AWS EXPERIENCE TEL AVIV ON FLOOR28

NoSQL Databases - Fitting the right database for your use-case

Yogev Hen, Solutions Architect, AWS

November 17th , 2022

SQL vs. NoSQL

RELATIONAL VS. NON-RELATIONAL

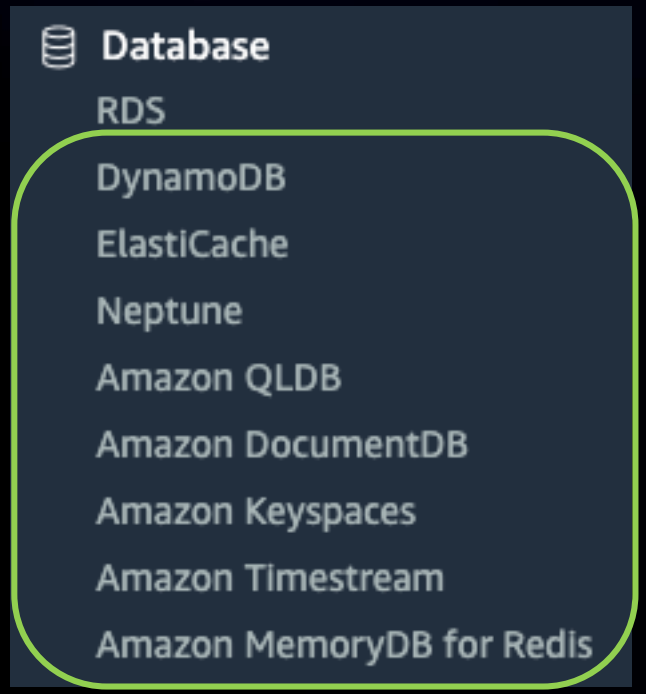

Normalized vs. denormalized data

Strong vs. eventual consistency

Vertical vs. horizontal scaling

Entity/relationship-based vs. access-based modeling

NoSQL database types

Type	Key-Value	Wide-Column	Document	In-Memory	Graph	Time-Series	Ledger
Strengths	Low-latency retrieval, extremely high performance	Query column data, high throughput	Complex / nested data, query any attribute including nested	Microseconds latency, specialized data structures	Creating and navigating data relations easily and quickly	Analytics over time-series data	Immutable, sequenced, verifiable history of changes
Use Cases	Media metadata, inventory tracking	Trade monitoring, fleet management	Content management, gaming	Leaderboards, caching	Fraud detection, social networking	IoT, health and usage metrics	Supply chain tracking, financial records

More General Purpose **More Specialized**

aws

Fully-managed NoSQL databases

Document

Amazon DocumentDB (with MongoDB compatibility)

Wide-Column

Amazon Keyspaces (for Apache Cassandra)

Key-Value

Amazon DynamoDB

What is Amazon DocumentDB?

Amazon DocumentDB (with MongoDB compatibility)

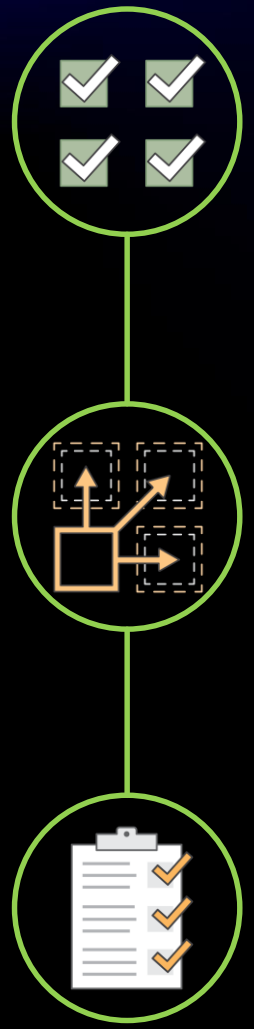

Fully managed

Scalable

MongoDB API compatible

Fully managed and scalable document database service that supports MongoDB workloads

Amazon DocumentDB (with MongoDB compatibility)

Fully managed

- Built-in high availability
- Backups enabled by default
- Durable by default
- Security best practices by default
- Automatic patching
- Monitoring and alerting

Amazon DocumentDB (with MongoDB compatibility)

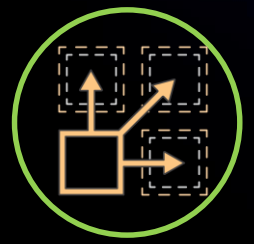

Scalable

Scale compute in minutes

Storage and IO autoscaling

Storage scales to 64TiB

Scale out to 15 replicas for millions of reads

Amazon DocumentDB (with MongoDB compatibility)

MongoDB API compatible

Applications, drivers, and tools can be used with Amazon DocumentDB with little or no change

Supports hundreds of APIs, operators, and stages

Continually working backward from customers to deliver the capabilities they need

Amazon DocumentDB architecture

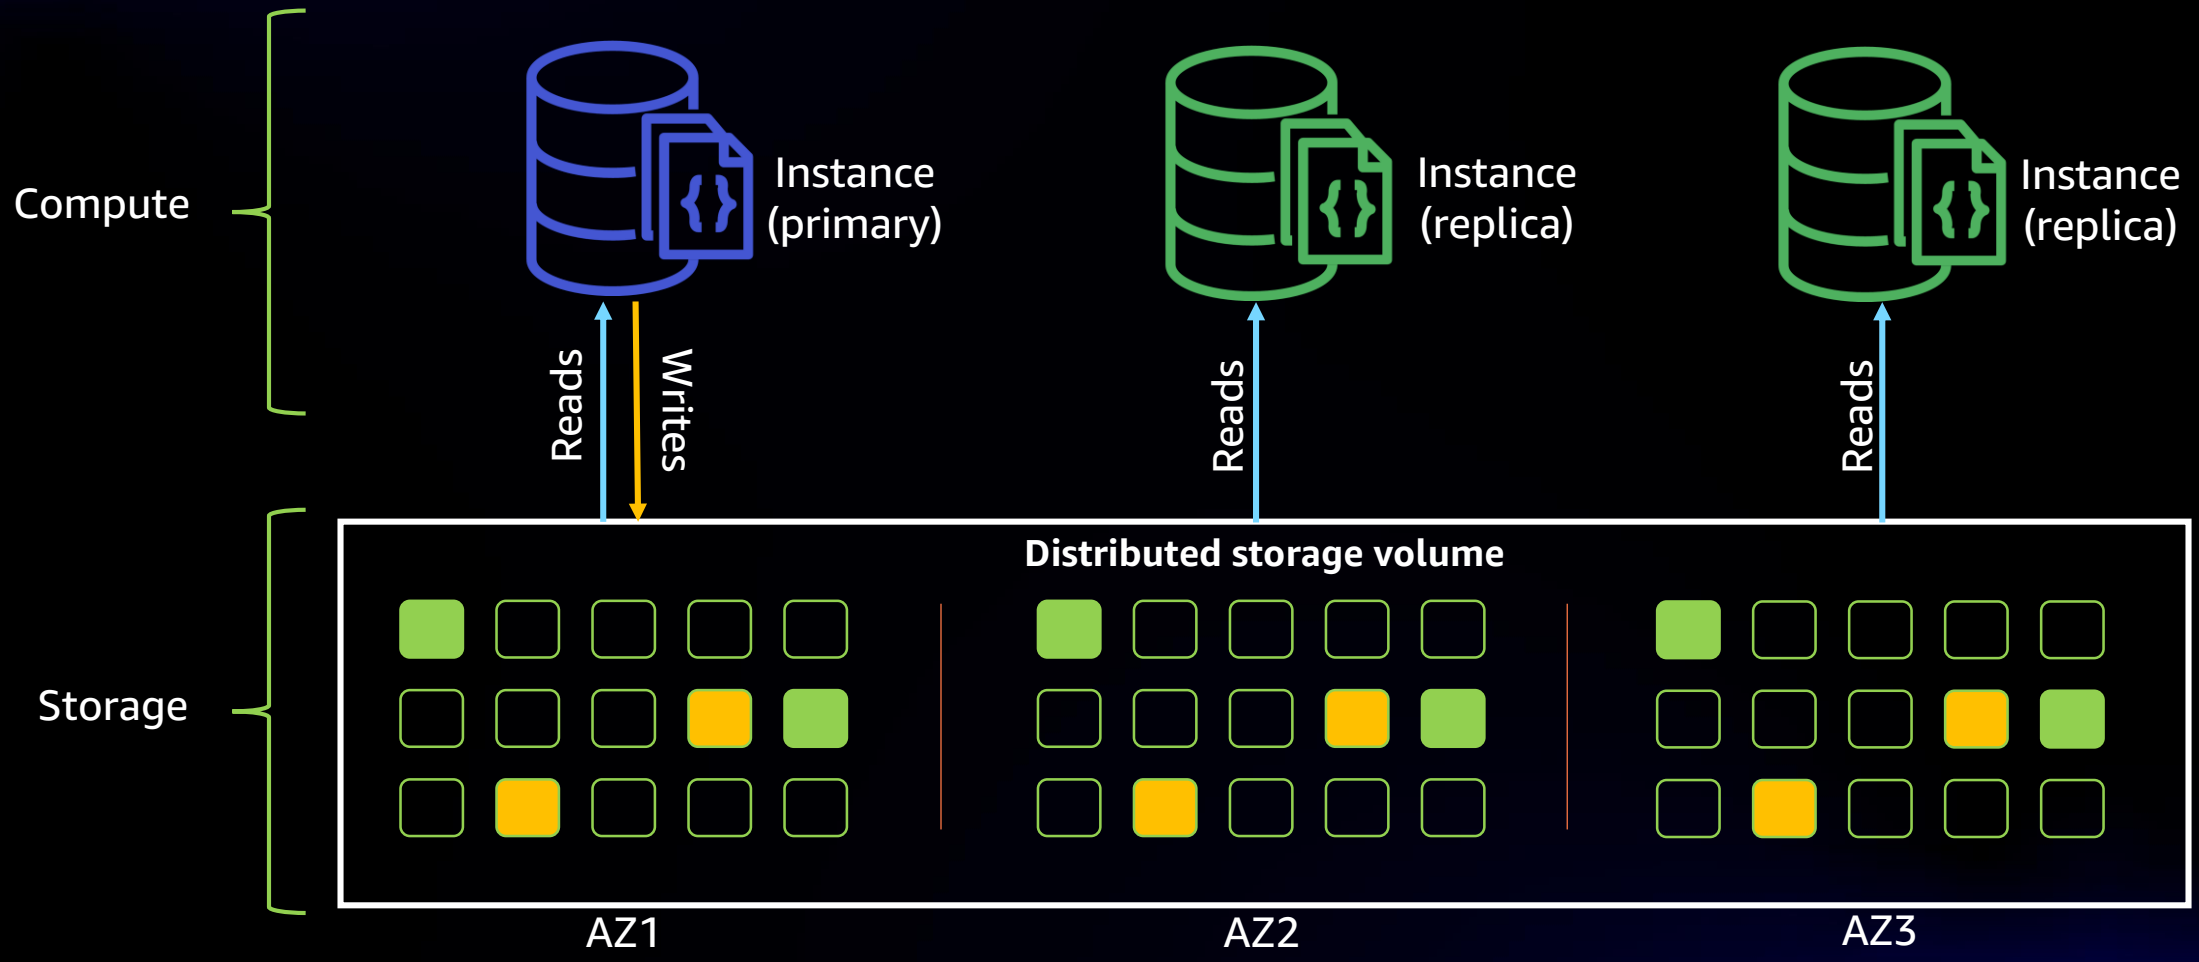

When should you use a Amazon DocumentDB?

JSON data

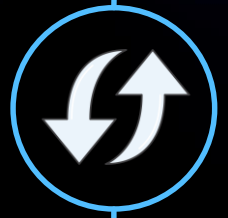

Flexible schema for fast iteration

Ad hoc query capabilities

Flexible indexing

Operational and analytics workloads

Amazon DocumentDB makes it easy to store, query, and index JSON data

aws

Industries

Web publishing

Product catalog

IoT

Regulatory documents

Gaming

Content management

What is Amazon DynamoDB?

Amazon DynamoDB

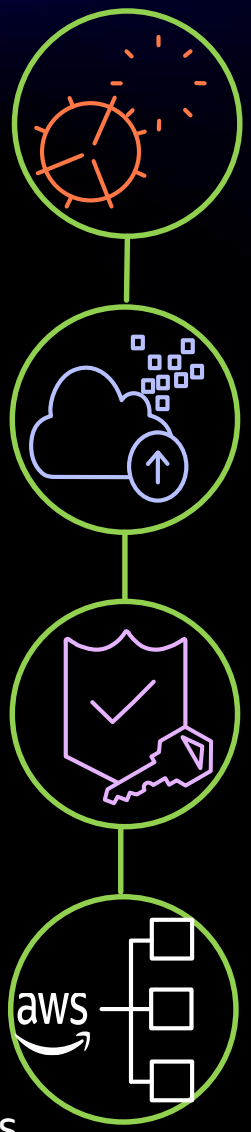

Performance at any scale

Serverless

Secure and Resilient

Built-in integration with other AWS services

Fast and flexible NoSQL database service for any scale

Amazon DynamoDB

Performance at any scale

Millions of requests per second

Single-digit-millisecond latency

Nearly unlimited throughput and storage

Amazon DynamoDB

Serverless

No hardware provisioning

scale from zero to infinity

No downtime maintenance

No maintenance windows

Amazon DynamoDB

Secure and Resilient

Data encryption at rest by default

Global replication

Up to 99.999% availability SLA

Fine-grained access control at a table, item or attribute level

aws

Amazon DynamoDB

Built-in integration with other AWS services

Amazon DynamoDB Streams and Kinesis Data Streams support

CloudWatch Contributor Insights

Import and export data to Amazon S3

Serverless & Event Driven Architecture

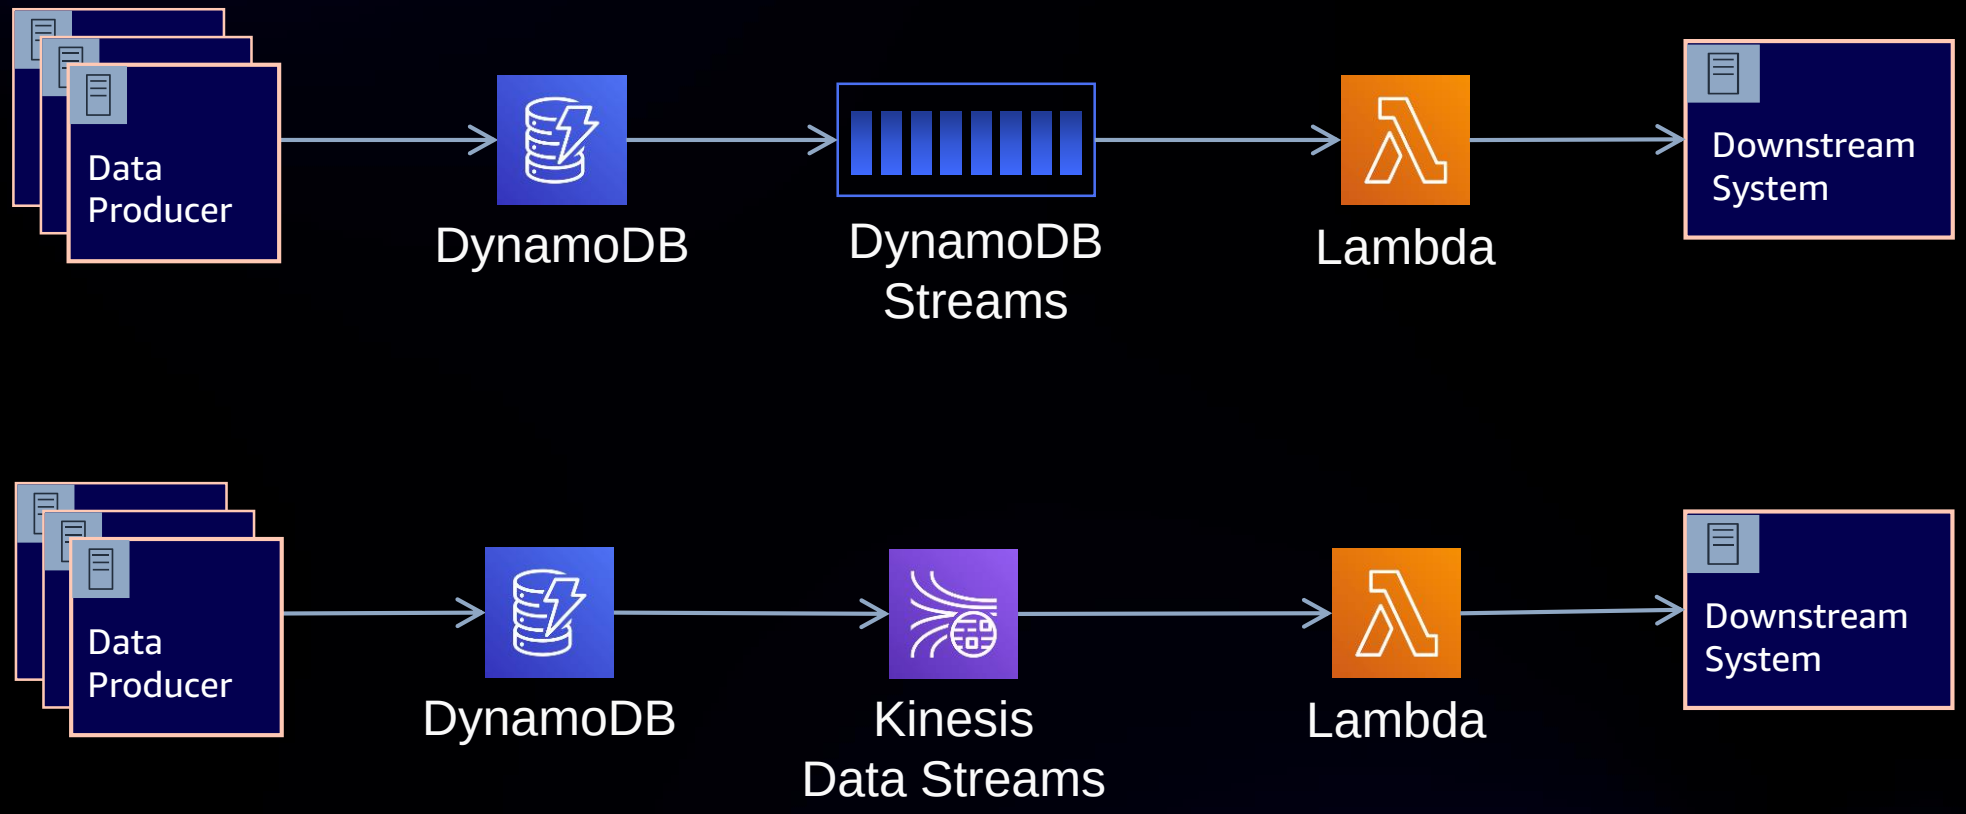

Amazon DynamoDB – Backup and Restore

On-demand backups for long-term data archival and compliance

Point in time restore for short term retention and data corruption protection (35 days)

Point in time recovery with restore times in a few hours depending on table size

Time-to-live (TTL)

TTL Attribute

ID	Name	Size	Expiry
1234	A	100	1456702305
2222	B	240	1456702400
3423	C	150	1459207905

Features

- Automatically delete items from a table based on expiration timestamp
- User defined TTL attribute in epoch time format
- TTL activity recorded in DynamoDB Streams

Key Benefits

- Reduce costs by deleting items no longer needed
- Optimize application performance by controlling table size growth
- Trigger custom workflows with Streams and Lambda

Amazon DynamoDB Standard-Infrequent Access (Standard-IA) table class

Lower storage costs

- The Standard-IA table class offers 60% lower storage costs than DynamoDB Standard tables.

No performance trade-offs

Standard-IA tables offer the same performance, durability, data availability, and massive scalability as existing DynamoDB Standard tables.

No developer overhead

Switch between table classes with a single click in the DynamoDB console, or using the AWS CLI or AWS SDK. Also, use the same DynamoDB APIs and service endpoints.

DynamoDB Accelerator (DAX)

In-memory cache

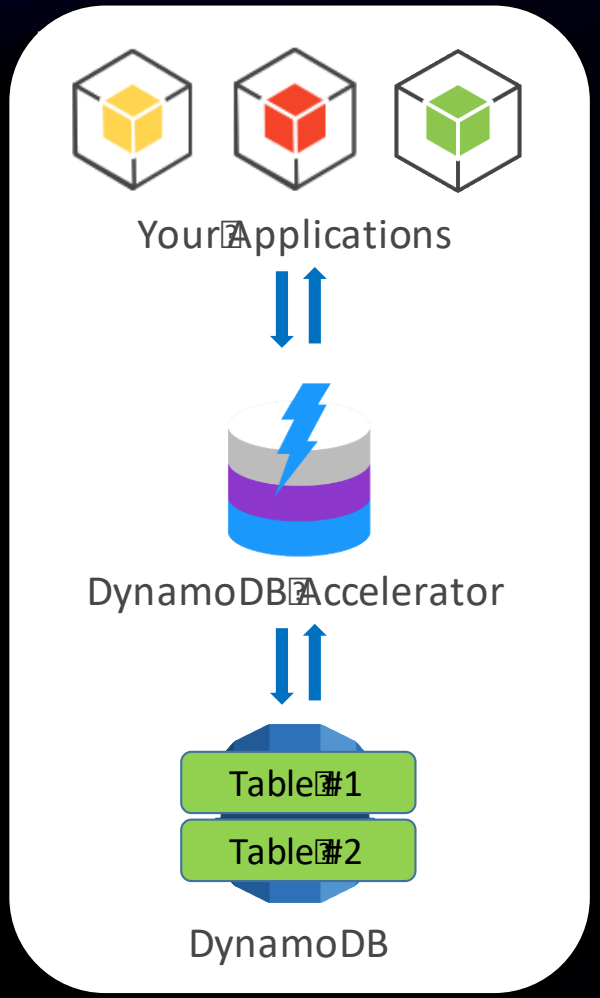

Features

- **Fully managed, highly available**
- **DynamoDB API compatible**
- **Write-through**
- **Flexible**
- **Scalable**
- **Security**

DynamoDB use cases by industry

Finance
Fraud detection
User transactions
Mainframe offloading
(Capital One, Vanguard, Fannie Mae)

Gaming
Game states
Leaderboards
Player data stores
(Riot Games, Electronic Arts, PennyPop)

Software and Internet
Metadata caches
Ride-tracking data stores
Relationship graph data stores
(Uber, Lyft, Swiggy, Snap, Duolingo)

Ad tech
User profile stores
Metadata stores for assets
Popular-item cache
(AdRoll, GumGum, Branch, DataXu)

Retail
Shopping carts
Workflow engines
Customer profiles
(Nordstrom, Nike, Zalando, Mercado Libre)

Media and Entertainment
User data stores
Media metadata stores
Digital rights management stores
(Airtel Wynk, Amazon Prime, Netflix)

Industires

Media and Entertainment

Software and Internet

Software and Internet

Banking and Finance

Manufacturing

Education

What is Amazon Keyspaces?

Amazon Keyspaces

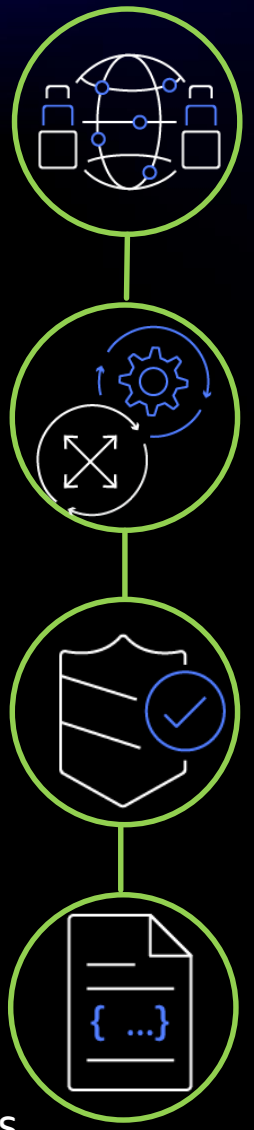

Performance at scale

Serverless

Highly available and secure

Compatible with Apache Cassandra

Scalable, highly available, and managed Apache Cassandra–compatible database service

Amazon Keyspaces

Performance at scale

Tables scale up and down automatically

Single-digit-millisecond latency

Virtually unlimited throughput and storage

Amazon Keyspaces

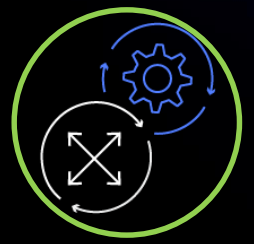

Serverless

No need to provision, patch, or manage servers

On-demand or provisioned capacity mode

Fully managed Time To Live (TTL)

Amazon Keyspaces

Highly available and secure

Fully managed and highly available data storage

Encryption at rest and in transit

99.99% availability SLA

Point-in-time recovery

Access management

Secure network connectivity

Amazon Keyspaces

Compatible with
Apache Cassandra

Compatible with CQL

Same drivers and developer tools

DocumentDB indicators

- MongoDB-compatibility
- Maximum data model flexibility
- The ability to query nested data
- Read-heavy workloads that can tolerate eventual consistency
- Workloads up to 64 TB of data

DynamoDB indicators

- Workloads with well-defined query patterns
- Single-digit millisecond latency
- Infinite scalability
- Mature, rich feature set
- Tight integration with AWS services

Keyspaces indicators

- Cassandra compatibility
- Schema validation
- Multiple independent sorts in a query

Q&A

Thank You!