DAT324

Introducing Amazon Managed Apache Cassandra Service

Arturo Hinojosa
Senior Product Manager
Amazon Web Services

Akshat Vig
Principal Engineer
Amazon Web Services

AWS re:Invent

aws

Agenda

Overview of Apache Cassandra

Introduction to Amazon Apache Managed Cassandra Service

NoSQL databases: Why it is hard?

Getting started

Questions

Apache Cassandra

- Open-source, wide-column data store
- Fast write throughput
- Distributed architecture
- Cassandra Query Language (CQL)
- Common use cases include
  - Transaction logging
  - Time-series data
  - Event history

Cassandra data model

Cassandra is a column database that can store key-value data

A keyspace is the outermost data container

Within a keyspace, you can create multiple tables

Each table has a primary key; primary keys uniquely identify items

Primary keys are composed of one or more partition keys and, optionally, one or more clustering columns

Clustering columns are used for sorting

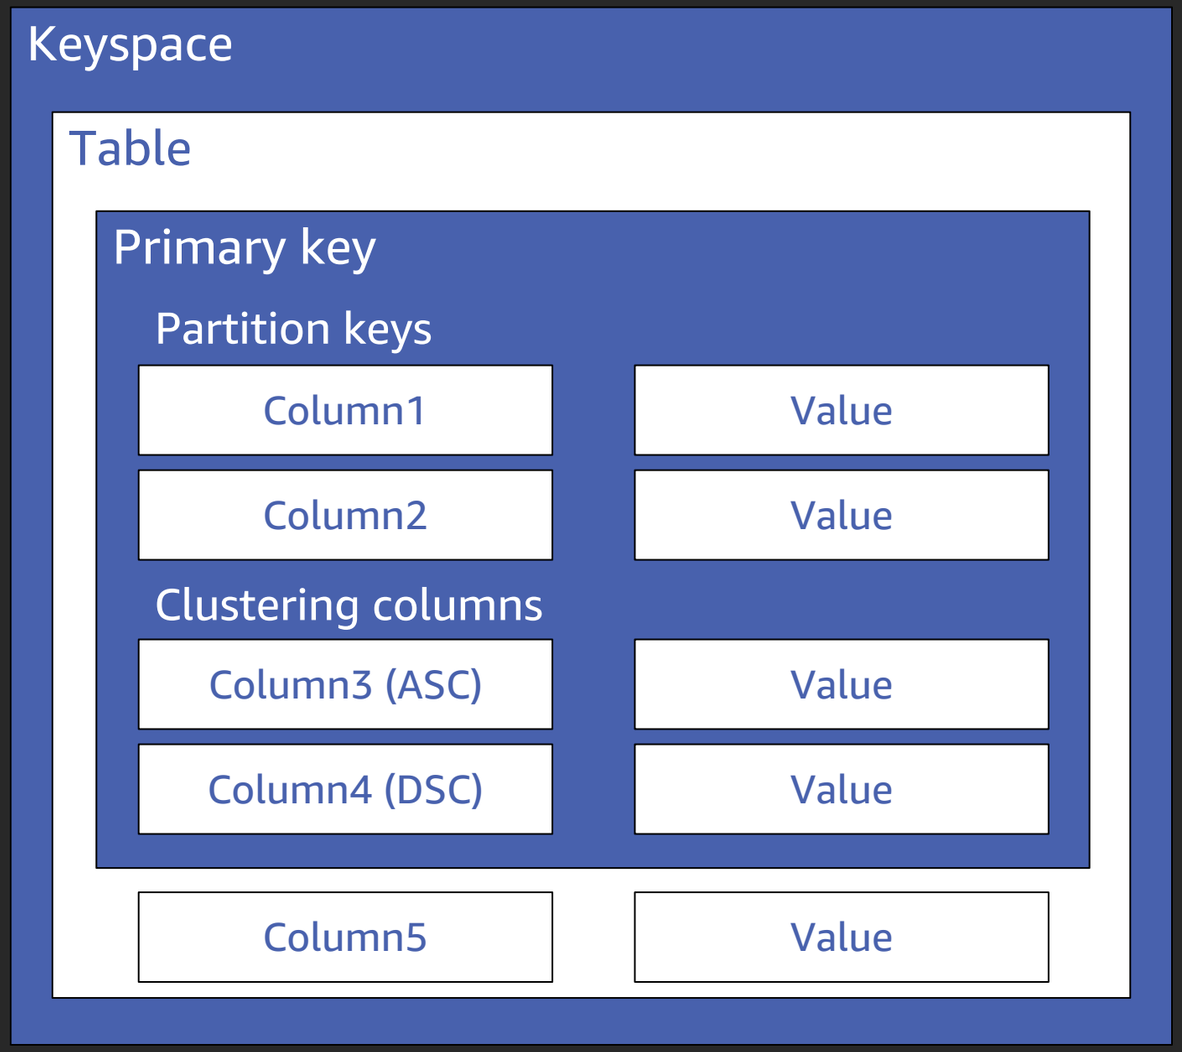

Challenges of self-managing Cassandra

- Requires specialized expertise to deploy, configure, and manage Cassandra software
- It’s complex to scale Cassandra clusters up and down
- Customers often build clusters for peak load, leading to higher infrastructure costs and paying for unused capacity
- Upgrading Cassandra to new versions is difficult
- Backing up and restoring data is complex
- Customers are responsible for patching, updating, and maintaining infrastructure

Introducing Amazon Managed Cassandra Service

A **scalable**, **highly available**, and **managed** Apache Cassandra–compatible database service

Demo

Technical deep dive

NoSQL databases @ Amazon

Secure

Durable

Highly available

Scalable

Fully managed

NoSQL databases @ Amazon

- Security
  - Authorization and authentication
  - Encryption @ rest
  - Encryption in transit

NoSQL databases @ Amazon

Scalability challenges

- Consistent
- Performance @ scale

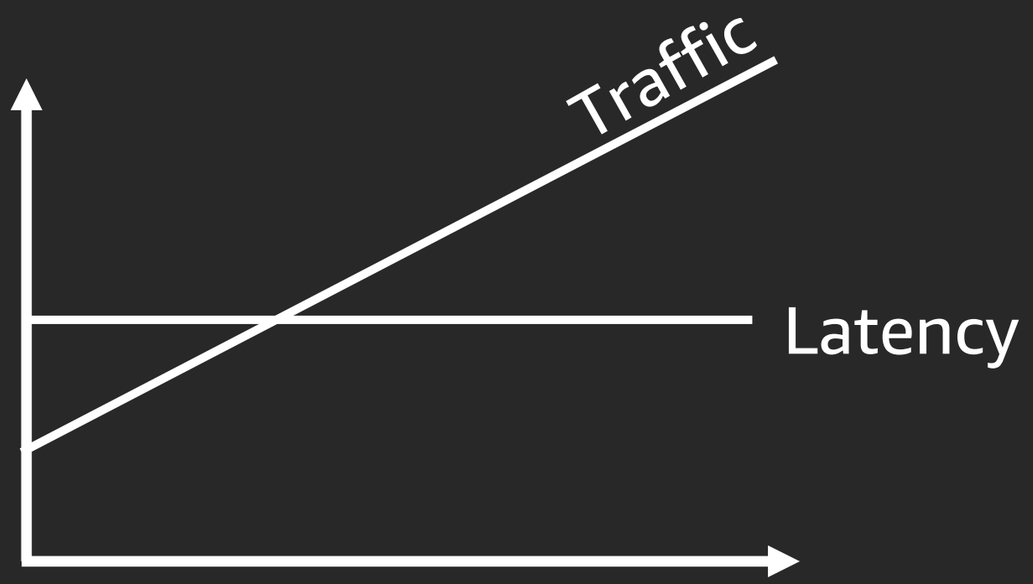

NoSQL databases @ Amazon

Scalability challenges

Uneven access, partition management, large partitions, capacity management

NoSQL databases @ Amazon

Scalability challenges: Node failures

NoSQL databases @ Amazon

- Availability and durability
  - Replicate the data across multiple AWS Availability Zones
  - Backup and restore

NoSQL databases @ Amazon

Fully managed magical experience

- Zero downtime
- Scale as you need, pay as you go

Benefits

Apache Cassandra–compatible

No servers to manage

Consistent performance @ scale

Highly available and secure

Getting started

Getting started

Create an AWS Identity and Access Management (IAM) user

Create and attach MCS policy to the entity

Generate server side specific credentials

Access using console or programmatically using CQLSH or Casandra drivers

Managed Cassandra Service Policies

Manage access to keyspaces and tables

Scoped to individual resources

Admin and read-only managed IAM policies available

Service-specific credentials

```
aws iam create-service-specific-credential \
    --user-name alice \
    --service-name cassandra.amazonaws.com
```

Using CQLSH

Download certificate

```
curl https://www.amazontrust.com/repository/AmazonRootCA1.pem
```

Edit $HOME/.Cassandra/cqlshrc

```
[connection]
port = 9142
factory = cqlshlib.ssl.ssl_transport_factory

[ssl]
validate = true
certfile =  $HOME/AmazonRootCA1.pem
```

Connect

```
cqlsh cassandra.us-east-2.amazonaws.com 9142 -u "alice-at-111122223333" -p "wJalrXUtnFEMI/K7MDENG/bPxRfiCYEXAMPLEKEY
```

Programmatic access

1. Download the Amazon digital certificate using the following command and save it to your home directory:

```
curl https://www.amazontrust.com/repository/AmazonRootCA1.pem
```

2. Convert the Amaozn digital certificate to a trustStore file:

```
openssl x509 -outform der -in AmazonRootCA1.pem -out temp_file.der
keytool -import -alias cassandra -keystore cassandra_truststore.jks -file temp_file.der
```

3. Attach the trusStore file in the JVM arguments:

```
-Djavax.net.ssl.trustStore=path_to_file/cassandra_truststore.jks
-Djavax.net.ssl.trustStorePassword=amazon
```

Architecture

Cassandra

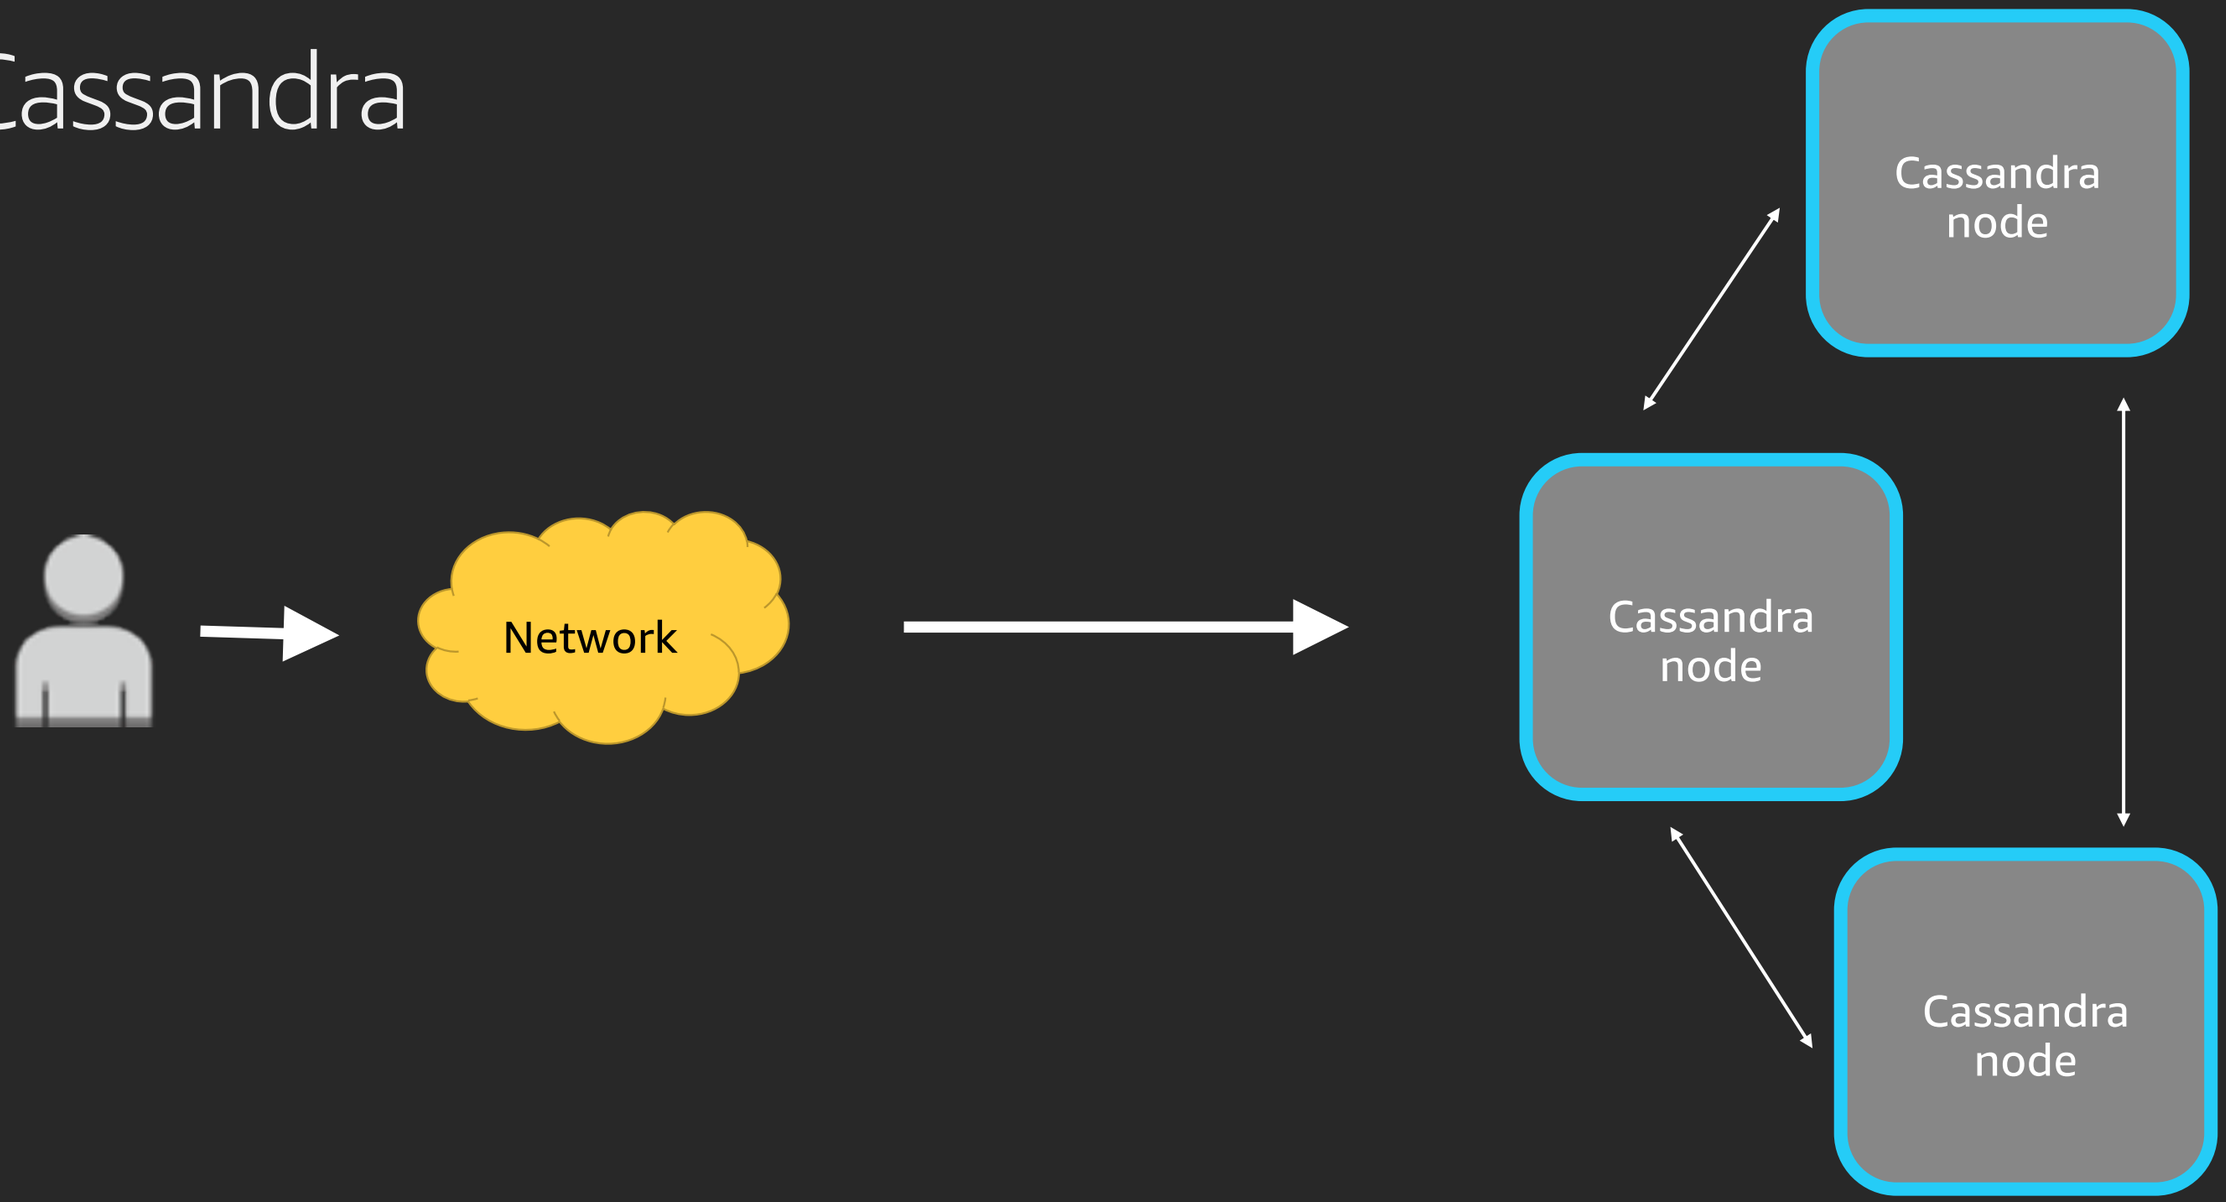

Pluggable storage

Amazon Managed Cassandra Service

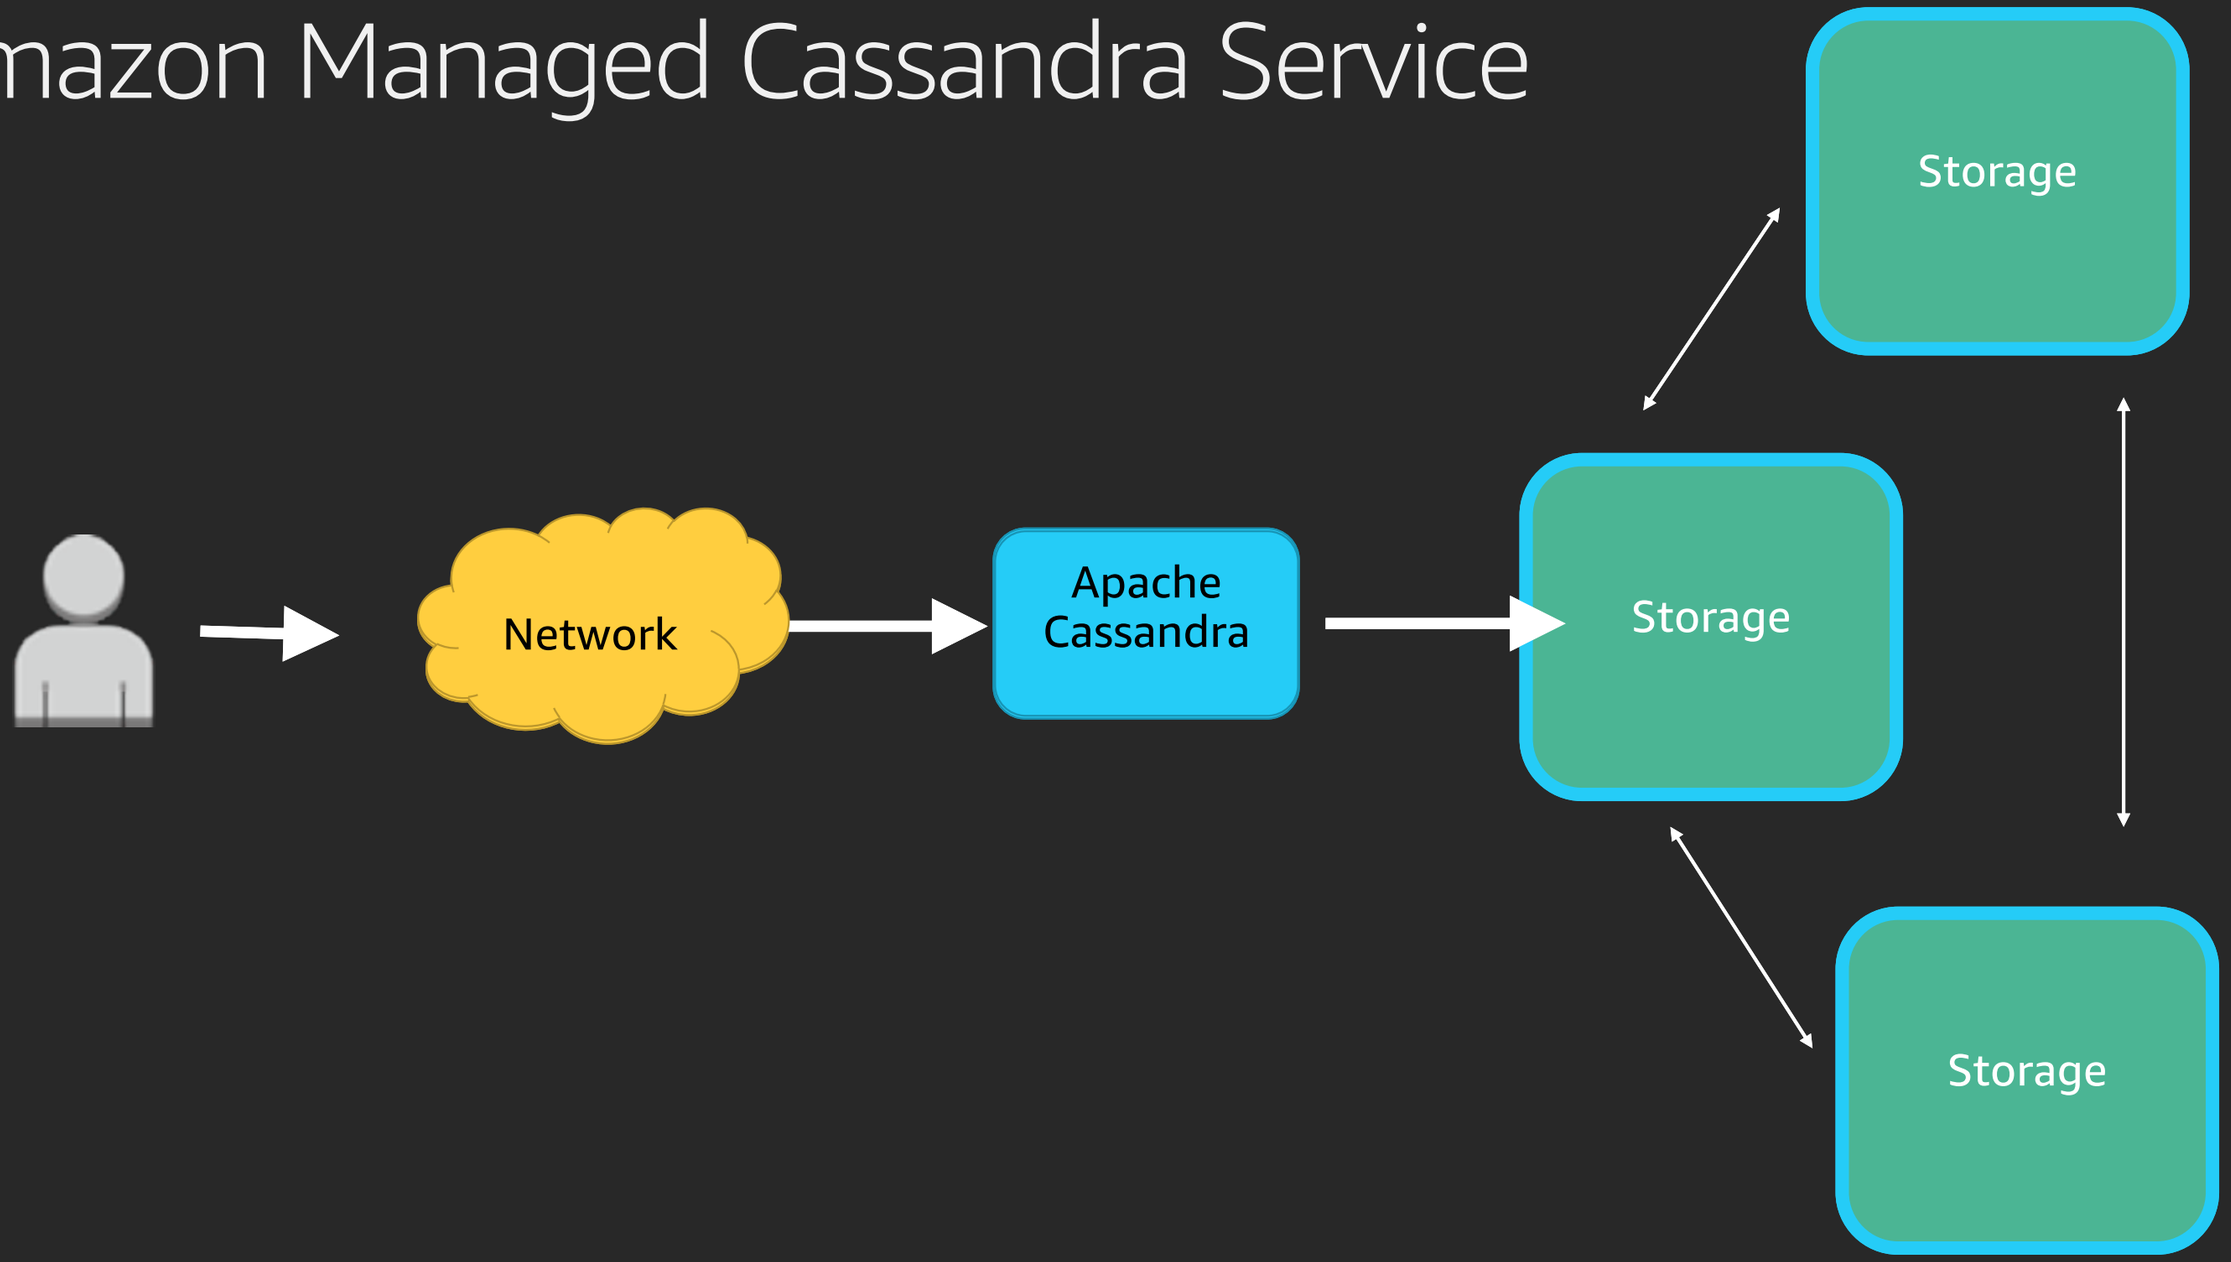

Questions

Please complete the session survey in the mobile app.

Thank you!

Arturo Hinojosa

 @SenorHinojosa

Akshat Vig

 @akshatvig

AWS re:Invent

aws