




Guerrilla Tactics for Scalable E-commerce Services(React, Spring, NoSQL)

with Apache Cassandra® and Apache Pulsar®




What's the connection?

- E-commerce?
- Guerilla tactics?
- Cyberpunk 2077?



DBRE/Developer Advocate DataStax



Aaron Ploetz





@aploetz / @aar0np





@aploetz

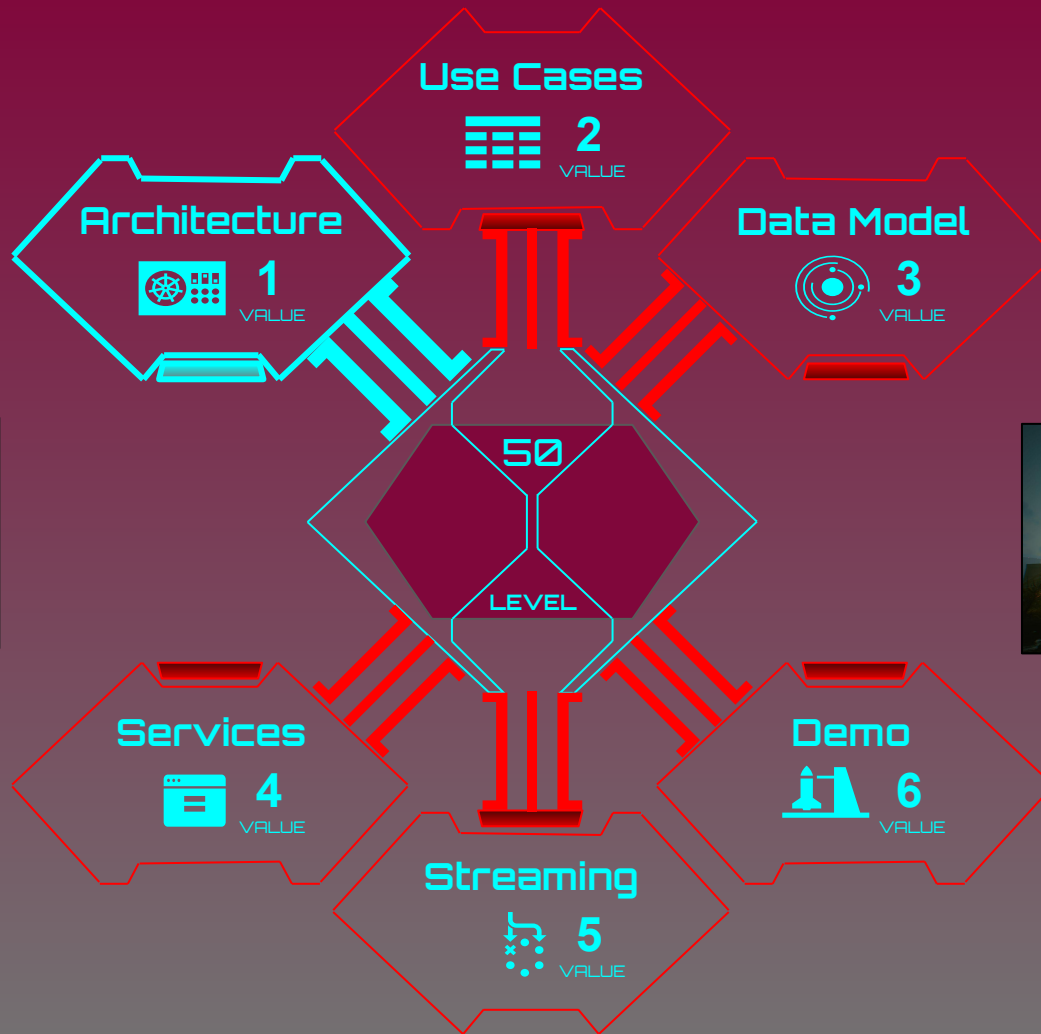


@aaronploetz

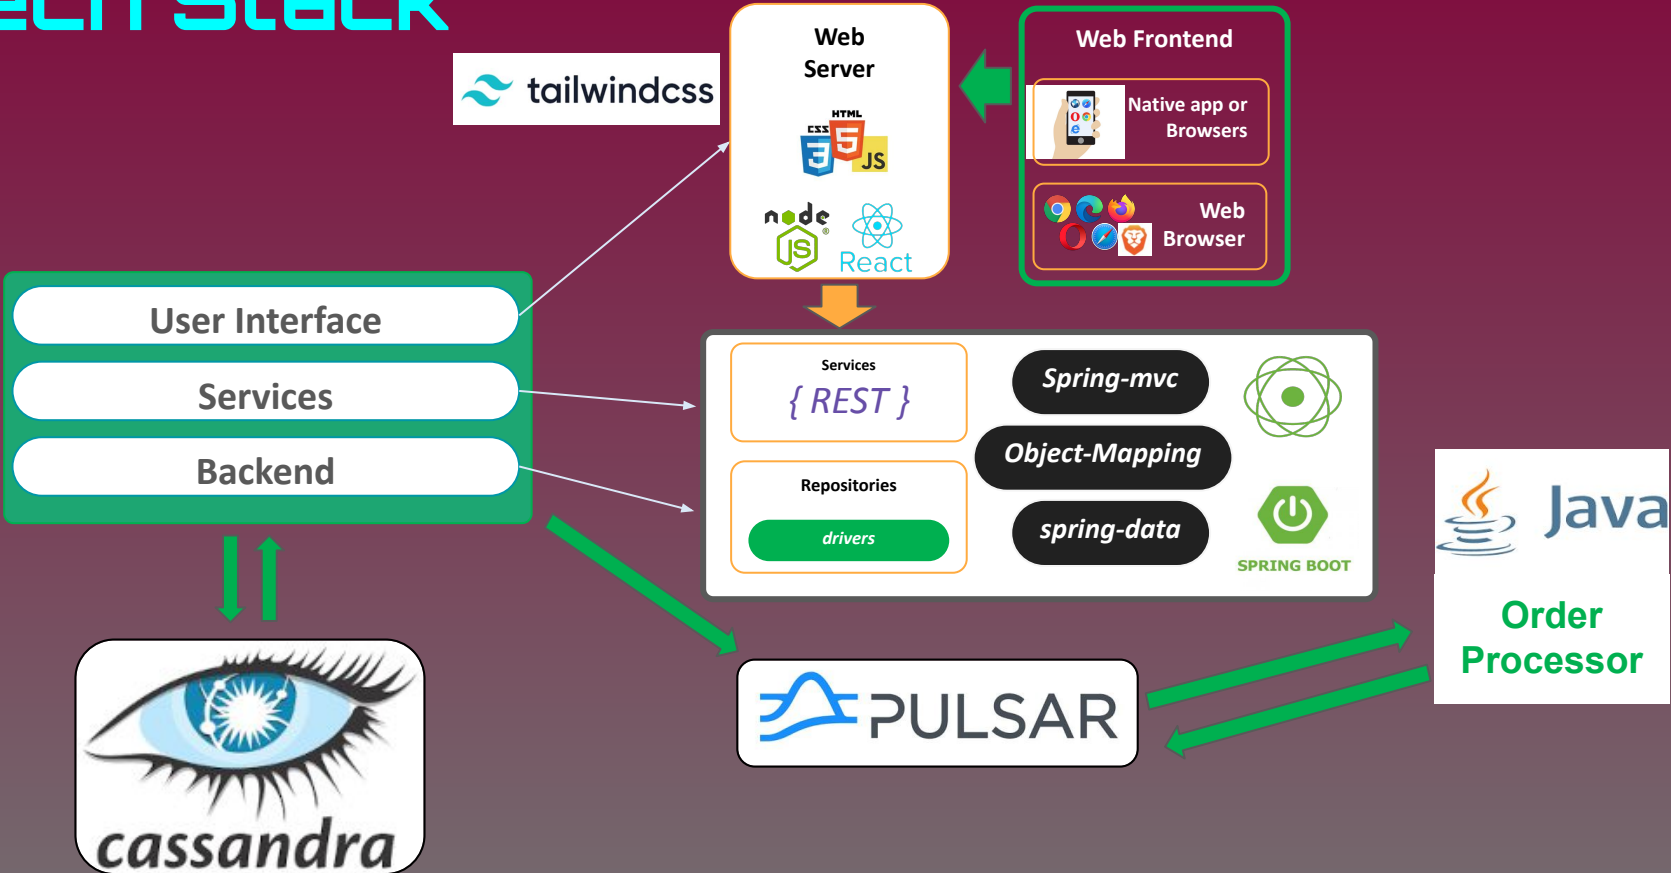
- Former SWE/DevOps/DB Lead @  & 
- Host - Apache Cassandra Corner podcast
- Cassandra® MVP
- Worked as an author on:
 - Mastering Apache Cassandra 3.x
 - Seven NoSQL Databases in a Week




Agenda



Tech Stack



Apache Cassandra



- Read / Write Performance
- Linear Scalability
- High Availability
- Geographical Distribution
- Platform Agnostic



Spring Data Cassandra




- Easy access to connection properties.
 - `application.yml`
- Reduces boilerplate code.
- Provides useful CRUD methods:
 - `save(entity)`
 - `findById(id)`
 - `delete(id)`



Spring Data Cassandra

- `count()`
- `findAll()`
- `saveAll(Iterable<>)`

**Many
containers
died to bring
us this
information.**



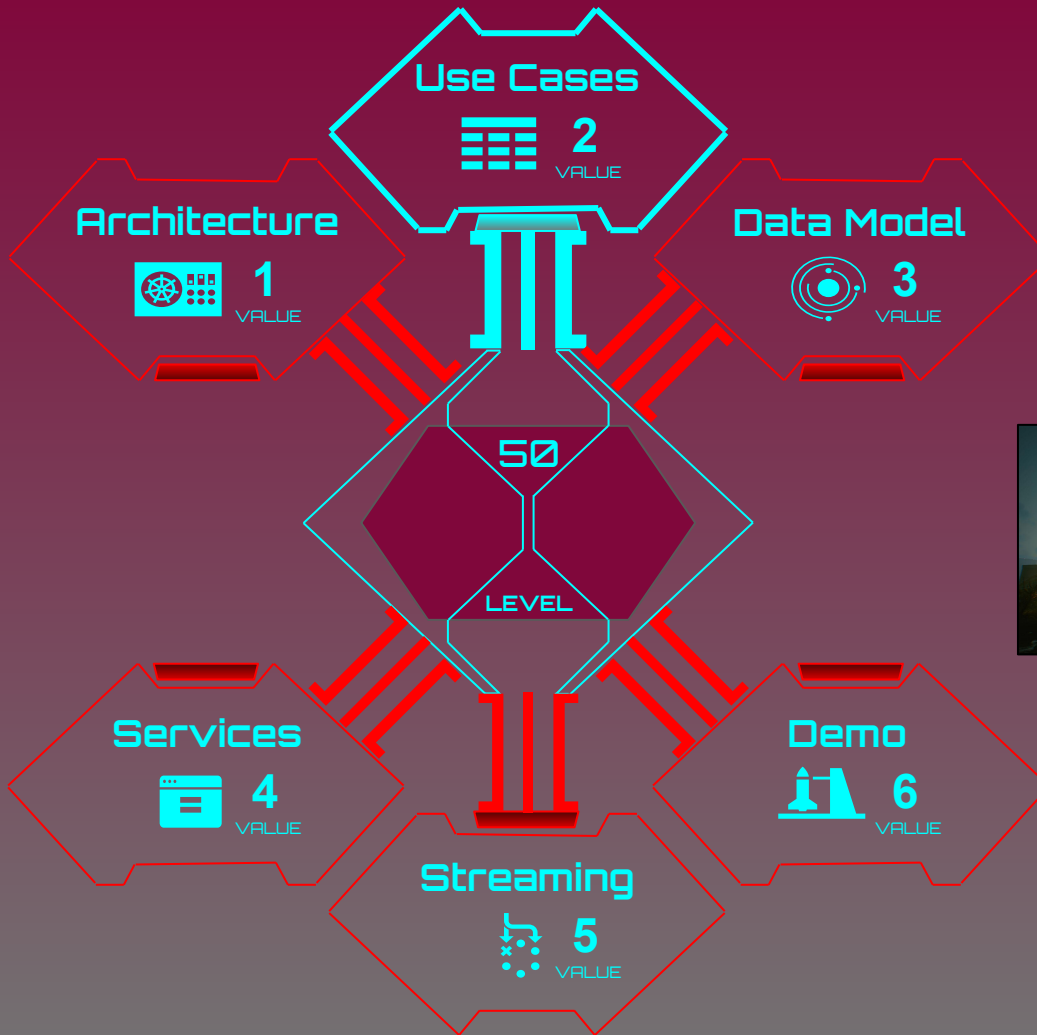
Apache Pulsar



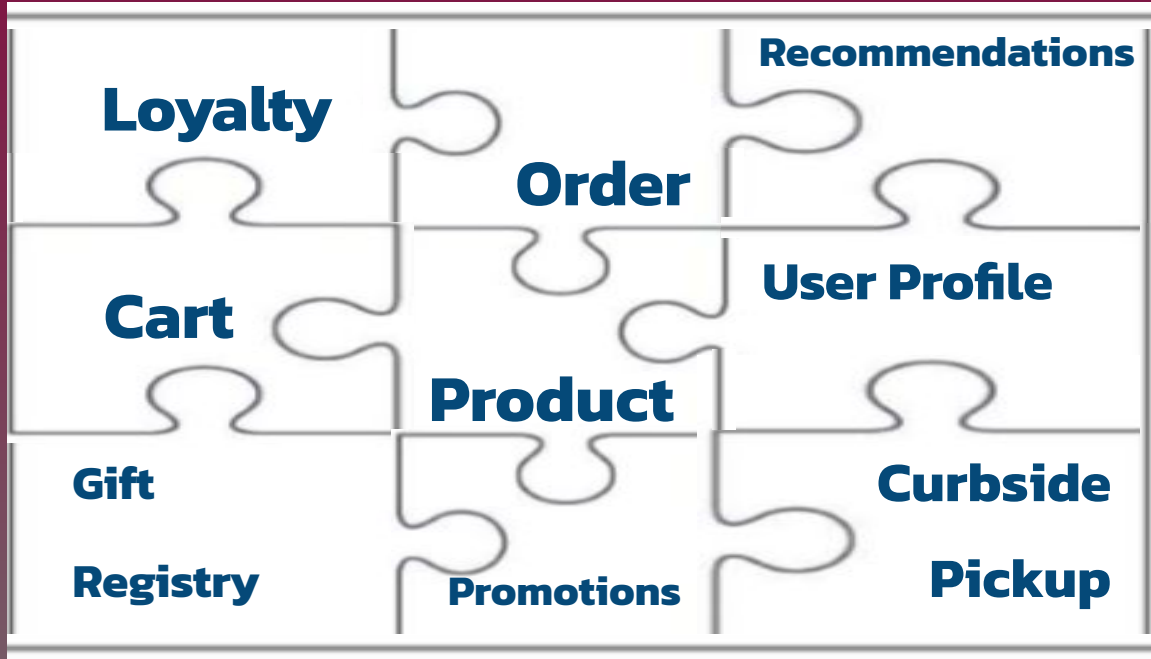
- Pub-sub, queue, stream – All in one!
- Designed with the cloud in mind.
- Geographic awareness.
- Highly scalable
- automatic rebalancing for partitioned topics
- Lightweight Pulsar function library



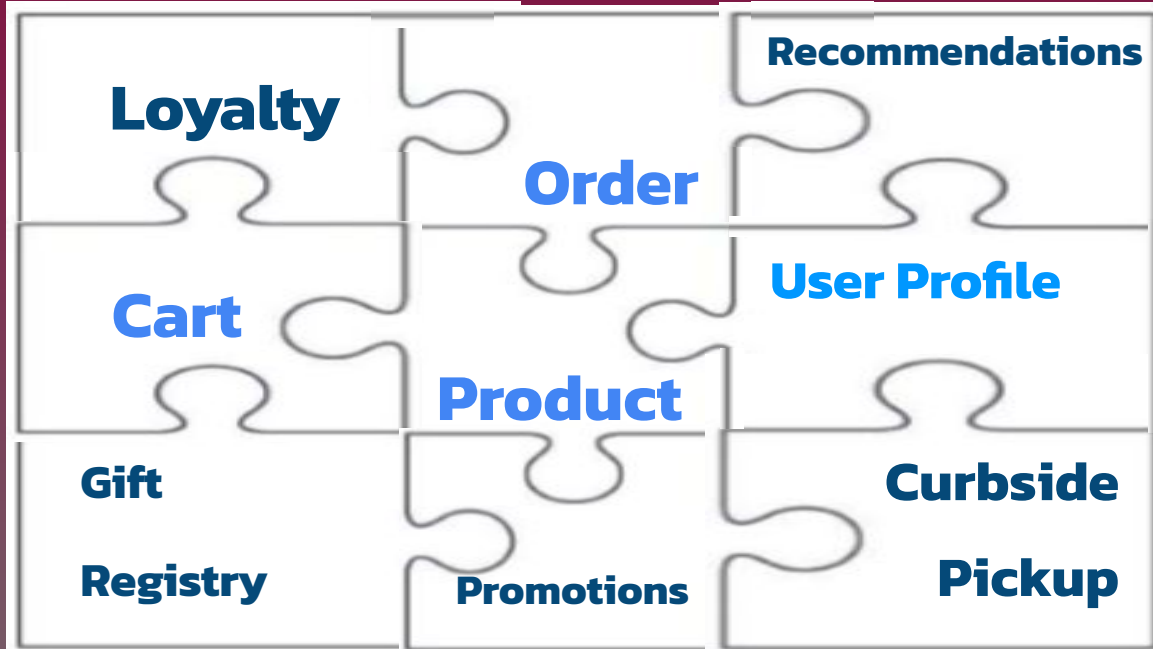
Agenda



E-commerce Subsystems

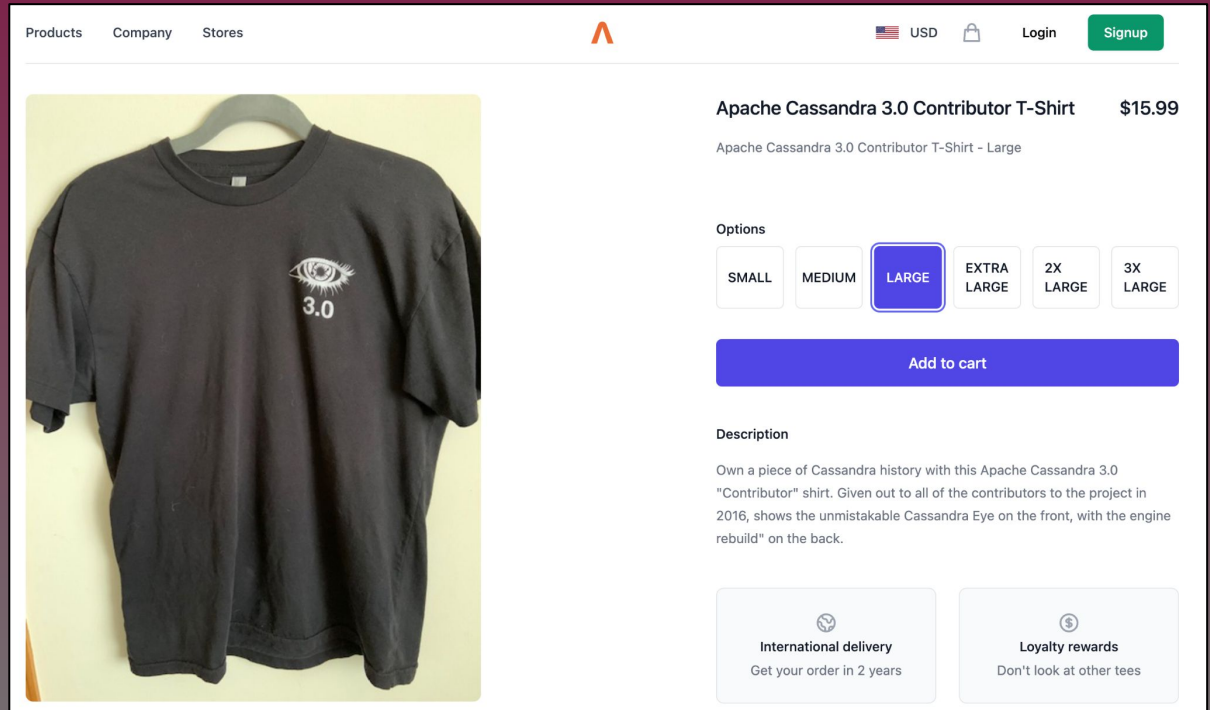


E-commerce Subsystems



Product Catalog

- Category (Navigation)
- Product (Data)
- Pricing



The screenshot shows a product page for an Apache Cassandra 3.0 Contributor T-Shirt. The page layout includes a navigation bar with links for 'Products', 'Company', and 'Stores', a logo, and utility links for 'USD', 'Login', and 'Signup'. The product image shows a black t-shirt with a white 'Cassandra Eye' logo and the number '3.0'. The product title is 'Apache Cassandra 3.0 Contributor T-Shirt' with a price of '\$15.99'. Below the title, there are size options: 'SMALL', 'MEDIUM', 'LARGE' (selected), 'EXTRA LARGE', '2X LARGE', and '3X LARGE'. A blue 'Add to cart' button is present. The description states: 'Own a piece of Cassandra history with this Apache Cassandra 3.0 "Contributor" shirt. Given out to all of the contributors to the project in 2016, shows the unmistakable Cassandra Eye on the front, with the engine rebuild on the back.' At the bottom, there are two promotional boxes: 'International delivery' (Get your order in 2 years) and 'Loyalty rewards' (Don't look at other tees).

Products Company Stores

USD Login Signup

Apache Cassandra 3.0 Contributor T-Shirt \$15.99

Apache Cassandra 3.0 Contributor T-Shirt - Large

Options

SMALL MEDIUM **LARGE** EXTRA LARGE 2X LARGE 3X LARGE

Add to cart

Description

Own a piece of Cassandra history with this Apache Cassandra 3.0 "Contributor" shirt. Given out to all of the contributors to the project in 2016, shows the unmistakable Cassandra Eye on the front, with the engine rebuild on the back.



International delivery
Get your order in 2 years

Loyalty rewards
Don't look at other tees

Shopping Cart

- User Cart Data
- Cart Products

Shopping Cart


	DataStax Gray Track Jacket gray Large ✓ In stock	\$44.99 Remove
	Apache Cassandra 3.0 Contributor T-Shirt black Large ✓ In stock	\$15.99 Remove
Subtotal		\$60.98

Shipping and taxes will be calculated at checkout.




[Checkout](#)

User Profile

- User Data
- Single Sign-On


Products Company Stores 

Sign in to your account

   Logout

Personal Information

Use a permanent address where you can receive mail.

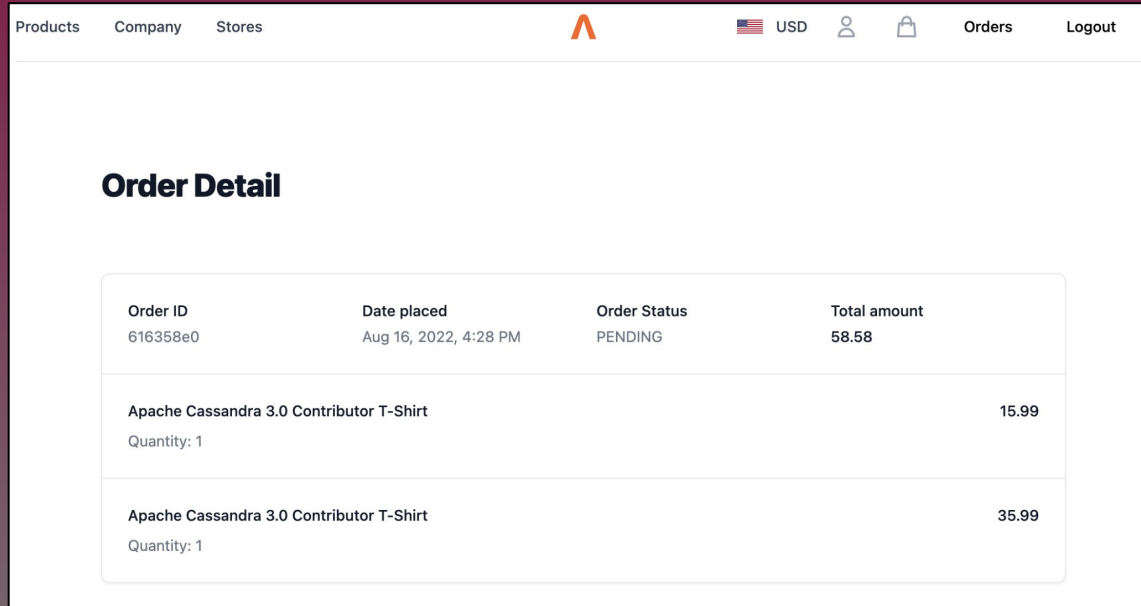
First name  Last name

Email address

Country

Ordering System

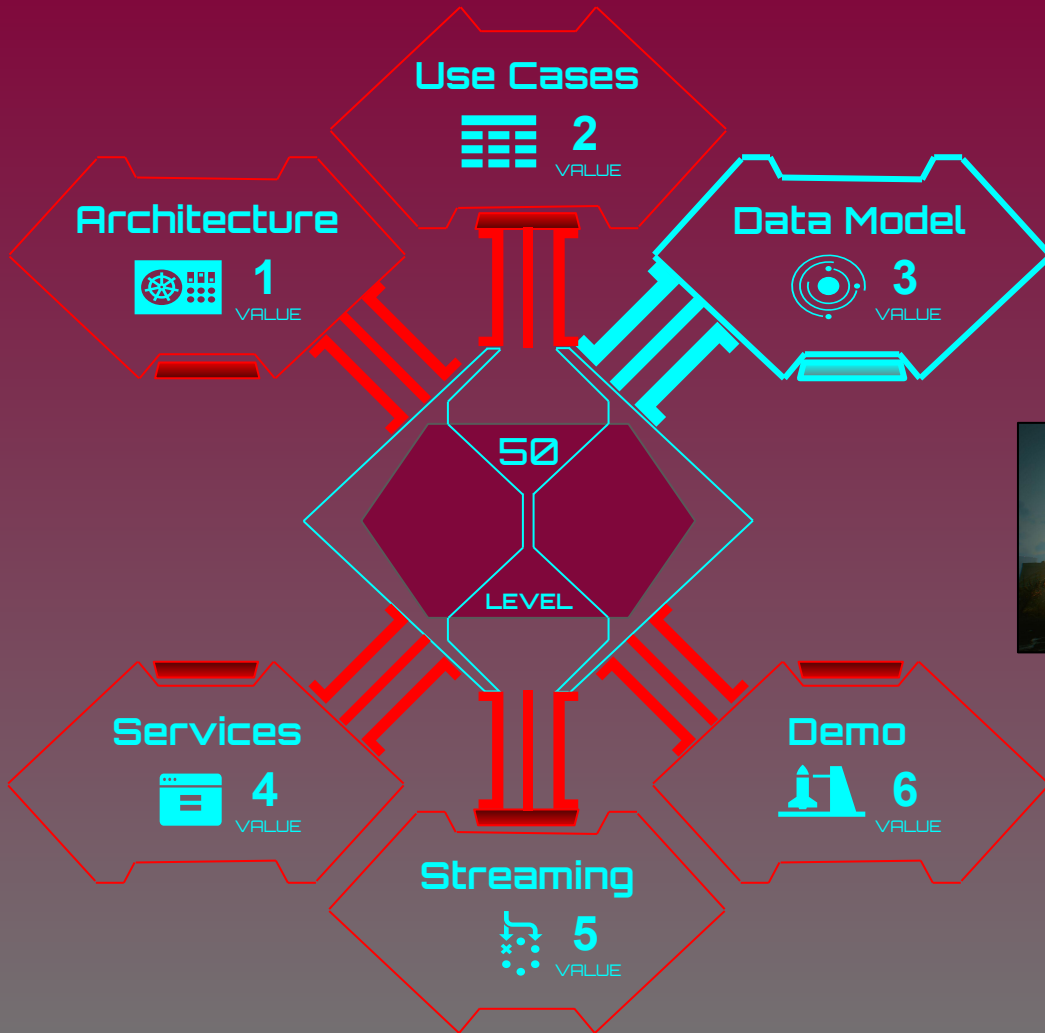
- Order Data Processing
- Order History



The screenshot displays a web application interface for an ordering system. At the top, there is a navigation bar with links for 'Products', 'Company', and 'Stores' on the left, and 'Orders' and 'Logout' on the right. A central logo is also present. Below the navigation bar, the main content area is titled 'Order Detail'. It features a table with the following data:

Order ID	Date placed	Order Status	Total amount
616358e0	Aug 16, 2022, 4:28 PM	PENDING	58.58
Apache Cassandra 3.0 Contributor T-Shirt Quantity: 1			15.99
Apache Cassandra 3.0 Contributor T-Shirt Quantity: 1			35.99

Agenda





Cassandra Data Modeling

Tables built to suit a query

Pros:

- Fast reads
- Simple query model

Cons:

- Data duplication (*but that's ok*)
- Manual integrity enforcement

Employees

userId	firstName	lastName	department
1	Edgar	Codd	Engineering
2	Raymond	Boyce	Engineering
3	Sage	Lahja	Math
4	Juniper	Jones	Botany

Data Modeling "Cassandra Style"

Cassandra Primary Keys



- PRIMARY KEY ((partitionK1,partitionK2),clusteringK1,clusteringK2);

Partitioning Key

- Determines which (or the cluster) the data is stored.
- Requires for proper query routing.



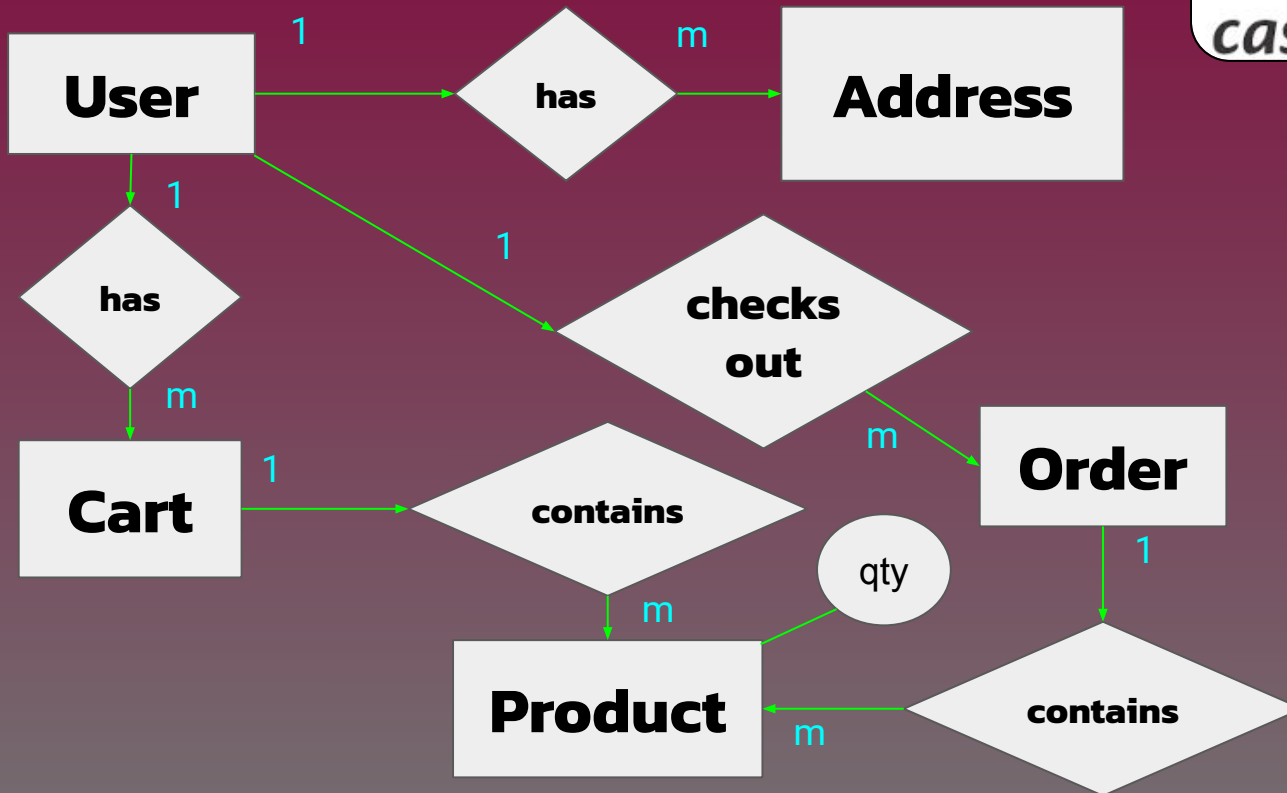


Tips for **Large Scale**

- Data **queried** together should be **stored** together.
- Use **high-cardinality** key values.
- **No** full table scans!
- Keep things **small!**
 - Partitions
 - Result sets



Data Modeling "Cassandra Style"



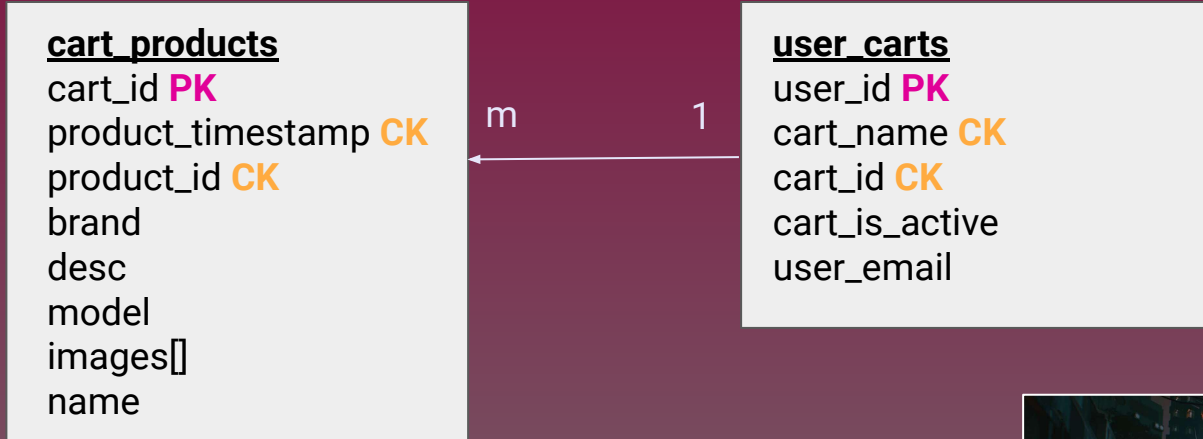


Queries

- Need to be able to query a user's "active" cart.
- Need to query all carts for a user.
- User needs to view all products in their cart.
- User needs to add a product to their active cart.
- User needs to remove a product from their active cart.



Data Model - Queries to Support



Partition Key **PK**
Clustering Key **CK**



Data Model - Physical

Cart Products



- Solving the cart query.
- Clustering on timestamp.
- Enforcing a TTL.

```
CREATE TABLE cart_products (  
    cart_id uuid,  
    product_timestamp timestamp,  
    product_id text,  
    product_description text,  
    product_name text,  
    quantity int,  
    PRIMARY KEY (cart_id,  
product_timestamp, product_id)  
    ) WITH CLUSTERING ORDER BY  
    (product_timestamp DESC, product_id ASC)  
    AND default_time_to_live = 5184000;
```




Anti-bot/malware measures:

- 60 day TTL (time to live).

```
default_time_to_live = 5184000
```

Old carts in the DB are avenues of attack!

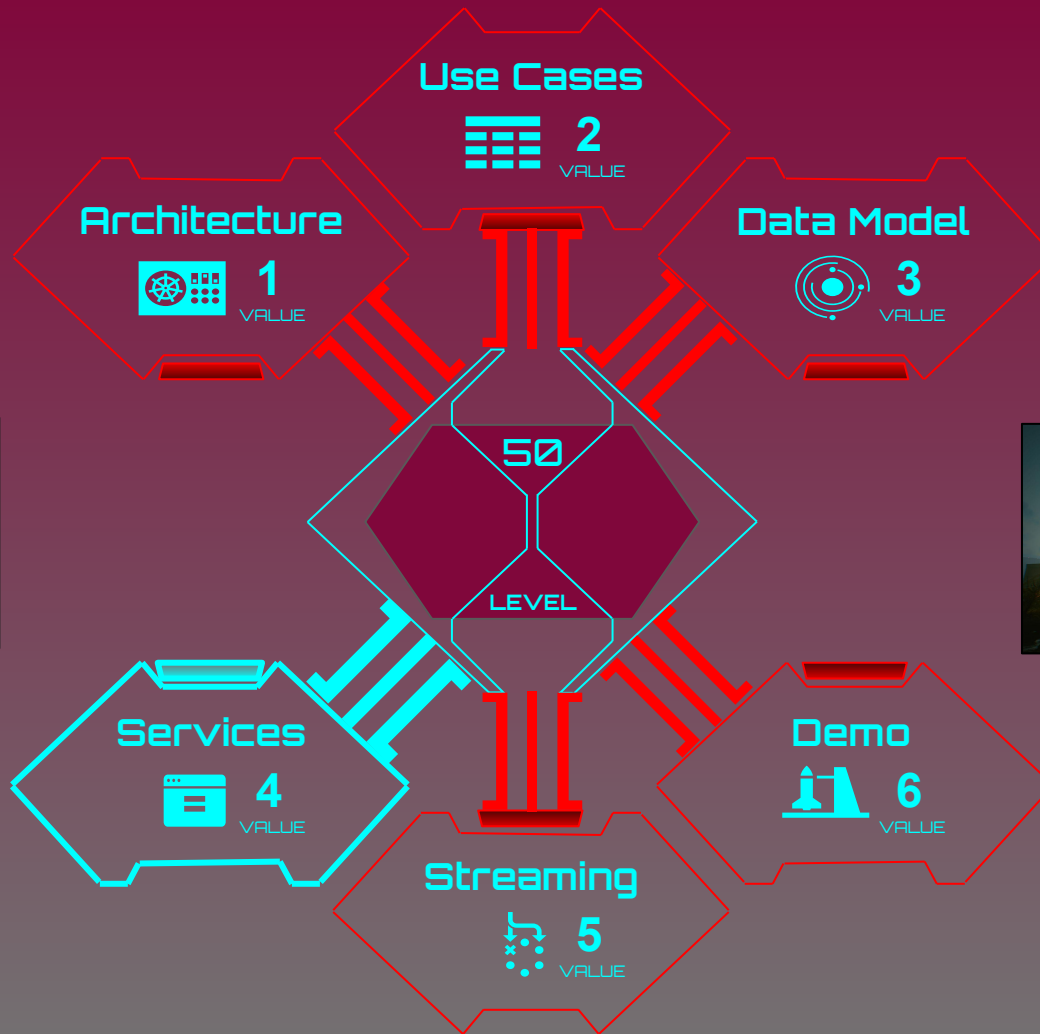
- Rate limiter on product-add service.

Bot attacks will add many, many products in a matter of seconds!



Shopping Cart - Special Considerations

Agenda



User Data



<http://localhost:8080/swagger-ui/index.html?configUrl=/v3/api-docs/swagger-config#/>

GET	<code>/api/v1/user/user</code>	Retrieve user by 3rd party login (Google, GitHub, etc)
GET	<code>/api/v1/user/user/{userid}</code>	Retrieve user by user_id
GET	<code>/api/v1/user/email/{email}</code>	Retrieve user by email

Service Endpoint - User Data GETs

User by Email

- A “manual” index.
- Returns a `user_id` for an email address.

```
CREATE TABLE user_by_email (  
    user_email TEXT PRIMARY KEY,  
    user_id UUID,  
);
```

User By Email CQL:



```
> SELECT * FROM user_by_email WHERE user_email='aaronploetz@gmail.com';
```

```
user_email          | user_id
```

```
-----+-----
```

```
  aaronploetz@gmail.com | dce3d828-4d27-40da-b48b-ef5096d1b113
```

```
(1 rows)
```

Service Endpoint - User By Email

getUserByEmail:

```
public ResponseEntity<User> getUserByEmail(HttpServletRequest req,
    @PathVariable(value = "email")
    @Parameter(name = "email", description = "email address", example = "bob.slydell@bobs.com")
    String email) {

    Optional<UserByEmailEntity> userByEmail = userByEmailRepo.findById(email);

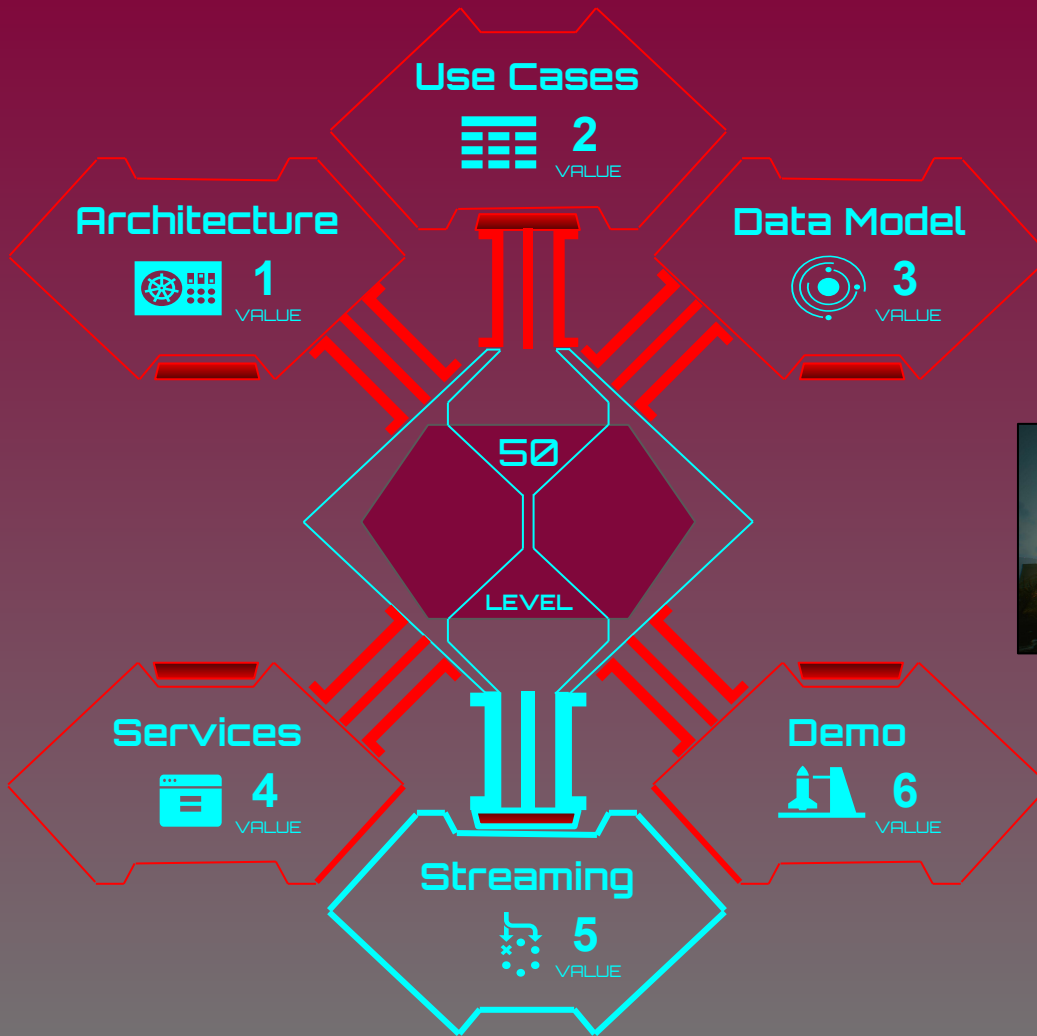
    if (userByEmail.isPresent()) {
        Optional<UserEntity> user = userRepo.findById(userByEmail.get().getUserId());

        if (!user.isPresent()) {
            return ResponseEntity.notFound().build();
        }

        return ResponseEntity.ok(mapUser(user.get()));
    } else {
        return ResponseEntity.notFound().build();
    }
}
```

Service Endpoint – User By Email

Agenda



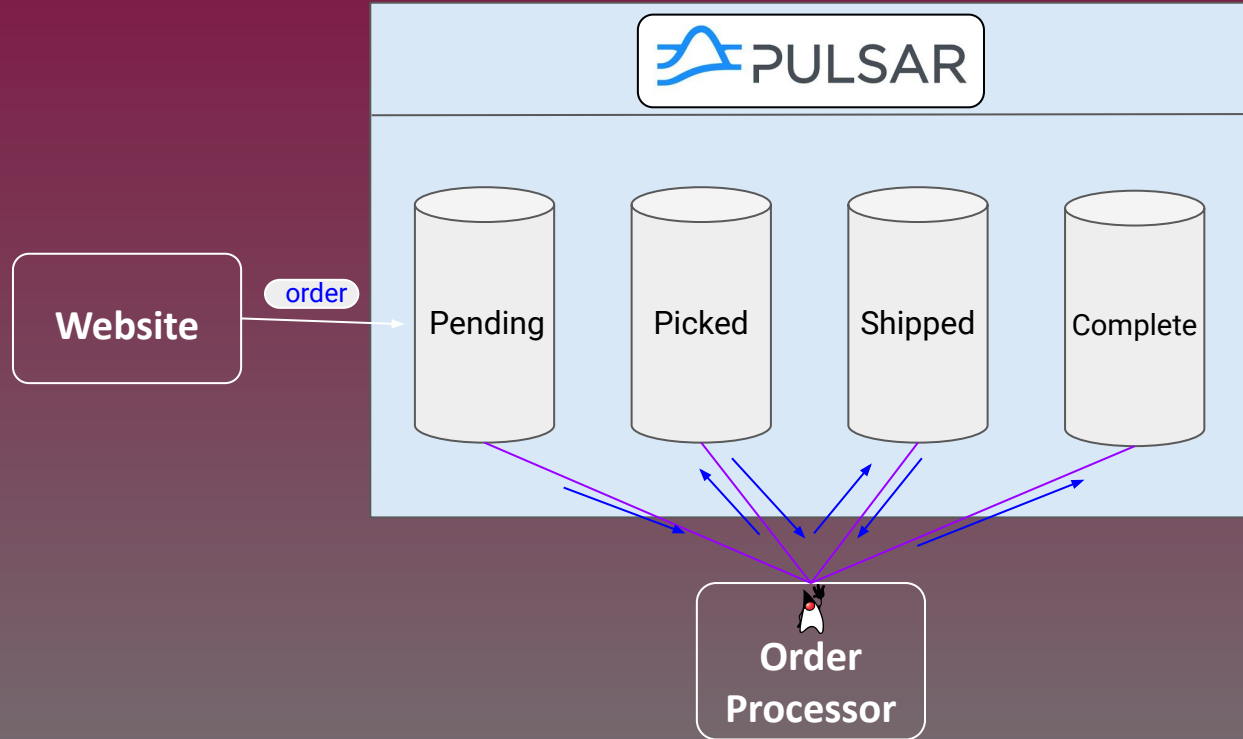
Messaging Requirements




- Place a user's order.
- Allow business functions on an order.
- Track the status timeline of an order.



Order Processing System



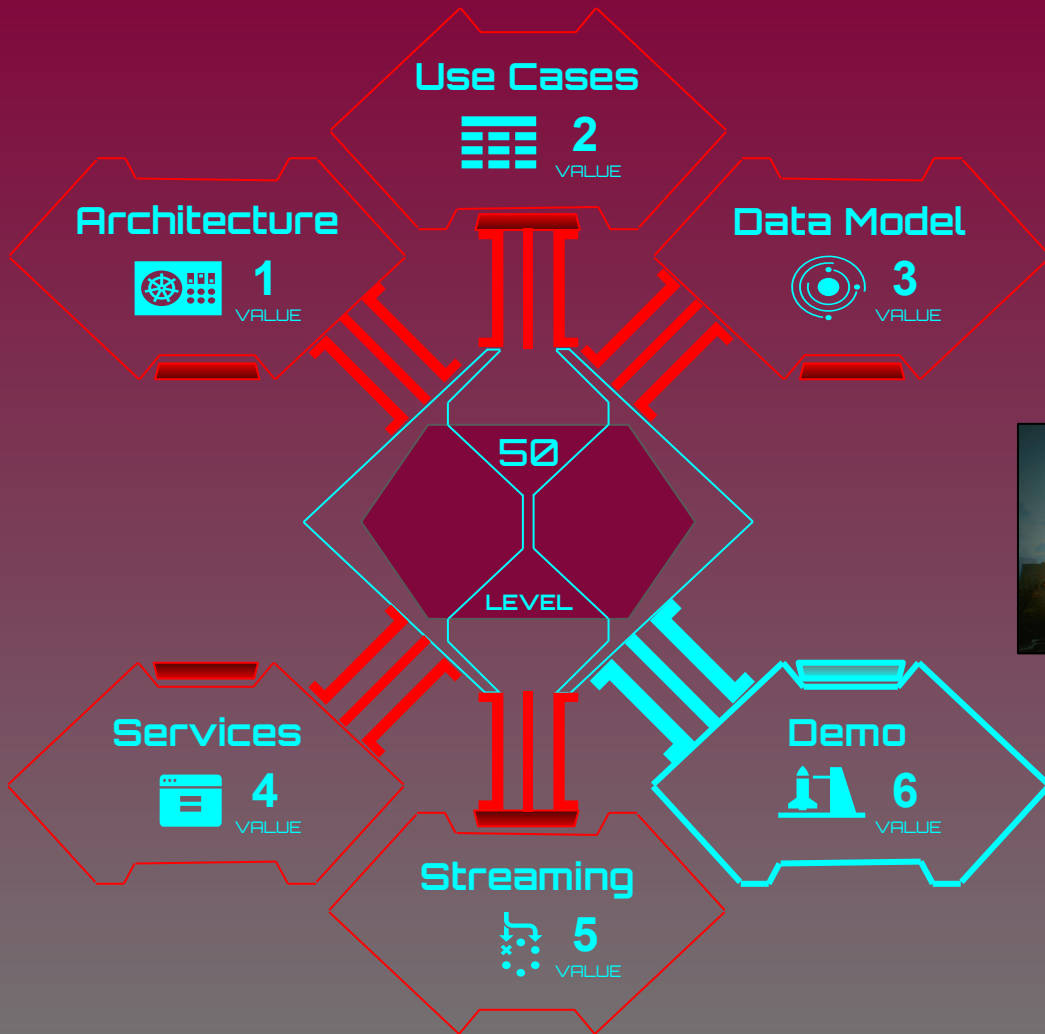
Events / Actions



- **Message ordering.**
- **Message delivery (guaranteed, exactly one – especially for payment).**
- **Scalability (high volumes, handling spikes, backpressure).**
- **Using the Publish–Subscribe pattern (Pub/Sub) for sending and receiving messages.**



Agenda



Demo

- Sign-in
- Navigate products
- Add to cart
- Place order
- Process order




Tips:

Database Migration

- Use BigDecimal or BigInteger type!
- Be sure to use software packages with invoice certificate!

<https://0.30000000000000004.com/>



Password Hashing

- Make sure that the password is **NEVER** stored in plain text!
- Storage:

```
BCryptPasswordEncoder pEncoder = new BCryptPasswordEncoder();  
String hashedPassword = pEncoder.encode(userData.getPassword());
```

- Comparison:

```
BCryptPasswordEncoder pEncoder = new BCryptPasswordEncoder();  
if (pEncoder.matches(rawPassword, hashedPassword)) {
```




CASSANDRA SUMMIT

MARCH 13-14, 2023 • SAN JOSE, CA

————— **SAVE THE DATE** —————

McEnergy Convention Center
San Jose, CA + Virtual

Resources



<https://github.com/datastaxdevs/workshop-ecommerce-app/>

- <https://www.youtube.com/c/DataStaxDevs>



<https://cassandra.apache.org/>

- <https://pulsar.apache.org/>



<https://discord.gg/c5NyPYXS>

DataStax



Thank you!

