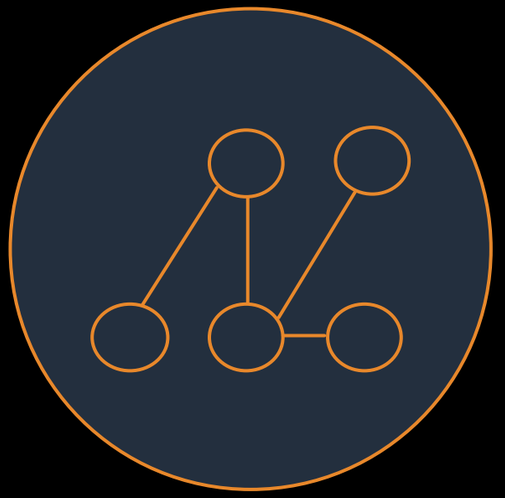
Let’s take a closer look at the data model



Relational



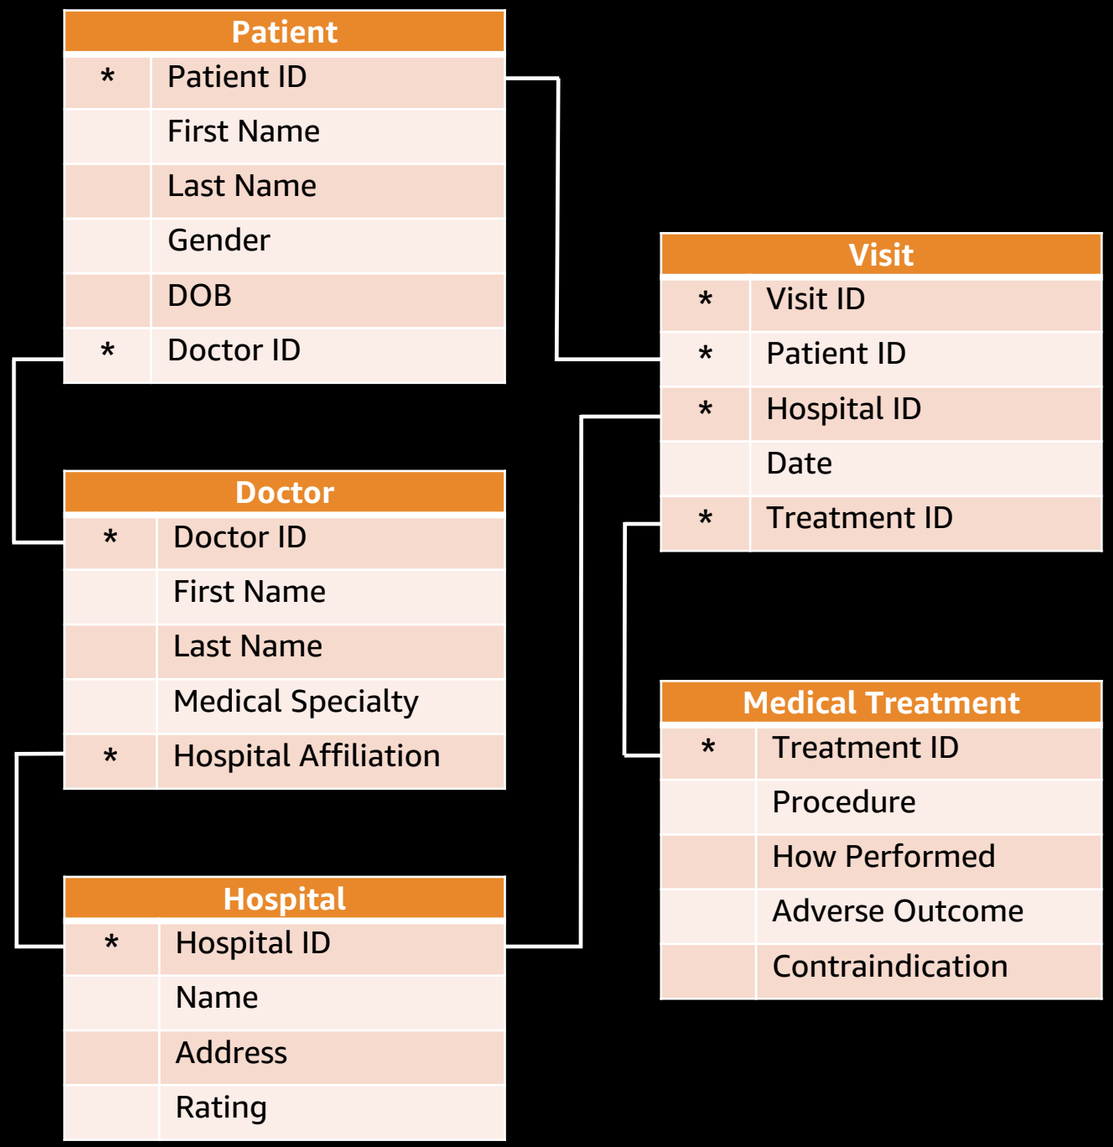
Key value



Graph

Relational data

- Divide data among tables
- Highly structured
- Relationships established via keys enforced by the system
- Data accuracy and consistency



Key-value data

- Simple key value pairs
- Partitioned by keys
- Resilient to failure
- High-throughput, low-latency reads and writes
- **Consistent performance at scale**

```
PUT {
  TableName:"Gamers",
   Item: {
  "GamerTag":"Hammer57",
  "Level":21,
  "Points":4050,
  "Score":483610,
  "Plays":1722
    } }


GET {
  TableName:"Gamers",
  Key: {
    "GamerTag":"Hammer57“,
    “ProjectionExpression“:”Points”
    } }
```

Gamers				
Primary Key	**Attributes**			
GamerTag	**Level**	**Points**	**High Score**	**Plays**
Hammer57	21	4050	483610	1722
FluffyDuffy	5	1123	10863	43
Lol777313	14	3075	380500	1307
Jam22Jam	20	3986	478658	1694
ButterZZ_55	7	1530	12547	66
…	…	…	…	…

Amazon DynamoDB

FAST AND FLEXIBLE KEY-VALUE DATABASE SERVICE FOR ANY SCALE



Performance at scale

Consistent single-digit millisecond response times at any scale; build applications with virtually unlimited throughput



Serverless architecture

No hardware provisioning, software patching, or upgrades; scales up or down automatically; continuously backs up your data



Enterprise security

Encrypts all data by default and fully integrates with AWS Identity and Access Management (IAM) for robust security

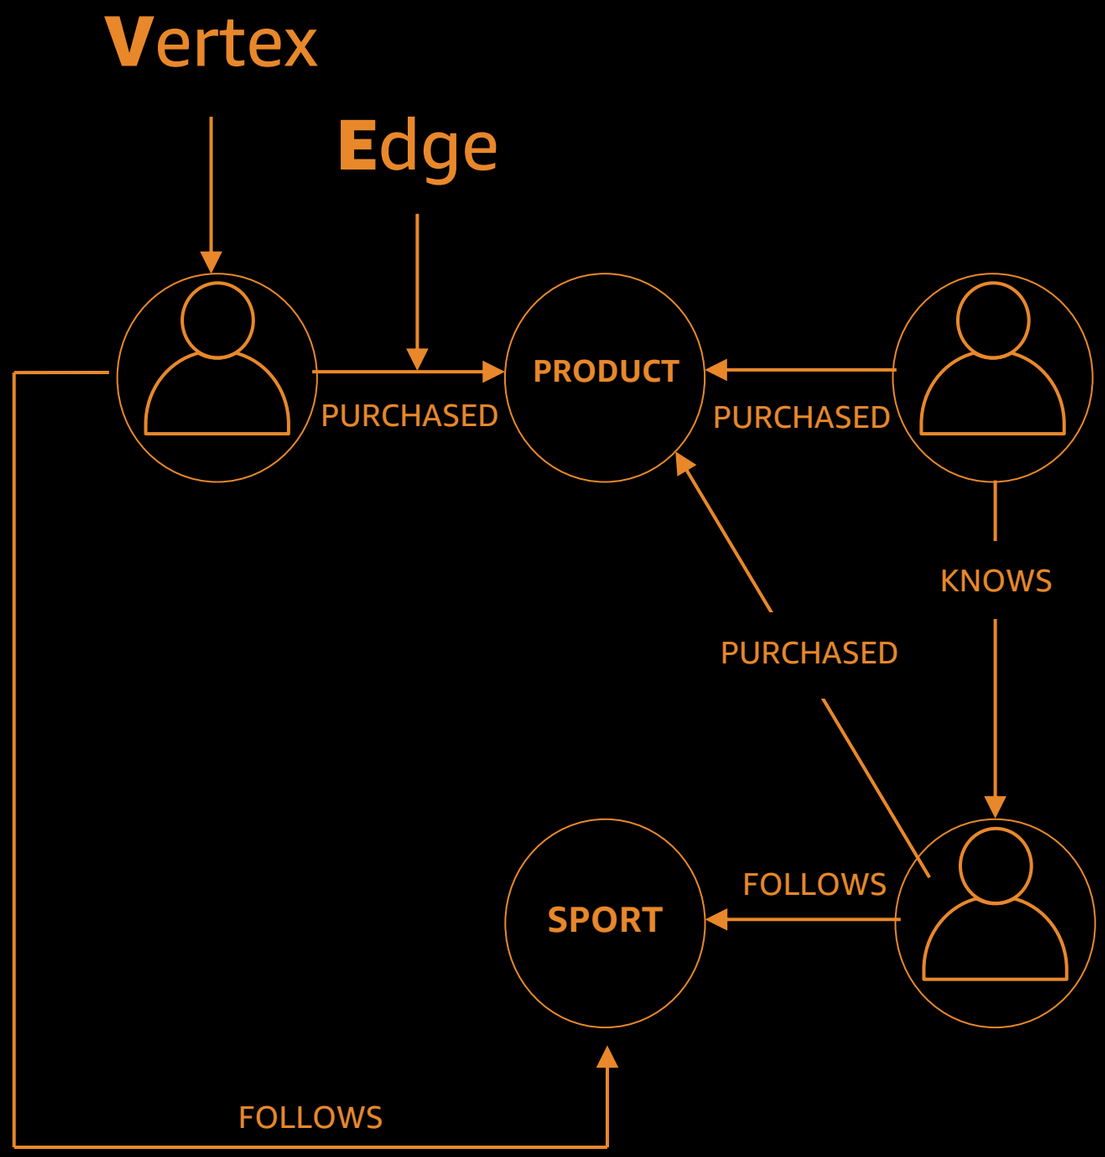


Global replication

Build global applications with fast access to local data by easily replicating tables across multiple AWS Regions

Graph data

- Relationships are first-class objects
- Vertices connected by Edges



Amazon Neptune

FAST, RELIABLE GRAPH DATABASE BUILT FOR THE CLOUD

Open



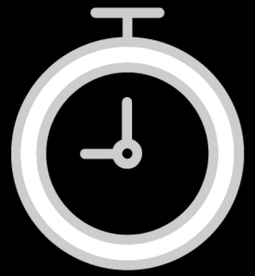
Supports Apache TinkerPop and W3C RDF graph models

Fast



Queries billions of relationships with millisecond latency

Reliable



6 replicas of your data across 3 AZs with full backup and restore

Easy



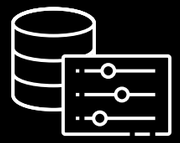
Build powerful queries easily with Gremlin and SPARQL

Amazon DocumentDB (with MongoDB compatibility)

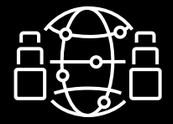
FAST, SCALABLE, HIGHLY AVAILABLE DATABASE SERVICE



Millions of requests per second; millisecond latency



Same code, drivers, and tools you use with MongoDB



Simple and fully managed



Secure and compliant



2x throughput of managed MongoDB services



Deeply integrated with AWS services

Amazon ElastiCache

MANAGED, REDIS, OR MEMCACHED-COMPATIBLE IN-MEMORY DATA STORE



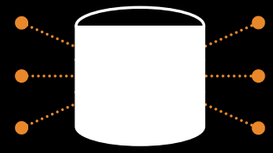
Unlimited scale

Read scaling with replicas; write and memory scaling with sharding; nondisruptive scaling



Consistent high performance

In-memory data store and cache for submillisecond response times



Fully managed

AWS manages all hardware and software setup, configuration, and monitoring

Amazon Timestream

FAST, SCALABLE, FULLY MANAGED TIME SERIES DATABASE

1,000x faster and 1/10th the cost of relational databases



Collect data at the rate of 10 million inserts per second

Trillions of daily events



Adaptive query processing engine maintains steady, predictable performance

Time series analytics



Built-in functions for interpolation, smoothing, and approximation

Serverless



Automated setup, configuration, server provisioning, and software patching

Amazon Quantum Ledger Database (Amazon QLDB)

FULLY MANAGED LEDGER DATABASE: TRACK AND VERIFY HISTORY OF ALL CHANGES MADE TO YOUR APPLICATION'S DATA

Immutable and transparent



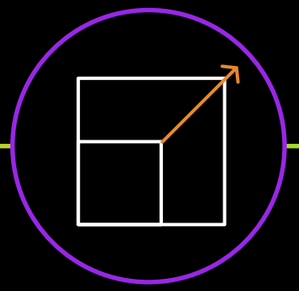
Append-only, immutable journal tracks history of all changes that cannot be deleted or modified; get full visibility into entire data lineage

Cryptographically verifiable



All changes are cryptographically chained and verifiable

Highly scalable



Executes 2–3x as many transactions as ledgers in common blockchain frameworks

Easy to use



Flexible document model; query with familiar SQL-like interface

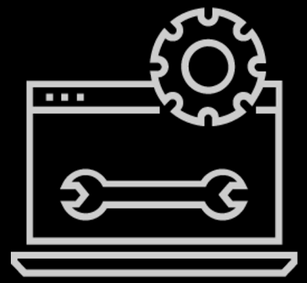
Amazon Keyspaces (for Apache Cassandra)

FAST, RELIABLE, WIDE-COLUMN DATABASE BUILT FOR THE CLOUD



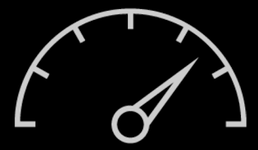
Apache Cassandra–compatible

Implements the Apache Cassandra Query Language (CQL) and the Apache Cassandra CQL API



No servers to manage

No need to provision, patch, or manage servers



Performance at scale

Consistent single-digit millisecond response times at any scale



Highly available and secure

Tables are encrypted by default and replicated 3 times in multiple AWS Availability Zones for high availability

How to leverage AWS purpose-built databases: Partners' perspective



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

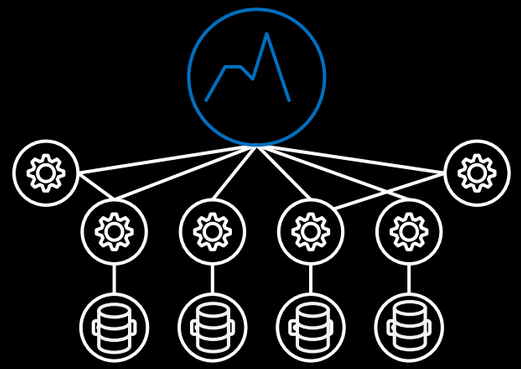
Rapid expansion of data requirements for modern workloads

Explosion of data



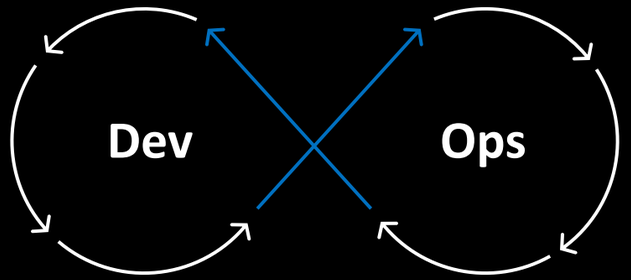
Data grows 10x every 5 years, driven by network-connected smart devices

Microservices changes data and analytics requirements



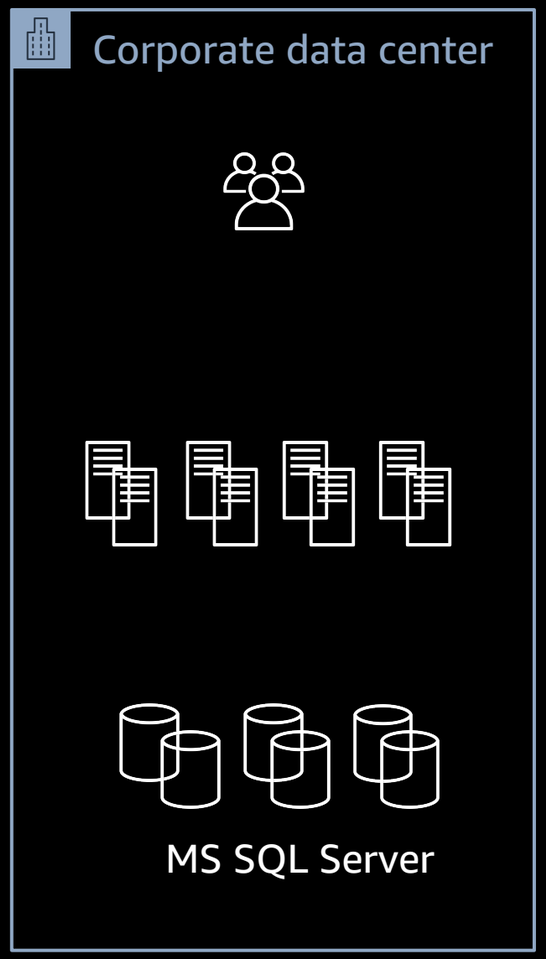
Microservices architecture decreases need for one-size-fits-all databases and increases need for real-time monitoring and analytics

Accelerated rate of change driven by DevOps



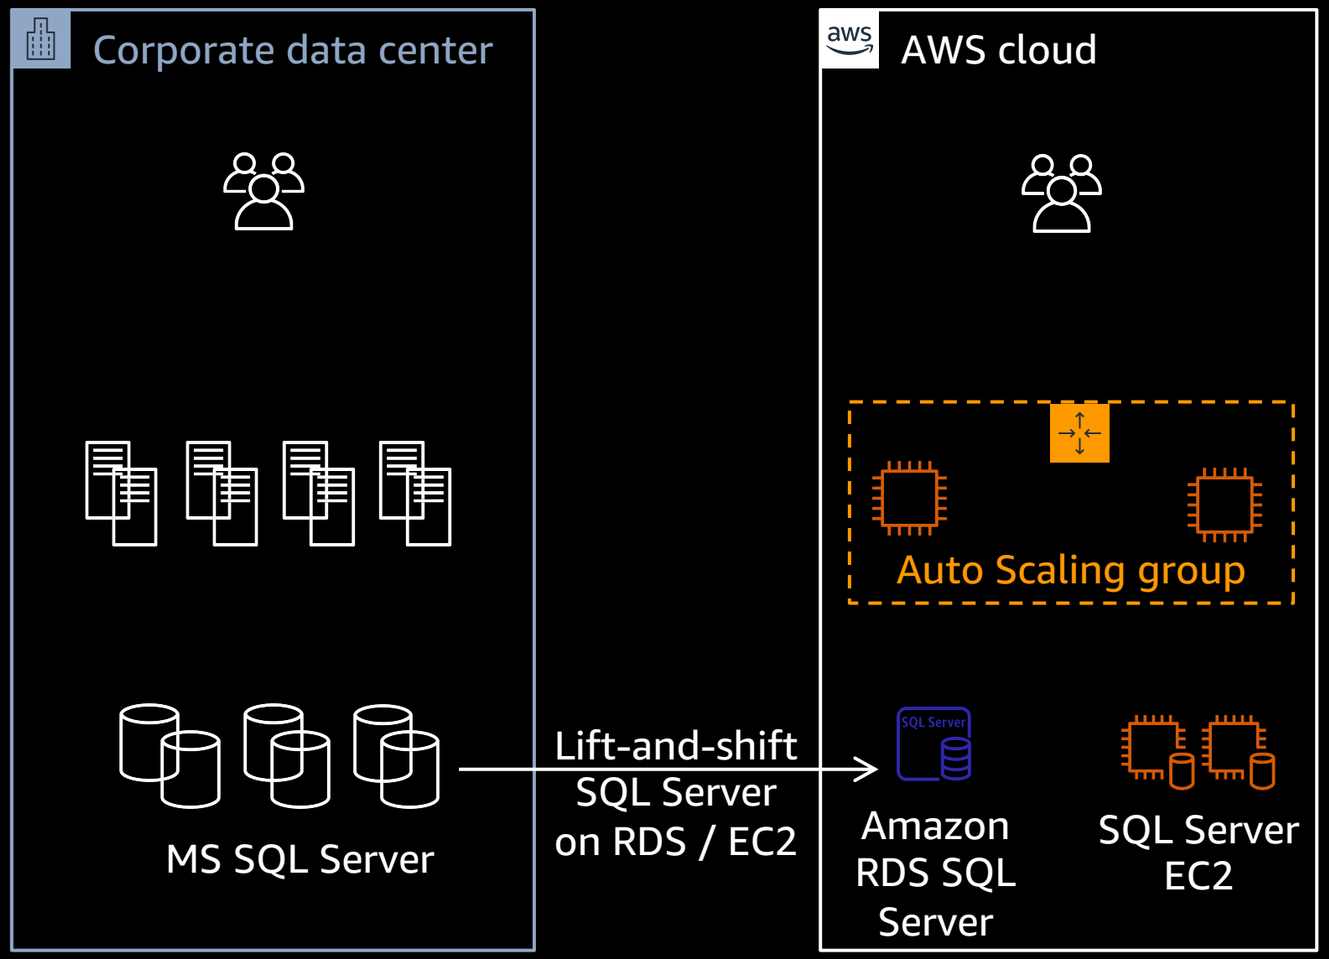
Transition from IT to DevOps increases rate of change

User journey



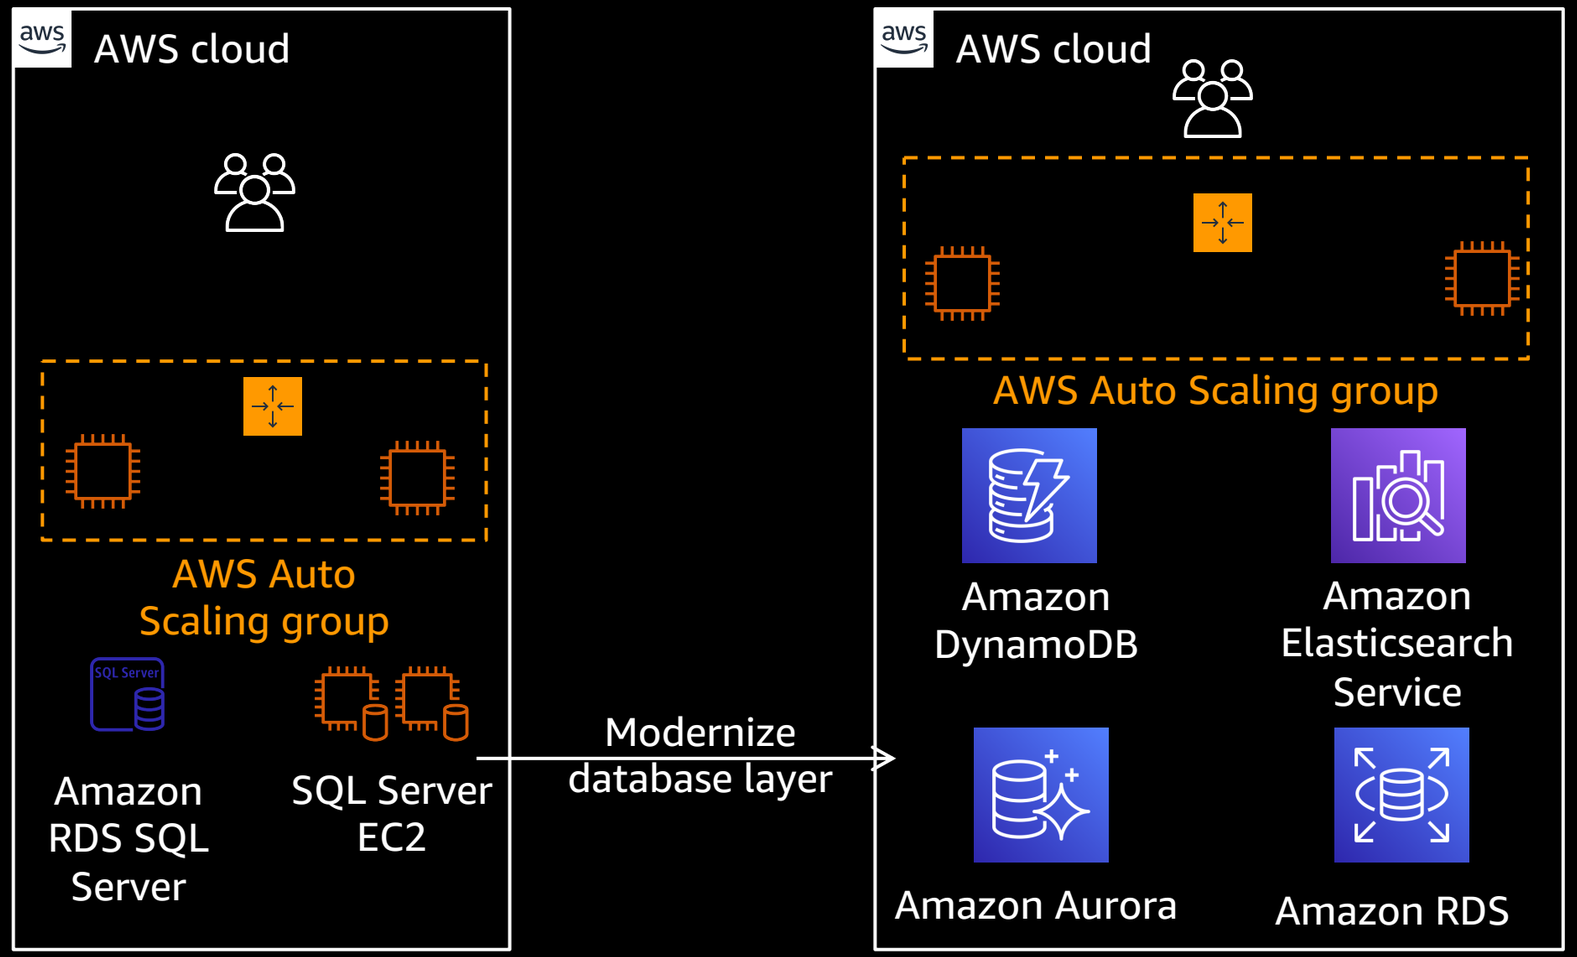
- ISV application with multiple MS SQL Server databases
- Databases were serving multiple functions
  - Session store
  - Logging information
  - Transactional store
  - Batch processing
  - Transactional reporting

Example user journey



- Start with a lift-and-shift to accelerate migration
- Consider an Amazon RDS SQL Server

Example user journey



- Analyze access patterns
- Migrate individual DB components to AWS purpose-built databases

Database management systems

$62.3 billion

MARKET OPPORTUNITY IN 2020

"By 2022, 75% of all databases will be deployed or migrated to a cloud platform."

Gartner

The Future of the Database Market Is Cloud, 2019

Customers are moving faster than ever to AWS managed database services

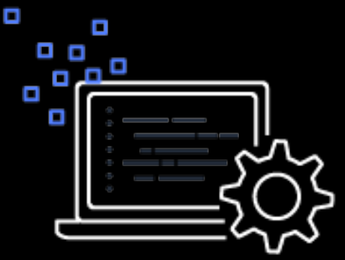
Customers are anxious to migrate from on-premises, legacy databases to AWS managed database services so they can



Spend less time managing infrastructure and more time innovating



Break free from the costs and restrictions of legacy databases



Build modern applications on AWS

Over 200,000 databases migrated to AWS

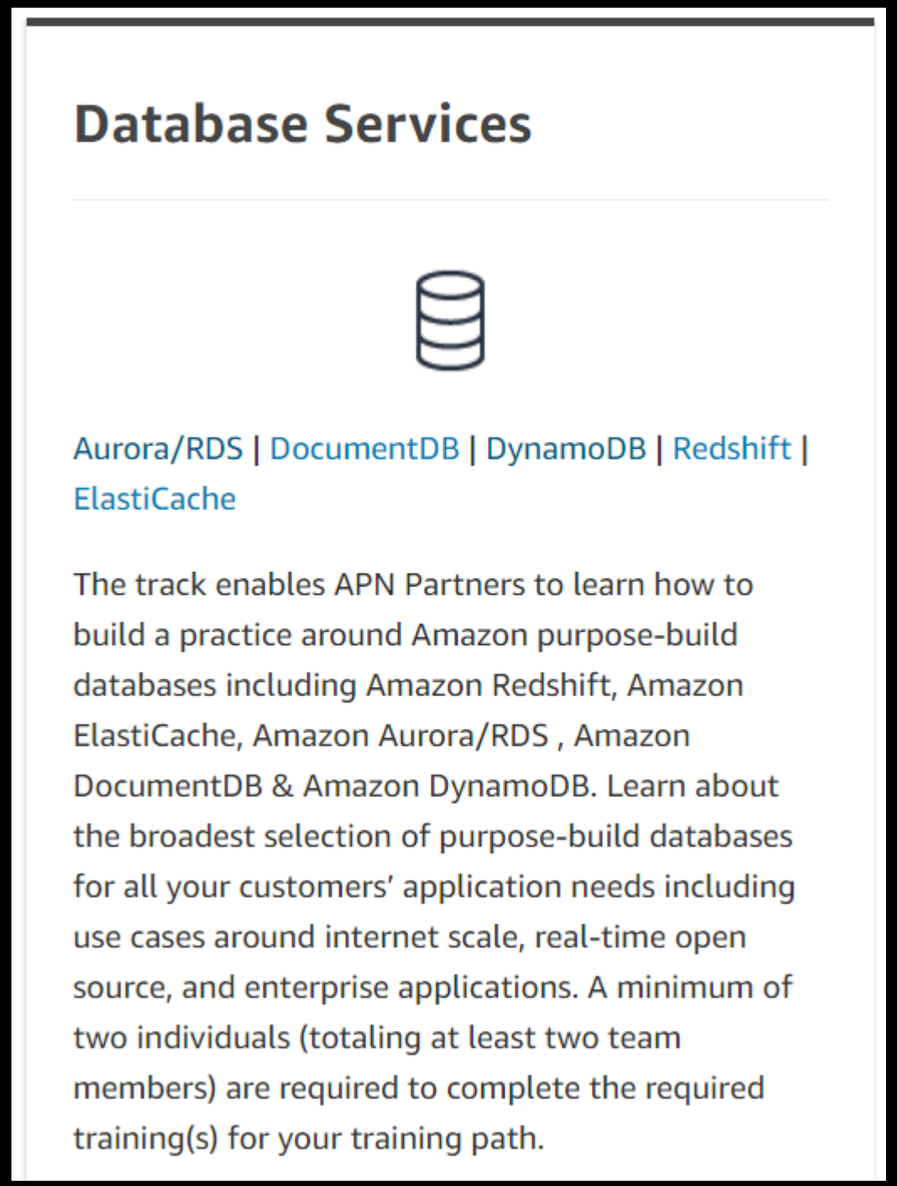


How to get started



© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.

APN Navigate program



- Self-paced online training at https://amzn.to/3jVLiHm
- Deep-dive workshops by AWS DB experts
  - Move to Managed for RDS Oracle / SQL Server Engines
  - DB Modernization with Amazon RDS and Amazon Aurora
  - DB Modernization with Amazon DynamoDB and Amazon DocumentDB
  - Database Modernization with AWS Database Services

"AWS Navigate is helping us to fast-track onboarding our newest headcount so they can provide working software sooner for our customers. In addition, Navigate is helping us to organize the collection of the collateral needed to apply for AWS Competencies."

APN Consulting Partner

Call to action

1. Learn



- Read about databases on AWS (https://aws.amazon.com/products/databases/) and user case studies (https://aws.amazon.com/solutions/case-studies/)
- Complete the DB Navigate training (https://amzn.to/2TUSDfE)
- Complete Database – Specialty certification

2. Join



- Join the APN Partner Program
- Become a database Service Delivery Designation Partner (https://amzn.to/3679m5g)
- Achieve migration competency

3. Engage



- Engage AWS team early and often
- Leverage AWS Partner Programs

Thank you!



AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.



Please complete the session survey

AWS re:Invent

© 2020, Amazon Web Services, Inc. or its affiliates. All rights reserved.