

SPRINGONE2GX

WASHINGTON, DC

# Cassandra and Grails

By Jeff Beck

@beckje01



springone 2GX



# Agenda

- whoami
- What is Cassandra
- Cassandra Terminology
- Basic Data Modeling
- What is Available for Grails
  - Cassandra ORM
  - Cassandra GORM
  - Astyanax
  - Java Native Driver
- Which to Use
- Basic Configuration
- Tools to Use
- Advanced Data Modeling
- TokenStore Example
- Cleaning Up Mistakes

# \$ whoami

## Jeff Beck

- Engineer at SmartThings
- Ratpack Team Member
- Organize Minneapolis Cassandra Meetup



 beckje01 



# What is Cassandra

- Fast Writing Data Store
- Highly-Available, Distributed, Tuned Consistency
- Master-less Replication
- Redundancy Configurable
- Cluster can span Data Centers

# When Cassandra

- High write throughput needs
- Cross datacenter replication required
- Mature data access patterns

# When Not Cassandra

- When you depend on transactions
- Not sure how you will access the data
- Lack the shared responsibility on the running of Cassandra

# Cassandra Terminology

- Thrift - The old way to interface with Cassandra. Now longer getting enhancements.
- CQL / CQL v3 - The new native protocol that replaced Thrift and the new way to model data.
- Replication Factor - How many nodes will have a copy of the data.
- Consistency Level - When executing a query how many nodes need to agree.
- Cluster vs Ring

# Cassandra - Cluster vs Ring





# Data Modeling

It's hard you are going to make mistakes.

# Basic Data Modeling

- You need to focus on Fan Out on Write, not the more traditional Fan Out on Read.
- With Cassandra you really need to focus on the query not the thing being represented.

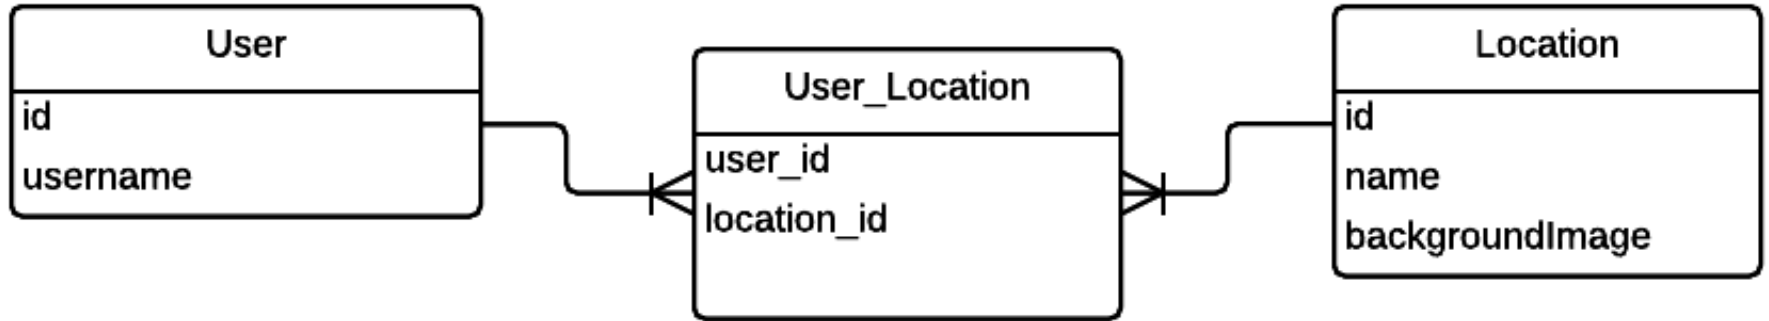
# Basic Data Modeling Hints

- Duplicate Data
- Natural Key
- Careful around adding an index
- Get a Second Opinion

# Data Modeling Example Problem

Users have access to a location we want to get the details of a location for a particular user.

# SQL Solution



# CQL Solution

```
1 CREATE TABLE user_location (  
2     username text,  
3     location_id text,  
4     shard_id text,  
5     name text,  
6     background_image text,  
7     PRIMARY KEY (username, location_id)  
8 );
```

Notice the fanout on write, given 2 users with access to a location we will write 2 rows. We have to do that on the application side.

# Idempotent Upserts

Model your interactions as all idempotent upserts if possible. This is the best case for Cassandra.

# What is Available for Grails

- ORMs
  - Cassandra ORM
  - Cassandra GORM
- Direct Cassandra Access
  - Astyanax Plugin
  - Java Native Driver



# Cassandra ORM

The older of the two ORMs this was modeled after GORM but doesn't fully participate in GORM.

Based on top of *Astyanax* and Thrift.

Provides many nice features out of the box that makes getting something working very quick.

# Advantages

- Counters
- Indices
- Relations
- Expando map

# Drawbacks

- Dealing with bootstrapping
- Based on Thrift
- Only easy data access is via the application

# Example Object

```
1 class AppScheduledEventHistory {  
2   UUID id  
3   String shard  
4   String appld  
5   Date scheduledTime  
6   Long executionTime  
7   AppScheduledEventStatus executionStatus  
8 }
```

```
9 static cassandraMapping = [  
10   unindexedPrimaryKey: "id",  
11   explicitIndexes: [ ["appld"],  
12     ["scheduledEventId"],  
13     ["timeBucket","shard"],  
14     ["timeBucket","shard","executionStatus"]  
15   ],  
16   counters: [  
17     [groupBy: ["scheduledTime","shard","executionStatus"]]  
18   ]  
19 ]  
20 }
```

# Example Query

```
1 AppScheduledEventHistory.get(params.id.toUUID())
2
3 AppScheduledEventHistory.findAllByAppId(appId)
4
5 AppScheduledEventHistory.findAllByTimeBucketAndShard(timeBucket,
shard)
```

# More Advanced Query

```
1 AppScheduledEventHistory
2 .findAllByTimeBucketAndShardAndExecutionStatus(
3     params.timeBucket,
4     params.id,
5     params.status,
6     [consistencyLevel: scheduledEventService.consistencyLevel])
```

# Query with Counters

```
1 AppScheduledEventHistory
2 .getCountsGroupByScheduledTimeAndShardAndExecutionStatus(
3     start: start,
4     finish: finish,
5     reversed: true,
6     consistencyLevel: scheduledEventService.consistencyLevel)
```

# Expando Map

```
1 class DeviceData {
2     String deviceId
3     Map data
4
5     static cassandraMapping = [
6         unindexedPrimaryKey: "deviceId",
7         expandoMap         : "data"
8     ]
9 }
```



# Using Expando Map

```
1 def d = new DeviceData()  
2  
3 d.first = 'Bob'  
4 d.last = 'Dole'  
5  
6 d.save()
```

# Cassandra GORM

A true GORM implementation for Cassandra based on Spring Data's Cassandra work which under the covers uses the Java Native Driver.

It is relatively new, so there are rough edges.

# Advantages

- Participates in GORM
- Can auto create the schema
- Same GORM syntax

# Disadvantages

- Allows use of Cassandra secondary indices
- All or nothing mapping\*

\**Implementation flaw, not an issue with GORM*

# Example Object

```
1 class Subscription {
2   UUID id
3   URI product
4   Date startDate
5   State state
6   Date endDate
7
8   static mapping = {
9     table "subscription"
10      id generator: 'assigned'
11      version false
12  }
13
14  static constraints = {
15    startDate nullable: false
16    endDate nullable: true
17    state nullable: false
18    product nullable: false
19
20  }
21 }
```

# Example Query

```
1 def subs = Subscription.get(uuid)
```

# With a Complex Primary Key

```
1 class Person {
2   String lastName
3   String firstName
4   Integer age = 0
5   String location
6
7   static mapping = {
8     id name:"lastName", primaryKey:[ordinal:0, type:"partitioned"], generator:"assigned"
9     firstName index:true, primaryKey:[ordinal:1, type: "clustered"]
10    age primaryKey:[ordinal:2, type: "clustered"]
11    version false
12  }
13 }
```

# Query Example

```
1 //Will work
2 def people = Person.findAllbyLastNameAndFirstName('Beck','Jeff')
3
4 //Will Fail
5 def people = Person.findAllbyFirstName('Jeff')
```



# Astyanax

This is a simple Grails plugin wrapping the Netflix Astyanax library.

A client for working with Cassandra.

# Advantages

- More direct access to Cassandra
- Works with legacy Thrift tables
- Provides features of modern Cassandra clients
  - Client Side Load Balancing
  - Token Aware
  - Node Discovery

# Drawbacks

- A slightly outdated view of Cassandra.
- Still supported but not seeing tons of new features.
- Lots of configuration needed to get a sensible starting point.

# Example Interaction

```
1 final SETTINGS_CF = new ColumnFamily("EventSettings",  
StringSerializer.get(), StringSerializer.get())  
2  
3 def m = astyanaxService.keyspace().prepareMutationBatch()  
4 m.setConsistencyLevel(consistencyLevel)  
5 m.withRow(SETTINGS_CF, "shards").putColumn(shard,  
DateUtils.timeStamp(date))  
6 m.execute()
```

# Java Native Driver

The new Driver put out by DataStax for Cassandra.

It doesn't have a plugin or anything yet so you must work with it directly.

# Advantages

- Supports Asynchronous Calls
- Default configuration is a sensible start
- Under active development
- Performance benefit for CQL prepared statements in C\* 2.1+
- Low level

# Drawbacks

- Low level
- Have to wire up and manage all your own connections

```

1 class CQLSessionService implements InitializingBean {
2     //...
3
4     @Override
5     void afterPropertiesSet() throws Exception {
6         keyspace = grailsApplication.config.defaultKeyspace
7         def clusterBuilder = new Cluster.Builder()
8             .withRetryPolicy(DowngradingConsistencyRetryPolicy.INSTANCE)
9             .withLoadBalancingPolicy(new TokenAwarePolicy(new
DCAwareRoundRobinPolicy()))
10            .withReconnectionPolicy(new ConstantReconnectionPolicy(100L))
11            .withSocketOptions(new SocketOptions().setKeepAlive(true))
12
13        def contacts = ["192.168.0.1", "192.168.0.3"]
14        contacts.each { String contact ->
15            clusterBuilder.addContactPoint(contact)
16        }
17
18        cluster = clusterBuilder.build()
19        cluster.getConfiguration()
20        session = cluster.connect()
21    }
22 }

```

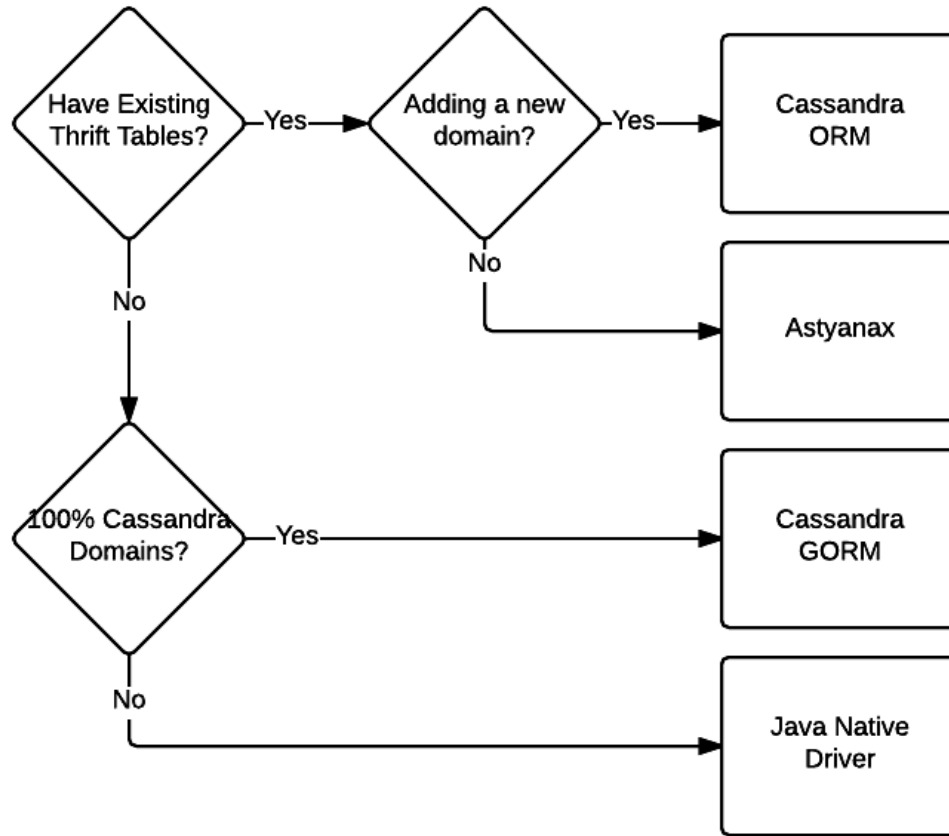


```

1 class LoginAttemptService {
2     def CQLSessionService
3
4     static transactional = false
5
6     private final String tableName = "LoginAttempts"
7     private PreparedStatement login_attempt_insert
8
9     @PostConstruct
10    private void prepareStatements() {
11        def session = CQLSessionService.session
12        def keyspace = CQLSessionService.keyspace
13
14        login_attempt_insert = session.prepare("""
15            INSERT INTO ${keyspace}.${tableName}
16            (userName, successful, ipAddress, userAgent, method, dateCreated)
17            VALUES ( ?, ?, ?, ?, ?, ?);""")
18    }
19 }

```

# Which to Use



# Client Config Options

- Seeds
- Datacenter *Aware*
- Token *Aware*
- Retries
- Speculative Execution
- Address Translator

# Dev Tools

- [CCM](#) - Local cassandra cluster
- [DevCenter](#) - Visual query tool
- [CQL Data Modeler](#) - Helps model CQL3 tables.
- [Pillar](#) - Migrations

# Server Tools

- [OpsCenter](#) - Visual display of clusters and nice dashboards
- Log aggregation - Cassandra generates good logs you need to monitor for increased ERROR and WARN rates
- [Cassandra Reaper](#) - Repair tool

# Data Modeling Example Problem

Store every login and login attempt to the system. We want to find all attempts per user for either successful or unsuccessful sorted by the date it happened.

# Data Modeling Example Solution

```
1 CREATE TABLE login_attempts (  
2   user_name text,  
3   successful boolean,  
4   ip_address text,  
5   user_agent text,  
6   method text,  
7   date_created timestamp,  
8  
9   PRIMARY KEY(user_name, successful, date_created)  
10 );
```

# CQL3 Primary Key

- Defined as `PRIMARY KEY (last_name, first_name)`
- Can query parts going left to right but no skipping.

1 **PRIMARY KEY** (last\_name, first\_name)

2 **PRIMARY KEY** ((last\_name), first\_name)

One is just short hand for two. This is important because the first parameter is the partition key, the rest is the cluster key.



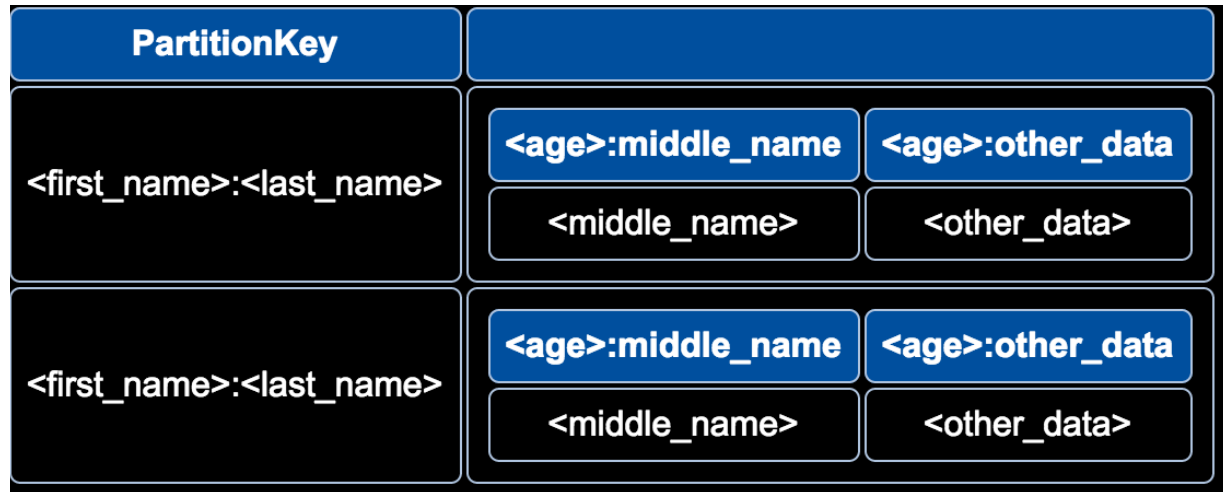
# CQL3 Primary Key

1 PRIMARY KEY (last\_name, first\_name, age)

2 //Not Equal

3 PRIMARY KEY ((last\_name, first\_name), age)

The first part of the primary key is the partition key, you will always need to provide this part.



# Token Store Example

Implement TokenStore interface from SpringSecurity OAuth.

# Token Store Example

It takes four tables to take the place of what is done in two in the JDBC version of the token store.

It is an easy process to support this as we have a known interface that defines all our queries and we can look at the SQL statements used to verify the access patterns.

# Token Store Example

```
1 CREATE TABLE oauth_access_token_by_client_with_token (  
2   client_id text,  
3   client_mod int,  
4   user_id text,  
5   oauth_token text,  
6   access_token text,  
7   authentication text,  
8   last_modified timestamp,  
9   PRIMARY KEY ((client_id, client_mod), user_id, oauth_token)  
10 );
```

PartitionKey			
<client_id>:<client_mod>	<user_id>:<oauth_token>:access_token	<user_id>:<oauth_token>:authentication	<user_id>:<oauth_token>:last_modified
	<access_token>	<authentication>	<last_modified>
<client_id>:<client_mod>	<user_id>:<oauth_token>:access_token	<user_id>:<oauth_token>:authentication	<user_id>:<oauth_token>:last_modified
	<access_token>	<authentication>	<last_modified>

# Token Store Example

```
1 session.prepare("""
2     BEGIN BATCH
3         INSERT INTO ${keyspace}.oauth_access_token
4             (oauth_token, access_token, authentication, last_modified)
5             VALUES (?, ?, ?, ?);
6         INSERT INTO ${keyspace}.oauth_access_token_by_refresh
7             (refresh_token, access_token, authentication, last_modified)
8             VALUES (?, ?, ?, ?);
9         INSERT INTO ${keyspace}.oauth_access_token_by_authentication
10            (authentication_id, access_token, authentication, last_modified)
11            VALUES (?, ?, ?, ?);
12        INSERT INTO ${keyspace}.oauth_access_token_by_client
13            (client_id, client_mod, user_id, access_token, authentication,
last_modified)
14            VALUES (?, ?, ?, ?, ?, ?);
15        APPLY BATCH;
16        """).setConsistencyLevel(ConsistencyLevel.ONE)
```

# Bucketing

For some items we want to query on such as Client (ex: iOS, Android) there are too many entries to depend on a single wide row. So we create buckets, when searching for things by Client we will run a query for each bucket.

```
1 Hashing.consistentHash(  
    Hashing.murmur3_32().hashString(principal, Charsets.UTF_8), CLIENT_BUCKETS)
```

# Search by client

```
1 Collection<OAuth2AccessToken> findTokensByClientId(String clientId) {
2     def tokens = []
3     for (int mod = 0; mod < CLIENT_BUCKETS; mod++) {
4         ResultSet rs =
5             session.execute(getAccessTokenByClient.bind(clientId, mod))
6
7         for (Row row : rs.all()) {
8             def tokenJson = row.getString("access_token")
9             tokens << objectMapper.readValue(tokenJson, OAuth2AccessToken.class)
10        }
11    }
12    return tokens
13 }
```

# Cleaning Up Mistakes

- Write two ways until you can drop old data, pairs well with TTLed data.
- Expect new queries to change how you insert the data.
- Use an external system to migrate all the data
  - Spark
  - The application itself



**SPRINGONE2GX**

WASHINGTON, DC

# Learn More. Stay Connected.



@springcentral



Spring.io/video

Starting with Cassandra make sure you know your queries!

[Richer Data History with Event Sourcing](#) by Steve Pember