

Ever **changing**
data model -
Schema management
for the future

Vidya Venugopal
@venuvid 

Software engineer

Founder of thekofkanerd.io

Spent a lot of time on Enterprise Integration

Platforms & Distributed systems

Now, full on with Event Streaming

Senior Dev, Vanguard Australia.



Agenda

The problem
space

Intro to
Schemas

Era of Web
Services

Stepping into
Event-driven
Architecture

Rise of Event
Streaming

Schema
management -
way forward

**The problem
space**



Intro to
Schemas



Era of Web
Services



Stepping into
Event-driven
Architecture



Rise of Event
Streaming

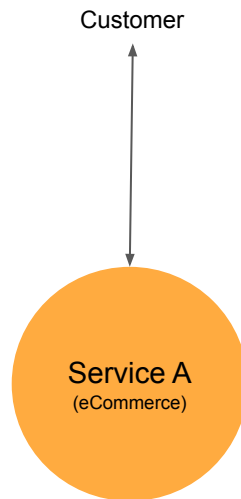


Schema
management -
way forward



The need for Schemas

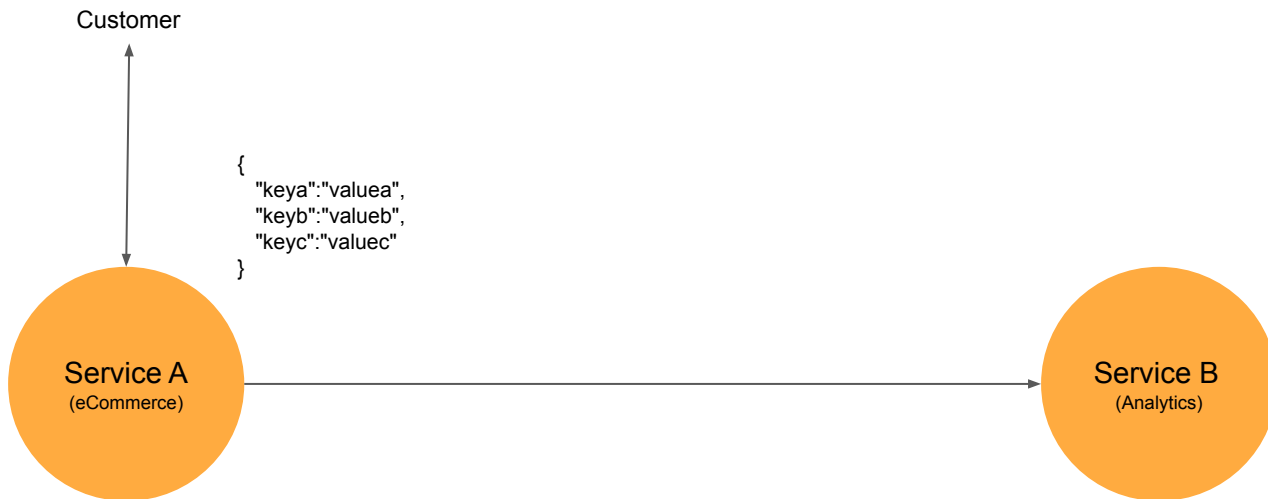
The need for schemas



The need for schemas



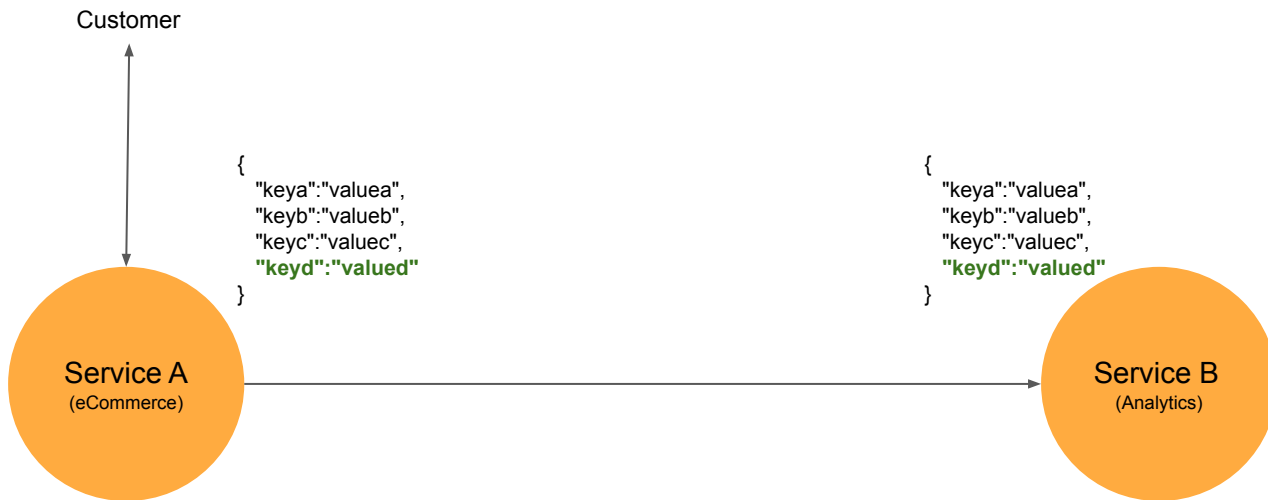
The need for schemas



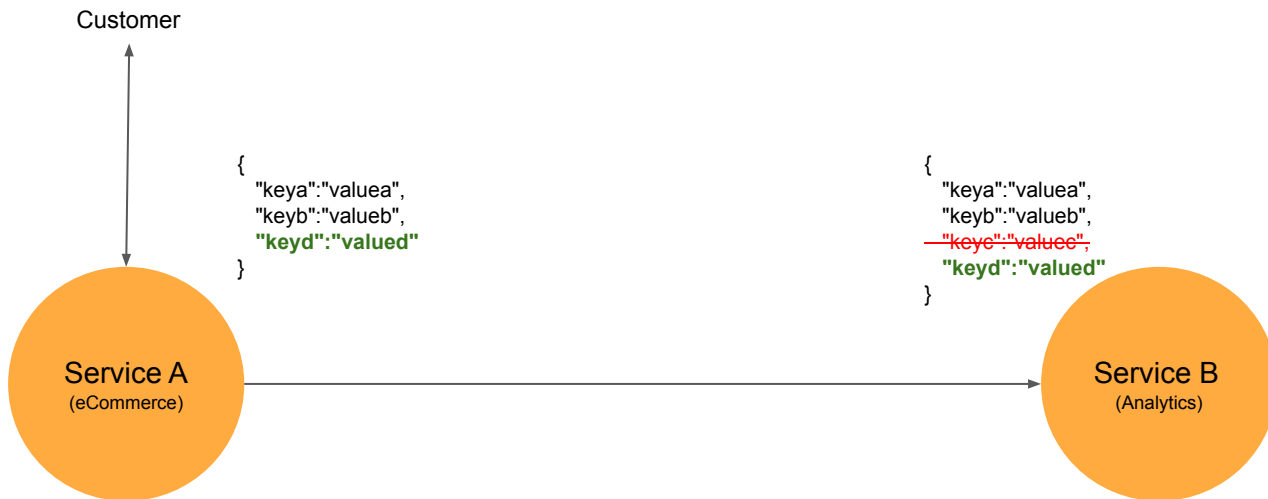
The need for schemas



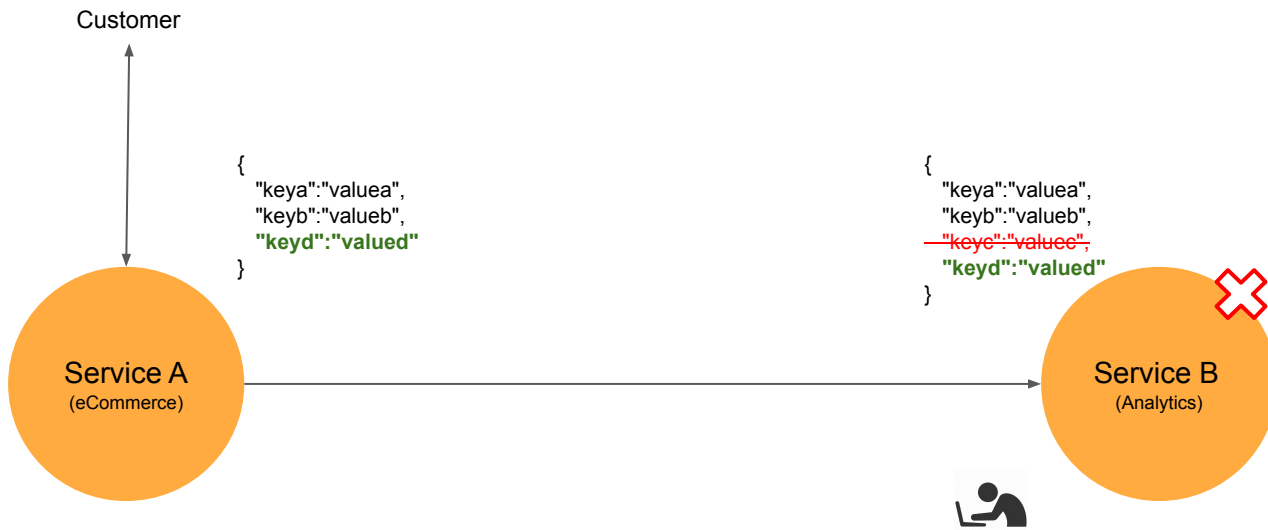
The need for schemas



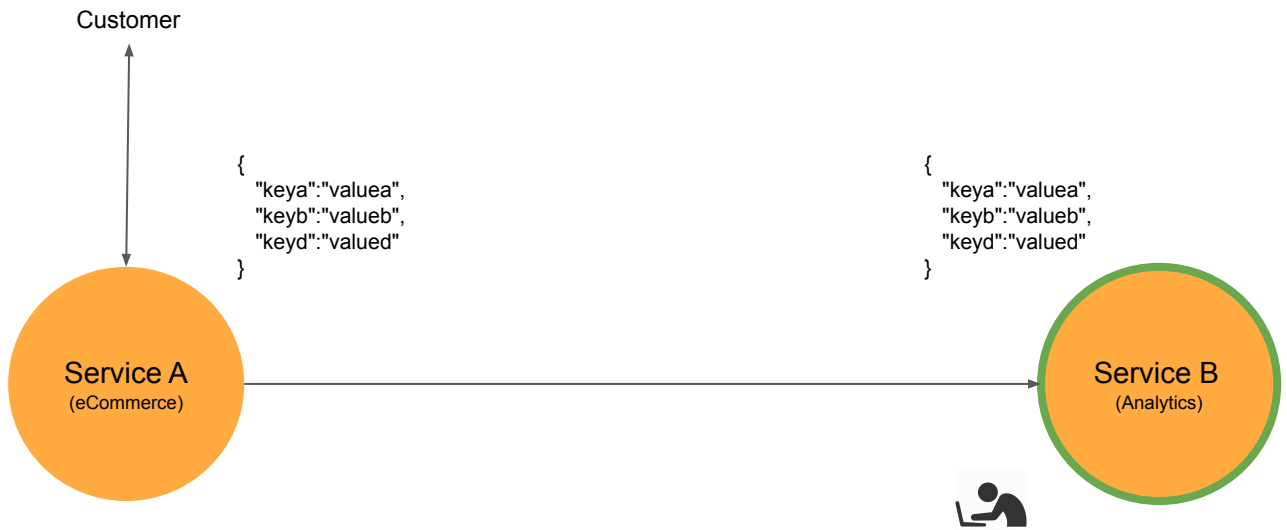
The need for schemas



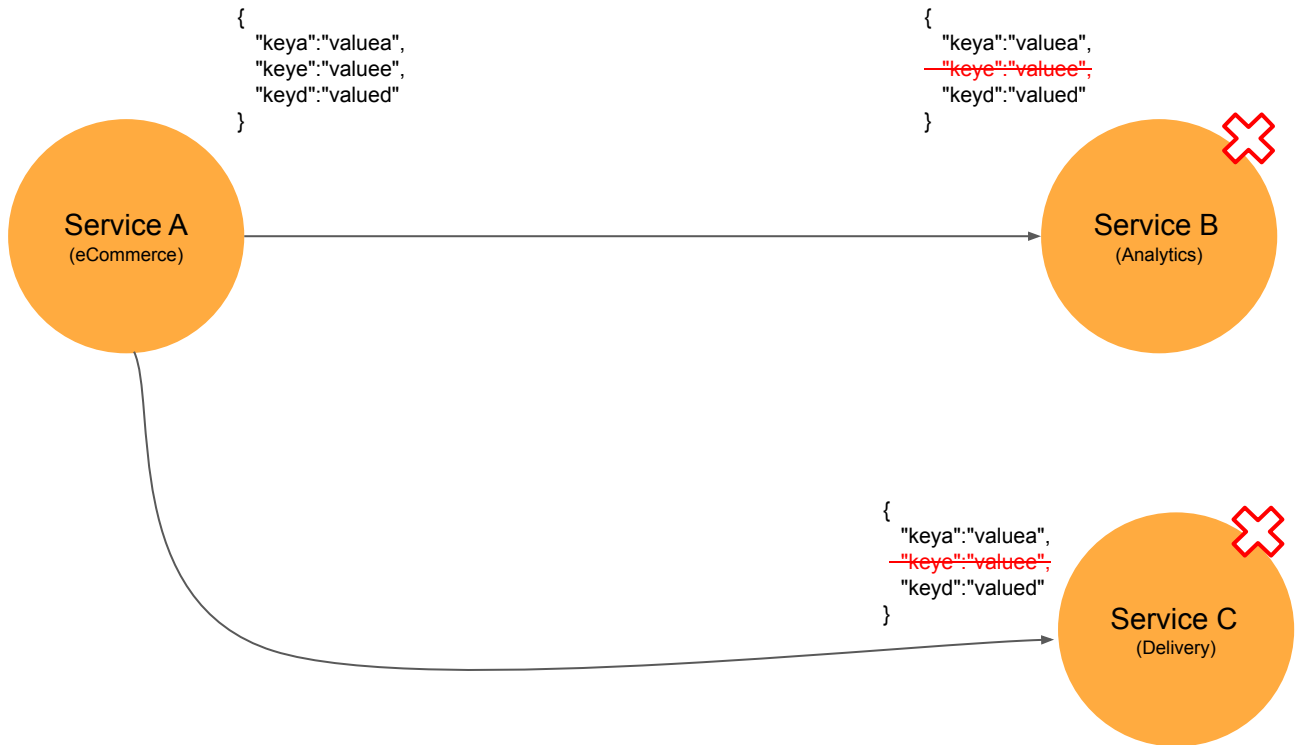
The need for schemas



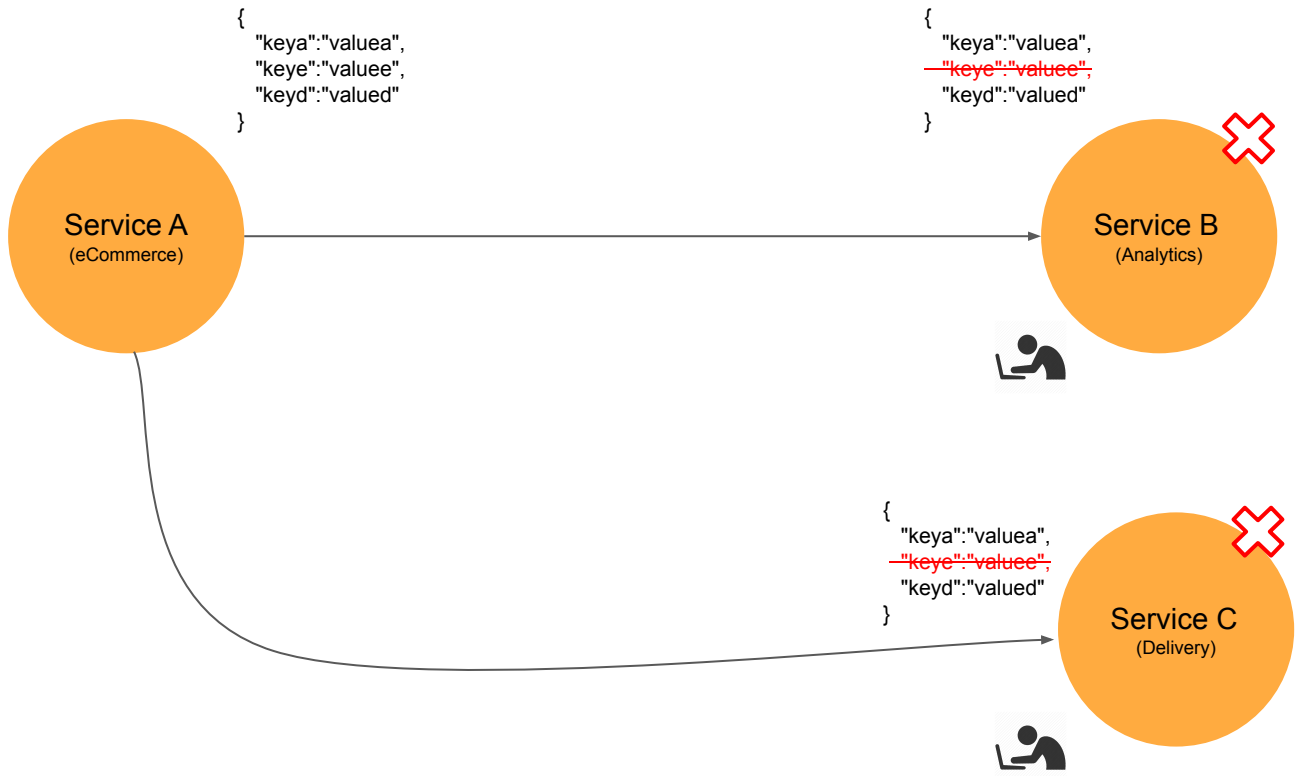
The need for schemas



The need for schemas



The need for schemas





The problem
space

**Intro to
Schemas**

Era of Web
Services

Stepping into
Event-driven
Architecture

Rise of Event
Streaming

Schema
management -
way forward

Schemas to the rescue

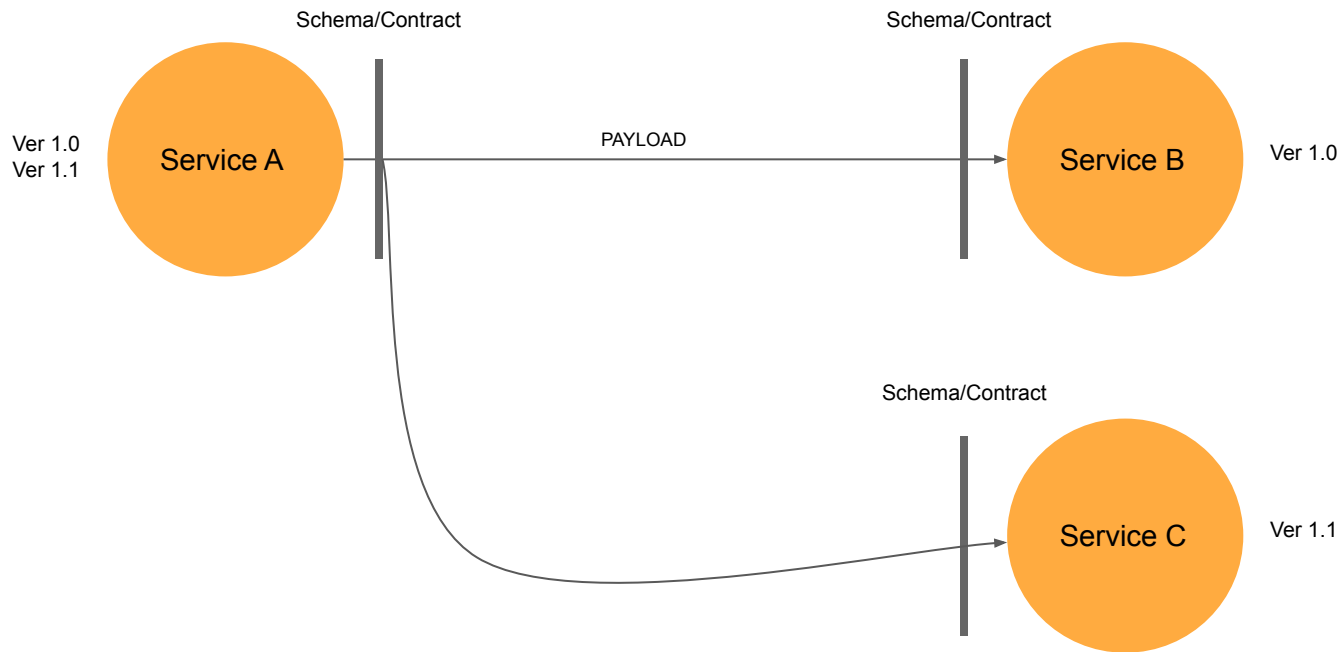


Schemas to the rescue (Examples)

```
person-schema.xsd X
Users > vidya > Desktop > person-schema.xsd
1 <xs:schema xmlns:xs="http://www.w3.org/2001/XMLSchema">
2   <xs:element name="person" type="personType"/>
3   <xs:complexType name="personType">
4     <xs:sequence>
5       <xs:element type="xs:string" name="firstName"/>
6       <xs:element type="xs:string" name="lastName"/>
7       <xs:element type="xs:integer" name="age"/>
8     </xs:sequence>
9   </xs:complexType>
10 </xs:schema>
```

```
person-schema.json •
Users > vidya > Desktop > person-schema.json > ...
1 {
2   "$id": "https://example.com/person.schema.json",
3   "$schema": "http://json-schema.org/draft-07/schema#",
4   "title": "Person",
5   "type": "object",
6   "properties": {
7     "firstName": {
8       "type": "string",
9       "description": "The person's first name."
10    },
11    "lastName": {
12      "type": "string",
13      "description": "The person's last name."
14    },
15    "age": {
16      "description": "Age in years",
17      "type": "integer",
18      "minimum": 0
19    }
20  }
21 }
```

Schemas to the rescue



Summary



Schemas are the contracts between services

Schemas improve data quality

Schemas allows versioning

It's a way of documenting your interfaces

Enables contract testing

The problem
space

Intro to
Schemas

**Era of Web
Services**

Stepping into
Event-driven
Architecture

Rise of Event
Streaming

Schema
management -
way forward

Era of Web Services



Era of Web Services

The screenshot displays the SoapUI Pro 4.6.0 interface. The main window shows a test case named "GetCityWeatherByZIP - Request 1" with its SOAP request and response. The response is a SOAP envelope containing a "GetCityWeatherByZIPResponse" with a "City Found" status and weather details for Boston.

Annotations on the screenshot:


- Icon to add assertion:** A red box highlights the "Add Assertion" icon (a green circle with a plus sign) in the bottom left corner of the main window.
- Click assertions tab:** A red box highlights the "Assertions (1)" tab in the bottom left corner of the main window.

The bottom left pane shows the "TestRequest Properties" tab with the following properties:

Property	Value
Name	GetCityWeatherByZIP - Req...
Description	
Message Size	335
Encoding	UTF-8
Endpoint	http://wsf.cdyne.com/Weath...
Timeout	
Bind Address	
Follow Redirects	true
Interface	WeatherSoap
Operation	GetCityWeatherByZIP

The bottom right pane shows the "SOAP Response - VALID" status and a "Request Log (4)" tab.

Era of Web Services

 SMARTBEAR
SwaggerHub™

open-weather_ma... 2.5.2

Export

Info

Tags

Servers

Search

Current Weather Data ^

GET /weather

Schemas ^

SCHEMA 200

SCHEMA Coord

SCHEMA Weather

SCHEMA Main

SCHEMA Wind

SCHEMA Clouds

SCHEMA Rain

SCHEMA Snow

SCHEMA Sys

Read Only

```
125      default: "json"
126
127      schemas:
128      200:
129        title: Successful response
130        type: object
131        properties:
132          coord:
133            $ref: '#/components/schemas/Coord'
134          weather:
135            type: array
136            items:
137              $ref: '#/components/schemas/Weather'
138          description: (more info Weather condition codes)
139          base:
140            type: string
141            description: Internal parameter
142            example: cmc stations
143          main:
144            $ref: '#/components/schemas/Main'
145          visibility:
146            type: integer
147            description: Visibility, meter
148            example: 16093
149          wind:
150            $ref: '#/components/schemas/Wind'
151          clouds:
152            $ref: '#/components/schemas/Clouds'
```

OpenWeatherMap API

2.5.2 OAS3

Get current weather, daily forecast for 16 days, and 3-hourly forecast 5 days for your city. Helpful stats, graphics, and this day in history charts are available for your reference. Interactive maps show precipitation, clouds, pressure, wind around your location stations. Data is available in JSON, XML, or HTML format. **Note:** This sample Swagger file covers the **current** endpoint only from the OpenWeatherMap API.

Note: All parameters are optional, but you must select at least one parameter. Calling the API by city ID (using the **id** parameter) will provide the most precise location results.

[Terms of service](#)

[OpenWeatherMap API - Website](#)

[Send email to OpenWeatherMap API](#)

[CC Attribution-ShareAlike 4.0 \(CC BY-SA 4.0\)](#)

[API Documentation](#)

Authorize

Era of Web Services

CONSUMER
DATA
STANDARDS

Q Search

Introduction

Standards

Security Profile

Consumer Experience

Banking APIs

Get Accounts

Get Bulk Balances

Get Balances For Specific A...

Get Account Balance

Get Account Detail

Get Transactions For Account

Get Accounts

GET /banking/accounts

Obtain a list of accounts

Endpoint Version

Version 1

Parameters

Name	In	Type	Required	Description
product-category	query	string	optional	Used to filter results on the product category

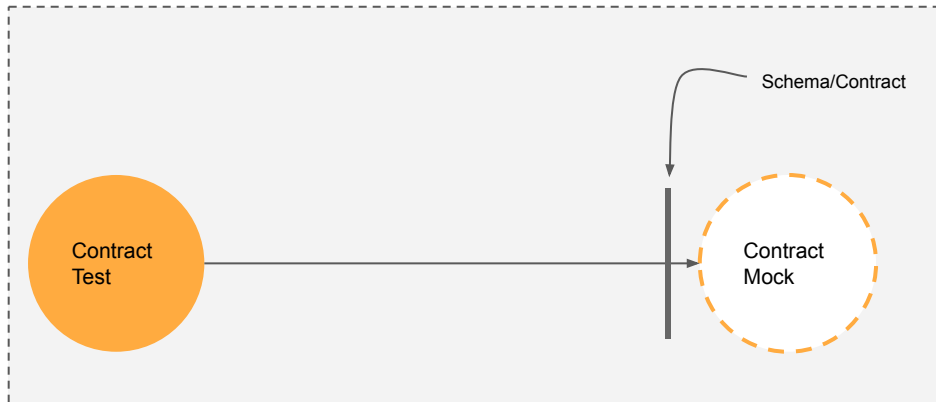
Code samples

```
GET https://data.holder.com.au/cds-au/v1/
Host: data.holder.com.au
Accept: application/json
x-v: string
x-min-v: string
x-fapi-interaction-id: string
x-fapi-auth-date: string
x-fapi-customer-ip-address: string
x-cds-client-headers: string

var headers = {
  'Accept': 'application/json',
```

Contract Testing

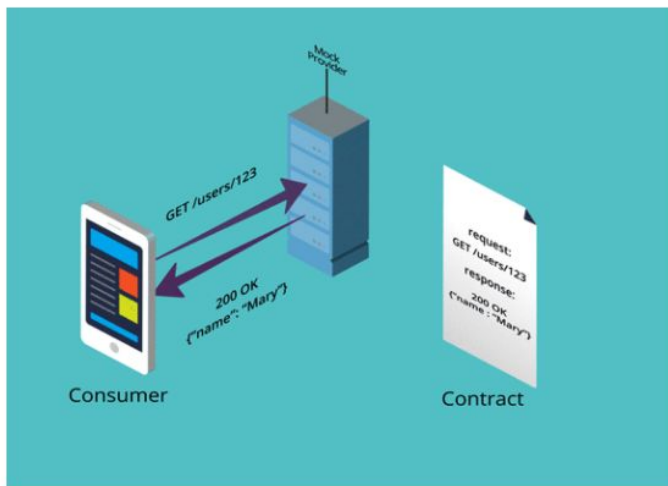
- Integration testing, but just with contracts
- Validate the end to end flow without building one.
- Find bugs upfront, don't wait till the actual integration test cycle.



Contract Testing - Example

How Pact works

How Pact helps - the consumer side



Testing a consumer using a Pact mock provider

Pact solves the problem of keeping the two sets of tests in sync by use of a "contract", also known as a "pact".

During the consumer tests, each request made to a Pact mock provider is recorded into the contract file, along with its expected response.



Replay animation

The problem
space

Intro to
Schemas

Era of Web
Services

**Stepping into
Event-driven
Architecture**

Rise of Event
Streaming

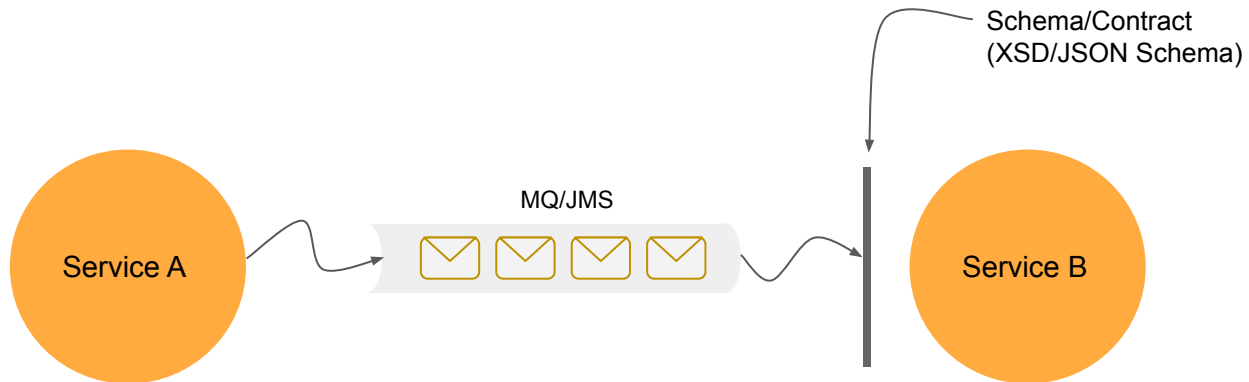
Schema
management -
way forward

Event Driven Architecture



Adoption of Message Queues = Event driven communication + Being real-time + Being asynchronous

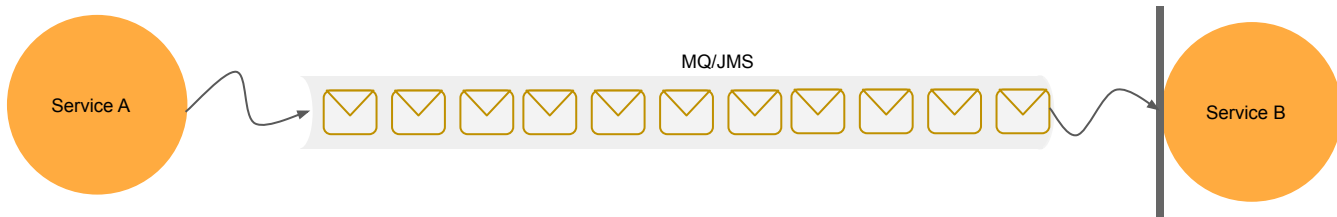
Event Driven Architecture



EDA - Schema Validation

```
ship-order.xml X
Users > vidya > Desktop > ship-order.xml
1 <?xml version="1.0" encoding="UTF-8"?>
2
3 <shiporder orderId="889923"
4 xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
5 xsi:noNamespaceSchemaLocation="shiporder.xsd">
6   <orderperson>John Smith</orderperson>
7   <shipto>
8     <name>Ola Nordmann</name>
9     <address>Langgt 23</address>
10    <city>4000 Stavanger</city>
11    <country>Norway</country>
12  </shipto>
13  <item>
14    <title>Empire Burlesque</title>
15    <note>Special Edition</note>
16    <quantity>1</quantity>
17    <price>10.90</price>
18  </item>
19  <item>
20    <title>Hide your heart</title>
21    <quantity>1</quantity>
22    <price>9.90</price>
23  </item>
24 </shiporder>
```

```
ship-order.xsd X
Users > vidya > Desktop > ship-order.xsd
1 <?xml version="1.0" encoding="UTF-8"?>
2 <xs:schema
3   xmlns:xs="http://www.w3.org/2001/XMLSchema"
4   <xs:simpleType name="stringtype">
5     <xs:restriction base="xs:string"/>
6   </xs:simpleType>
7   <xs:simpleType name="inttype">
8     <xs:restriction base="xs:positiveInteger"/>
9   </xs:simpleType>
10  <xs:simpleType name="decetype">
11    <xs:restriction base="xs:decimal"/>
12  </xs:simpleType>
13  <xs:simpleType name="orderidtype">
14    <xs:restriction base="xs:string">
15      <xs:pattern value="[0-9]{6}"/>
16    </xs:restriction>
17  </xs:simpleType>
18  <xs:complexType name="shiptotype">
19    <xs:sequence>
20      <xs:element name="name" type="stringtype"/>
21      <xs:element name="address" type="stringtype"/>
22      <xs:element name="city" type="stringtype"/>
23      <xs:element name="country" type="stringtype"/>
24    </xs:sequence>
25  </xs:complexType>
26  <xs:complexType name="itemtype">
27    <xs:sequence>
28      <xs:element name="title" type="stringtype"/>
29      <xs:element name="note" type="stringtype" minOccurs="0"/>
30      <xs:element name="quantity" type="inttype"/>
31      <xs:element name="price" type="decetype"/>
32    </xs:sequence>
33  </xs:complexType>
34  <xs:complexType name="shipordertype">
35    <xs:sequence>
36      <xs:element name="orderperson" type="stringtype"/>
37      <xs:element name="shipto" type="shiptotype"/>
38      <xs:element name="item" maxOccurs="unbounded" type="itemtype"/>
39    </xs:sequence>
40    <xs:attribute name="orderId" type="orderidtype" use="required"/>
41  </xs:complexType>
42  <xs:element name="shiporder" type="shipordertype"/>
43 </xs:schema>
```



The problem
space

Intro to
Schemas

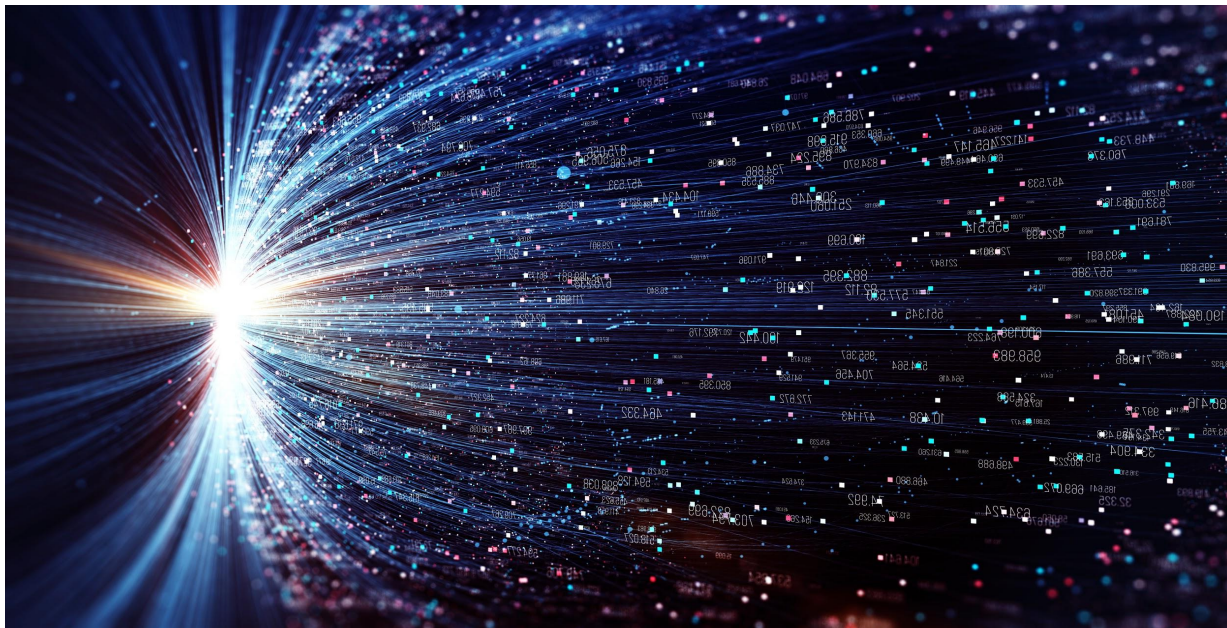
Era of Web
Services

Stepping into
Event-driven
Architecture

**Rise of Event
Streaming**

Schema
management -
way forward

Event Streaming



Rise of Event Streams = Capture **All** possible **Events** + at a **Scale** + **Process** them on the fly

Event Streaming - A brief

- Record all events at a scale, in and around an Enterprise
- Process/Interrogate/Analyse events on the fly (Event Stream Processing)
- Sophisticated real-time processing abilities: Windowing, re-ordering, etc.
- Good examples: Apache Kafka, AWS Kinesis, Google Pub-Sub

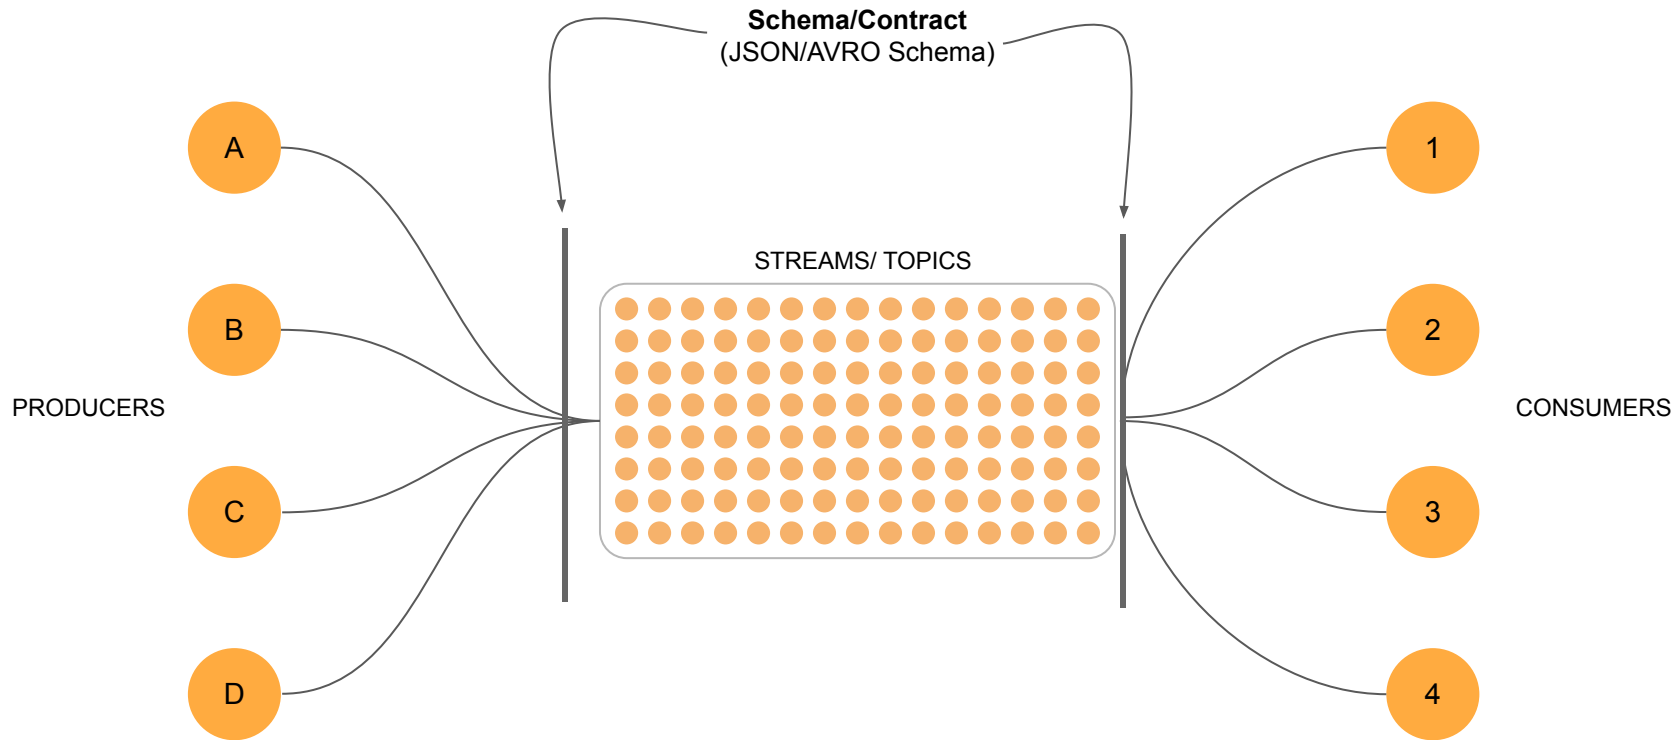


Google Cloud
Pub/Sub

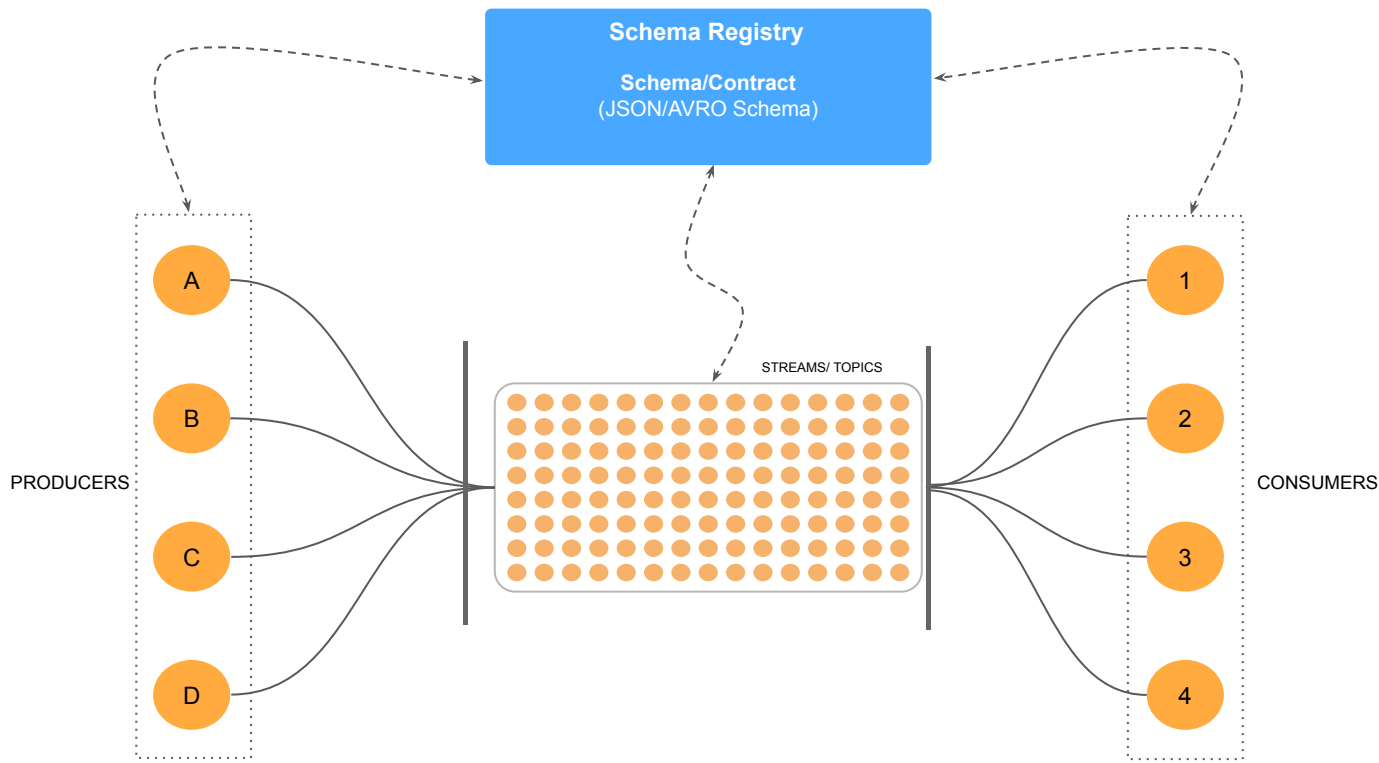


Rise of Schema Registry

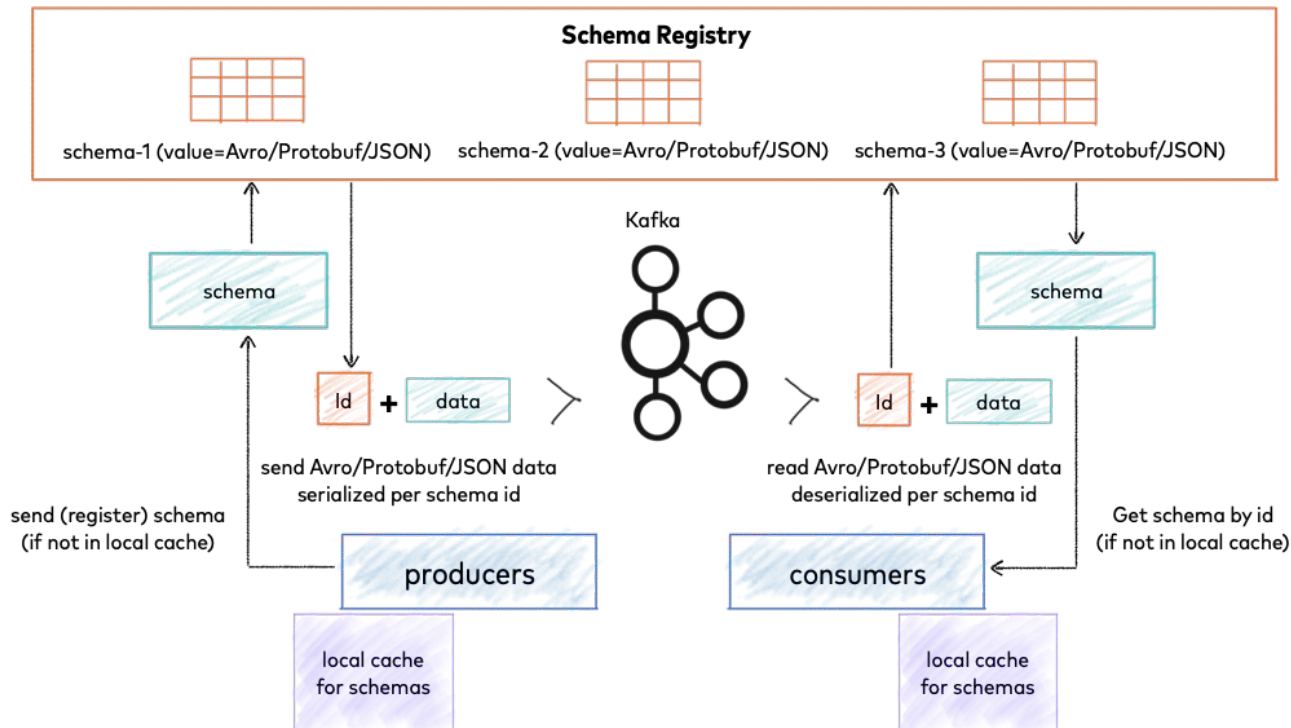
Schema Registry



Schema Registry - How it works



Schema Registry - How it works (Example)



Reference - <https://docs.confluent.io/current/schema-registry/index.html>

Schema Registry - How it works (Producer Example)

master schema-registry-example / src / main / resources / avro / ParkingAddress.avsc

vidyavenu producer and consumer changes

1 contributor

23 lines (10 sloc) 331 Bytes

```
1 {
2   "type": "record",
3   "namespace": "com.schemaregistryexamples",
4   "name": "ParkingAddress",
5   "version": "1",
6   "fields": [
7     { "name": "human_address", "type": "string", "doc": "Complete Address" },
8     { "name": "status", "type": "string", "doc": "Parking present or unoccupied" }
9   ]
10 }
```

1. Define Schema

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - How it works (Producer Example)

2. Register Schema to a topic

ALL TOPICS >

parking-events

Overview Messages **Schema** Configuration

Value Key

[Edit schema](#) [Version history](#) [Download](#)

```
1 {
2   "fields": [
3     {
4       "doc": "Complete Address",
5       "name": "human_address",
6       "type": "string"
7     },
8     {
9       "doc": "Parking present or unoccupied",
10      "name": "availability",
11      "type": "string"
12    }
13  ],
14  "name": "ParkingAddress",
15  "namespace": "com.schemaregistryexamples",
16  "type": "record",
17  "version": "1"
18 }
```

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - How it works (Producer Example)

```
public class ProducerExample {  
    private static final String TOPIC = "parking-events";  
  
    @SuppressWarnings("InfiniteLoopStatement")  
    public static void main(final String[] args) {  
  
        final Properties props = new Properties();  
        props.put(ProducerConfig.BOOTSTRAP_SERVERS_CONFIG, "localhost:9092");  
        props.put(ProducerConfig.ACKS_CONFIG, "all");  
        props.put(ProducerConfig.RETRIES_CONFIG, 0);  
        props.put(ProducerConfig.KEY_SERIALIZER_CLASS_CONFIG, StringSerializer.class);  
        props.put(ProducerConfig.VALUE_SERIALIZER_CLASS_CONFIG, KafkaAvroSerializer.class);  
        props.put(AbstractKafkaAvroSerDeConfig.SCHEMA_REGISTRY_URL_CONFIG, "http://localhost:8081");  
  
        try (KafkaProducer<String, ParkingAddress> producer = new KafkaProducer<>(props)) {  
  
            for (long i = 0; i < 20; i++) {  
                final String parkingID = "parkingid-" + Long.toString(i);  
  
                final ParkingAddress parkingAddress = ParkingAddress.newBuilder().setHumanAddress("Albert st, Melbourne").setStatus("Available").build();  
                final ProducerRecord<String, ParkingAddress> record = new ProducerRecord<>(TOPIC, parkingID, parkingAddress);  
                producer.send(record);  
                Thread.sleep(1000L);  
            }  
        }  
    }  
}
```

3. Produce data

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - How it works (Producer Example)

```
17:17:52.705 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.NetworkClient - [Producer clientId=producer-1] Created socket
17:17:52.705 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.NetworkClient - [Producer clientId=producer-1] Completed connect
17:17:52.705 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.NetworkClient - [Producer clientId=producer-1] Initiating API version
17:17:52.713 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.NetworkClient - [Producer clientId=producer-1] Recorded API version
Successfully sent 20 messages to parking-events
17:18:12.773 [main] INFO org.apache.kafka.clients.producer.KafkaProducer - [Producer clientId=producer-1] Closing the Kafka producer with timeoutMillis
17:18:12.773 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.producer.internals.Sender - [Producer clientId=producer-1] Beginning
17:18:12.784 [kafka-producer-network-thread | producer-1] DEBUG org.apache.kafka.clients.producer.internals.Sender - [Producer clientId=producer-1] Shut
17:18:12.785 [main] DEBUG org.apache.kafka.clients.producer.KafkaProducer - [Producer clientId=producer-1] Kafka producer has been closed
```

4. Messages sent !

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - How it works (Producer Example)



The screenshot shows a GitHub commit titled "Update schema - Introduced parking_id" on the master branch. The commit was made by vidyavenu 22 seconds ago. It shows a diff for the file src/main/resources/avro/ParkingAddress.avsc. The diff indicates that a new field "parking_id" was added to the schema, while the "status" field remains unchanged. The "human_address" field is also present.

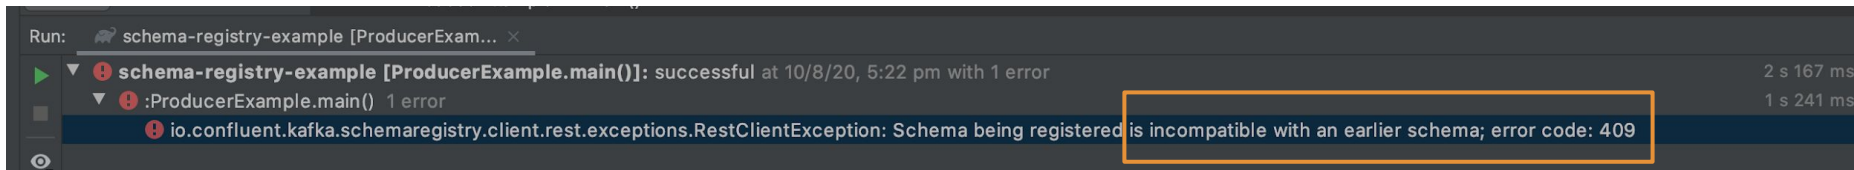
```
@@ -5,7 +5,8 @@
5      "version": "1",
6      "fields": [
7        { "name": "human_address", "type": "string", "doc": "Complete Address" },
8        { "name": "status", "type": "string", "doc": "Parking present or unoccupied" },
9        { "name": "status", "type": "string", "doc": "Parking present or unoccupied" },
10       { "name": "parking_id", "type": "string", "doc": "Parking ID" }
11     ]
12   }
```

6. Producer modifies the payload

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - How it works (Producer Example)

7. Validation fails :)



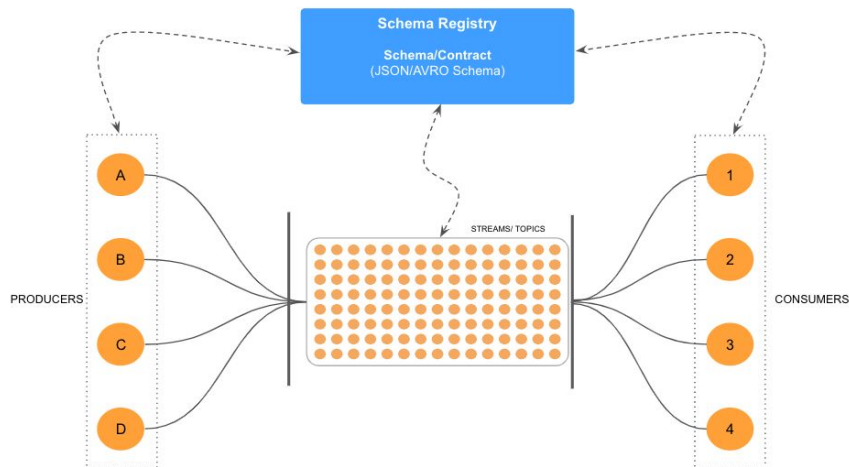
The screenshot shows the Run console of an IDE. The top line indicates a successful run of `schema-registry-example [ProducerExample.main()]` at 10/8/20, 5:22 pm, with 1 error, taking 2 s 167 ms. The second line shows the error details for `:ProducerExample.main()`, which took 1 s 241 ms. The error message is `io.confluent.kafka.schemaregistry.client.rest.exceptions.RestClientException: Schema being registered is incompatible with an earlier schema; error code: 409`. The error message is highlighted with an orange box.

```
Run: schema-registry-example [ProducerExam... x
▼ schema-registry-example [ProducerExample.main()]: successful at 10/8/20, 5:22 pm with 1 error 2 s 167 ms
  ▼ :ProducerExample.main() 1 error 1 s 241 ms
    io.confluent.kafka.schemaregistry.client.rest.exceptions.RestClientException: Schema being registered is incompatible with an earlier schema; error code: 409
```

Reference - <https://github.com/vidyavenu/schema-registry-example/>

Schema Registry - Summary

- Centralised schema registry between Producers, Consumers & Event Streams
- Simplified serialization and deserialization
- Allows versioning & compatibility checks (forward/backward compatibility)



The problem
space

Intro to
Schemas

Era of Web
Services

Stepping into
Event-driven
Architecture

Rise of Event
Streaming

Schema
management -
way forward

