



KubeCon



CloudNativeCon

Europe 2020

Virtual

Zero Downtime Data Relocation



Deepthi Sigireddi
Software Engineer

Liz van Dijk
Solutions Architect & Field Operations

Who are we?



Deepthi Sigireddi - Software Engineer

- Vitess maintainer @ PlanetScale
- @ATechGirl



Liz van Dijk - Solution Architect

- Vitess newbie, MySQL friend @ PlanetScale
- @lizztheblizz

The Problem


As governments pass data-locality laws, *jurisdiction-aware* database clusters are becoming important

Supporting data-locality often requires *re-architecting* the applications

Migrating existing databases into databases resident in multiple locations is an *operational challenge*



Vitess Architecture



APP

VITESS



Keyspaces & Shards

- **Keyspace** is an analog to what we call a *logical database*.
- Keyspaces consist of one or multiple **Shards**.
- Rows are assigned to a shard based on their computed Keyspace ID.
- Shards contain one or more replica tablets, of which one will be elected as master.



Vitess Tablet



KubeCon




CloudNativeCon


Europe 2020

Virtual

- Most basic “worker” unit of a Vitess Cluster
- **VTablet** is a *sidecar process*
- Tablets can fulfill multiple roles
 - Master (or Source)
 - Replica
 - (Analytics Replica)



- Infrastructural Failure Domain, think AZ's
- Examples:
 - a full data center
 - a subset of a data center
 - a Kubernetes cluster



- A way to compute Keyspace ID for any row in a table
- Vindex for a Table is defined by
 - Column name
 - **Sharding Function Name**
- $\text{KeyspaceIDForRow} = \text{ShardingFunction}(\text{ColumnValueForRow})$
- For example, table name: customer, sharding column: id sharding function: hash
- For a row where id is 123, $\text{KeyspaceId} = \text{hash}(123)$



Concepts: Sharding Functions



Virtual

<code>binary</code>	Identity
<code>binary_md5</code>	md5 hash
<code>hash</code>	3DES null-key hash
<code>numeric</code>	Identity
<code>numeric_static_map</code>	A JSON file that maps input values to keyspace IDs
<code>unicode_loose_md5</code>	Case-insensitive (UCA level 1) md5 hash
<code>reverse_bits</code>	Bit Reversal

Or, add your own custom sharding function!



The Solution



Virtual

Map jurisdiction to a set of keyranges, which in turn **maps a jurisdiction to a shard**

(Examples: EU, NAFTA, California)

Shard is located in a **cell bounded by jurisdiction**

Geographic **location** corresponds to a **data column** in database



Demo: Setup

- Three regions /Eight countries
 - Americas → USA, Canada
 - Europe → France, Germany
 - Asia → China, Japan, India, Indonesia
- Sharding Scheme
 - Region_Json vindex
 - Based on “multi column” vindex
 - Maps an (id, location) tuple to keyspace_id
 - Looks up region_byte for location using a map



Cluster Design

- 3 Vitess “cells” for each region
 - us-central1 → 1a, 1b, 1c
 - europe-west1 → 1b, 1c, 1d
 - asia-southeast1 → 1a, 1b, 1c
- “main” “keyspace” with three shards -40, 40-80, 80-
 - Each shard resident in three cells
- Split keyspace into 3 keyranges
 - -40 → us-central1 {US: 0x1, Canada: 0x2}
 - 40-80 → europe-west1 {France: 0x40, Germany: 0x41}
 - 80- → asia-southeast1 {China: 0x80, Japan: 0x81, India: 0xc0, Indonesia: 0xc1}





KubeCon



CloudNativeCon

Europe 2020

Virtual

Demo



The Solution



Geographic location corresponds to a data column in database

Lookup Vindex on primary key column

Primary Key can include encoded location information (Custom Vindex)



The Vitess Schedule



Virtual

Tuesday, 18 August

- **14:30 CEST**
 - Vitess Intro: How to Migrate a MySQL Database to Vitess - Sugu Sougoumarane & Morgan Tocker, PlanetScale
- **16:21 CEST**
 - Sponsored Keynote: Network, Please Evolve: Chapter 3, Stretching Out – Vijoy Pandey, Vice President and CTO of Cloud, Cisco

Wednesday, 19 August

- **14:30 CEST**
 - Zero Downtime Data Relocation with Vitess - Liz van Dijk & Deepthi Sigireddi, PlanetScale
- **16:55 CEST**
 - Vitess Deep Dive: Deployment Design - Dan Kozlowski & Andres Taylor, PlanetScale
 - Multicloud Vitess over Network Service Mesh - Tim Swanson, Cisco & John Watson, PlanetScale





KubeCon



CloudNativeCon

Europe 2020

Virtual

Questions?

Deepthi Sigireddi

@ATechGirl

Liz van Dijk

@lizztheblizz

 planetscale
@vitessio @planetscaledata

