

# CrateDB

**Kryštof Hrubý**



# Database Model

- SQL relational database
- Search engine
- Document store
- Time Series database



# SQL relational database

- Supports most SQL commands
  - Joins, Aggregation, Sort, ...
- Built as distributed database from the start
  - Uses Apache Lucene, Elasticsearch, Netty
- Each table split into shards
  - Shards are distributed uniformly across cluster
  - Replication factor can be chosen
- Eventual Consistency
- No support for transactions
  - Any operation on a row is atomic
- Write-ahead logging



# Search Engine

- Uses Apache Lucene
  - Storage, indexing, text and geospatial search
- Provides full-text and geospatial search



# Document Store

- Each table row is a semi structured document
- Document – nested structure of object and array types
- JSON data can be loaded to either Object or Array type
- Operations on documents are atomic
- Object similar to JSON Object



# Query – Create Table



```
1  INSERT INTO products (  
2    id,  
3    name,  
4    description,  
5    specification,  
6    cost,  
7    inStock  
8  ) VALUES (  
9    1,  
10   'Fieldmann FZR 70335-A 2x20V',  
11   'Akumulátorová sekačka na trávu značky FIE  
    LDMANN...',  
12   {  
13     "length" = 130.4,  
14     "width" = 50.4  
15   },  
16   3399.90,  
17   10  
18  );
```



```
1  CREATE TABLE products (  
2    id INT PRIMARY KEY,  
3    name TEXT,  
4    description TEXT,  
5    specification OBJECT(DYNAMIC) AS (  
6      length FLOAT,  
7      width FLOAT  
8    ),  
9    cost FLOAT,  
10   inStock INT,  
11   reviews ARRAY(TEXT),  
12   INDEX productNameFt USING FULLTEXT(name)  
13  ) CLUSTERED INTO 10 SHARDS;
```



# Query – Delete, Update



```
1 DELETE FROM products
2 WHERE id = 0;
```



```
1 UPDATE products SET
2   specification['length'] = 100.0
3 WHERE id = 1;
```



# Query – Text search



```
1 SELECT name, inStock, _score
2 FROM products
3 WHERE MATCH(productNameFt, 'Fieldmann -A')
4 ORDER BY _score DESC;
```

<b>name</b>	<b>inStock</b>	<b>_score</b>
Fieldmann FZR 70335-A 2x20V	10	0.26152915
Fieldmann FZR 70435-0	3	0.13076457





# Query – Joins



```
1 CREATE TABLE orders (  
2   id INT PRIMARY KEY,  
3   customer TEXT  
4 );  
5  
6 INSERT INTO orders (  
7   id,  
8   customer  
9 ) VALUES (0, 'Jan Novak'), (1, 'Katerina Novotna');
```



```
1 CREATE TABLE productsOrders (  
2   productId INT,  
3   orderId INT,  
4   count INT,  
5   PRIMARY KEY (productId, orderId)  
6 );  
7  
8 INSERT INTO productsOrders (  
9   productId,  
10  orderId,  
11  count  
12 ) VALUES (0, 0, 1), (1, 0, 2), (2, 1, 3);
```



# Query – Joins



```
1 SELECT o.customer, SUM(po.count * p.cost) AS paid
2 FROM products p
3 JOIN productsOrders po ON p.id = po.productId
4 JOIN orders o ON po.orderId = o.id
5 GROUP BY o.customer;
```

<b>name</b>	<b>paid</b>
Katerina Novotna	16199.699
Jan Novak	11299.699



# Advantages, Disadvantages

- ✓ SQL language with objects
  - It might lead user to use slow joins
- ✓ Easy to setup
- ✓ Provides free client with db management, console, ...
- ✓ Nice documentation
  - Although sometimes not clear what is supported and what not
- × Full-text indices can be created only when creating table
- × When data are inserted, subsequent select may not find them
  - × Even for small cluster with one node
- × Weird float arithmetics



**Thanks for your attention**

