

What's new IBM Event Streams 2019.4.1

Improving Apache Kafka for enterprise

IBM Software Engineer Chris Patmore



Chris Patmore

Chris is an IBM Software Engineer and is a member of the IBM Event Streams development team. Through his work on IBM Event Streams he has built up experience running enterprise level Kafka clusters on Kubernetes, and writing microservices for the cloud. He has a first class masters in physics and If he's not writing code, he's probably out bouldering.



Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



Questions



Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



Questions



What is a schema

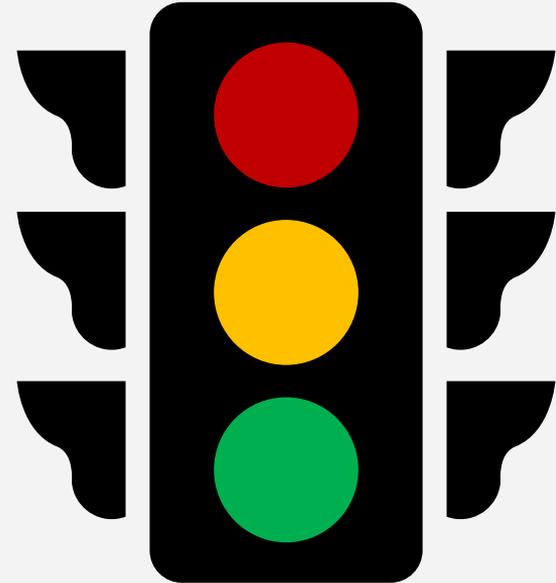
```
{
  "type": "record",
  "name": "Book",
  "namespace": "com.mycompany.schemas.materials",
  "fields": [
    {
      "name": "Title",
      "type": "string"
    },
    {
      "name": "Author",
      "type": "string"
    },
    {
      "name": "Blurb",
      "type": "string"
    },
    {
      "name": "Rating",
      "type": "int"
    },
    {
      "name": "Pages",
      "type": "int"
    }
  ]
}
```

Apache Avro

A schema defines the content and types of data in a Kafka message

Schemas Provide Governance

Schemas ensure message integrity between producers and consumers.



Schema Validation

Upload from UI

Upload from CLI

Get schema validation

Definition Preview

Schema Name	Book	Edit
Version	1.0.0	Edit
Type	Apache Avro Record	
Namespace	com.mycompany.schemas.materials	
Valid schema	Schema is valid	

Definition Preview

```
{
  "type": "record",
  "name": "Book",
  "namespace": "com.mycompany.schemas.materials",
  "fields": [
    {
      "name": "Title",
      "type": "string"
    },
    {
      "name": "Author",
      "type": "string"
    },
    {
      "name": "Blurb",
      "type": "string"
    },
    {
      "name": "Rating",
      "type": "int"
    },
    {
      "name": "Pages",
      "type": "string"
    }
  ]
}
```

Schema Versioning

Controlled compatible evolution

Definition Preview

 Schema file validation failed: Avro schema is not compatible with latest schema version: Compatibility type 'MUTUAL_READ' does not hold between 1 schema(s) in the chronology because: Schema[0] has incompatibilities: ['READER_FIELD_MISSING_DEFAULT_VALUE: Cost' at '/fields/5'].

Schema Name	Book
Version	1.1.0 Edit
Type	Apache Avro Record
Namespace	com.mycompany.schemas.materials
Valid schema	

Book

Versions

Manage schema

Add new version

Filter versions

Overview (all versions)

Version 1.1.0

Version 1.0.0

Version 1.1.0

Uploaded

23/10/2019, 11:46:02

Status

Active

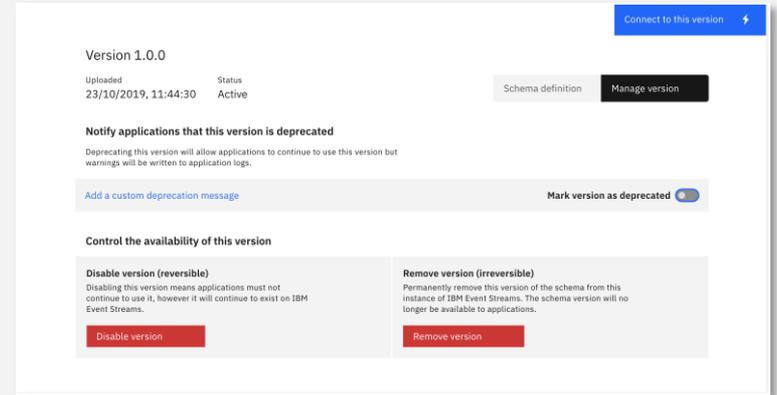
```
{
  "type": "record",
  "name": "Book",
  "namespace": "com.mycompany.schemas.materials",
  "fields": [
    {
      "name": "Title",
      "type": "string"
    },
    {
      "name": "Author",
      "type": "string"
    },
    {
      "name": "Blurb",
      "type": "string"
    },
    {
      "name": "Rating",
      "type": "int"
    },
    {
      "name": "Pages",
      "type": "string"
    },
    {
      "name": "Cost",
      "type": "int",
      "default": 0
    }
  ]
}
```

Schema Lifecycle

Manage the lifecycle of your schemas

Inform clients about changes

Stop outdated clients



Active

All clients running happily

Deprecated

Producers get a warning

Customise your warning

Consumers don't notice

Disabled

Producers stop

Consumers get warned

Removed

All clients stop

Application Development

Schema applications are Effortless to develop

1

Connect a Java application

[Not using Java?](#)

Configure the schema connection

Connect to schema	Version
Book	1.0.0

Message encoding

Binary
 JSON

Message deserialization behavior

If an application encounters messages that do not conform to the schema, what message deserialization behavior would you like?

Strict
 Permissive

Use generated or generic code

Your schema can use custom Java classes that are generated based on that schema, or can use generic code using the Apache Avro API.

Use schema-specific code
 Use generic Apache Avro schema code

Close Save

2

Configuration / Downloads and sample code

Follow these steps to set up your Java application

Connect to schema	Version
Book	1.0.0

1. Get the cluster certificate

 **Java truststore**
Download the Java truststore file which contains the server certificate. The truststore password is 'password'.

2. Download the schema registry dependencies or configure Maven to install dependencies in your project

Download the Event Streams schema registry JAR files and manually add them to your project. Alternatively, you can copy the configuration snippets for the Event Streams maven repository to your project Maven POM file and run the Maven install command to download and install project dependencies.

Use JARs Use Maven

 **Java dependencies**
Download all the schema registry JAR files in a ZIP

 **Schema JAR**
Download the Book schema JAR file

3. Copy this sample code

Copy these sample Java code snippets to use the schema registry in your application. Replace the <java_truststore_file_location>, <api_key>, <my_topic> and <my_consumer_group> items with real values.

Producer **Consumer**

Imports Connection properties Producer code

```
import java.util.Properties;

// Import the specific schema class
import com.mycompany.schemas.materials.Book;

import org.apache.kafka.clients.CommonClientConfigs;
import org.apache.kafka.common.config.SaslConfigs;
import org.apache.kafka.common.config.SslConfigs;
import com.ibm.eventstreams.serdes.SchemaRegistryConfig;
import com.ibm.eventstreams.serdes.SchemaInfo;
import com.ibm.eventstreams.serdes.SchemaRegistry;

import org.apache.kafka.clients.producer.KafkaProducer;
import org.apache.kafka.clients.producer.ProducerRecord;
```

Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



Questions

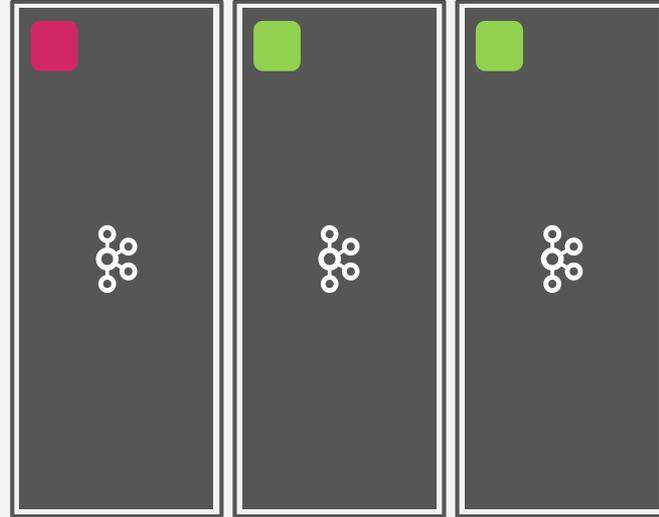


Multizone Provides Resilience

Spread your enterprise resources

Native Kubernetes feature

Tolerate failure

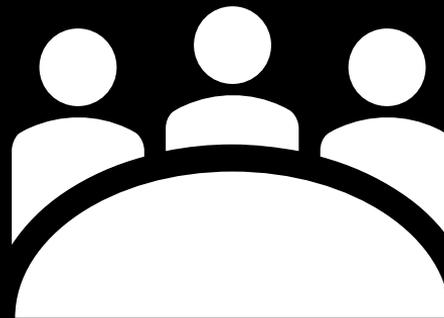


Consensus

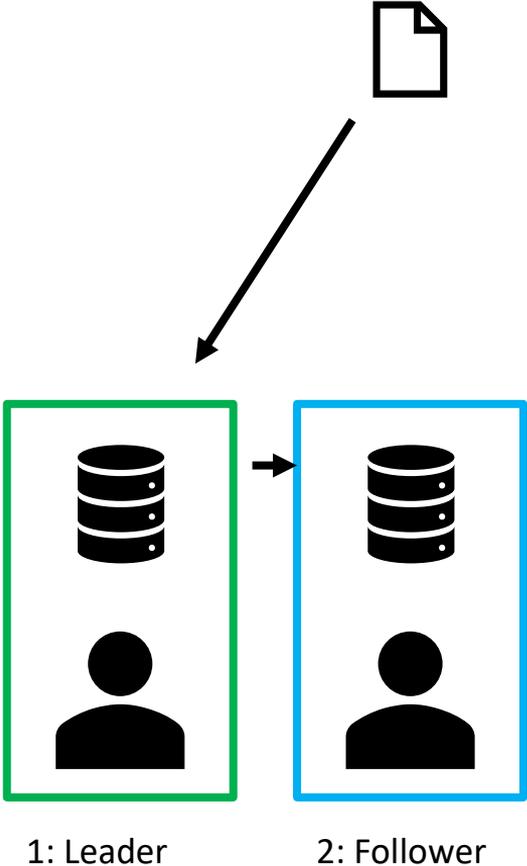
Kubernetes is asynchronous

Who's right

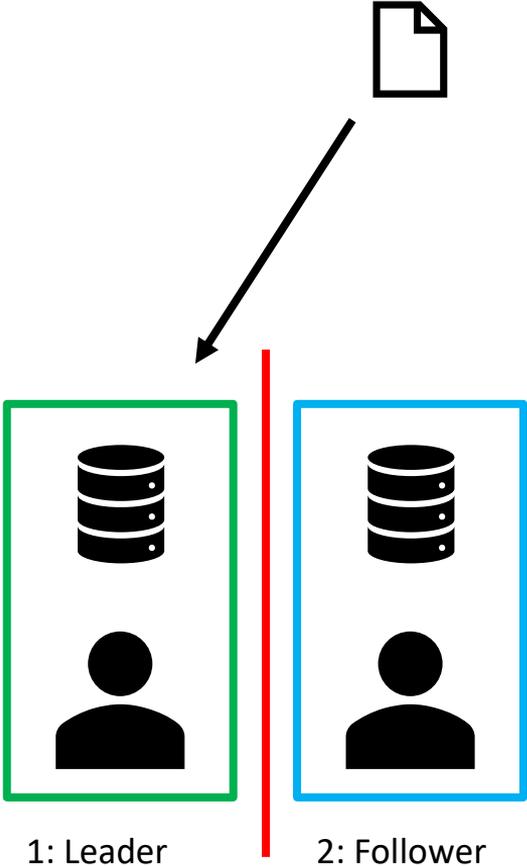
Consensus required



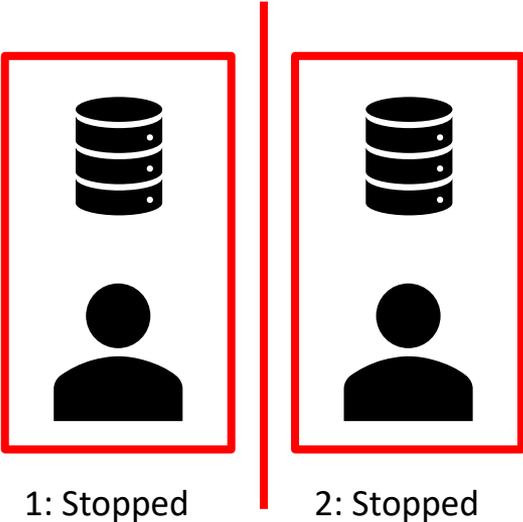
Consensus continued



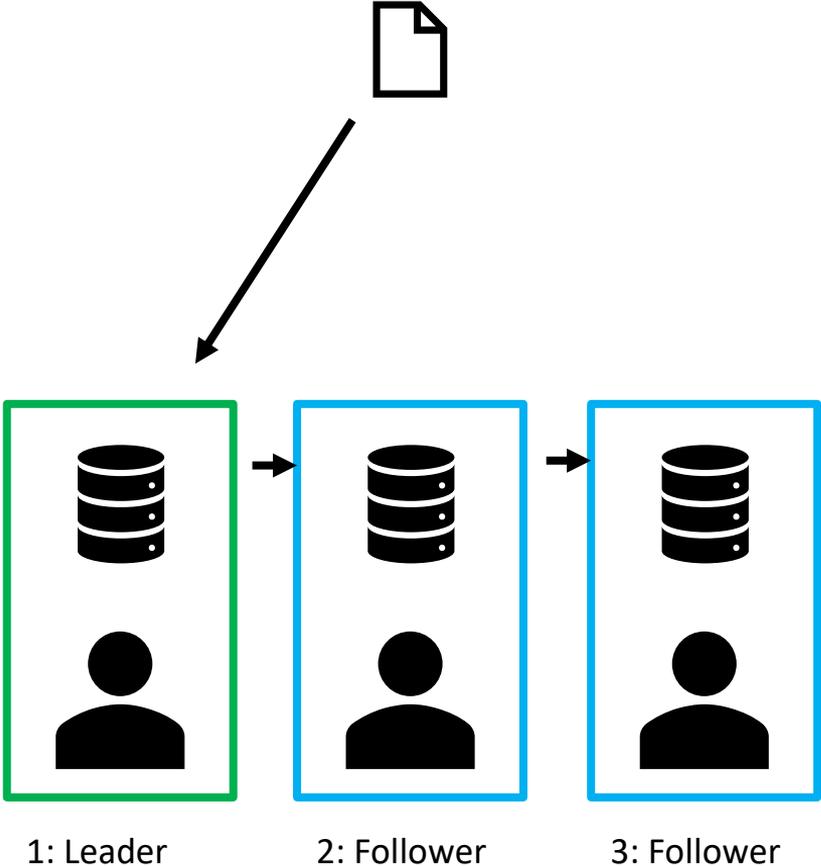
Consensus continued



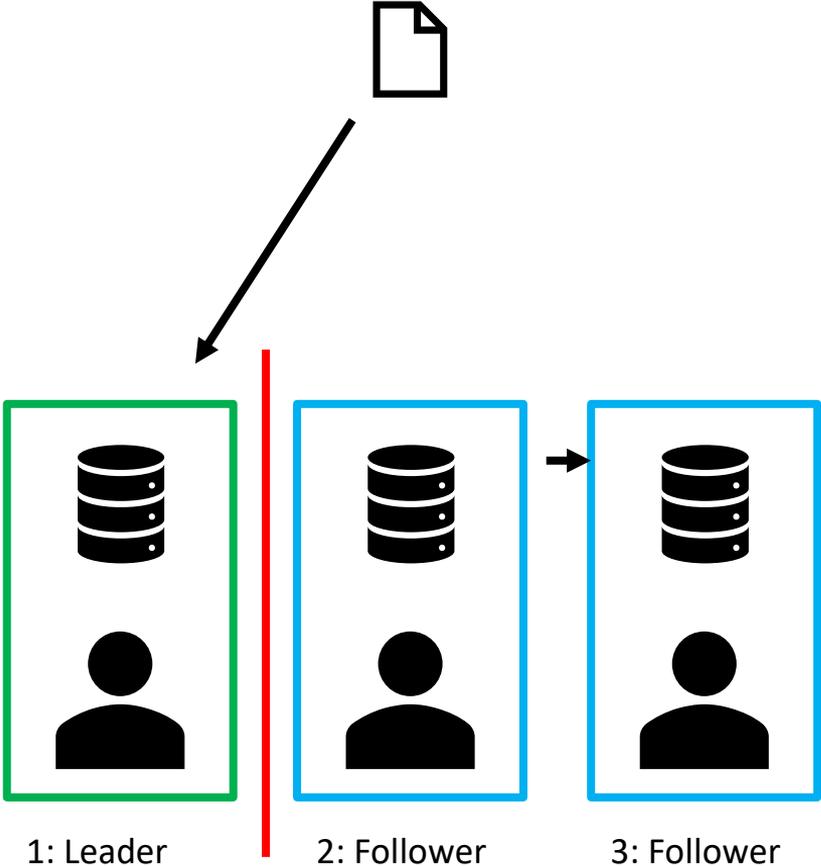
Consensus continued



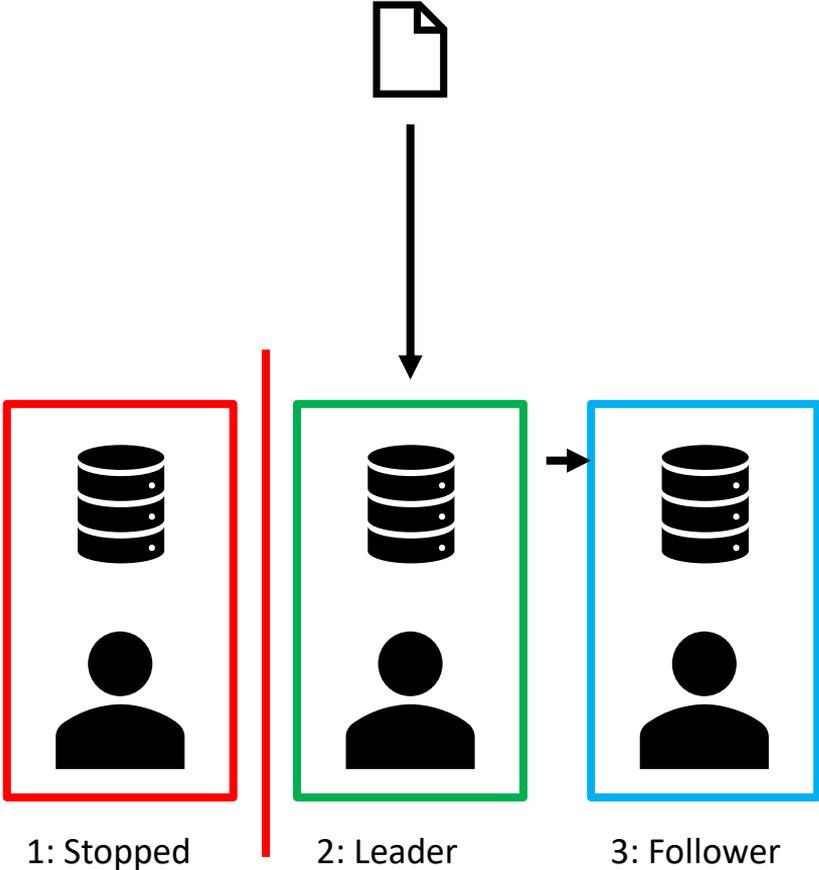
Consensus continued



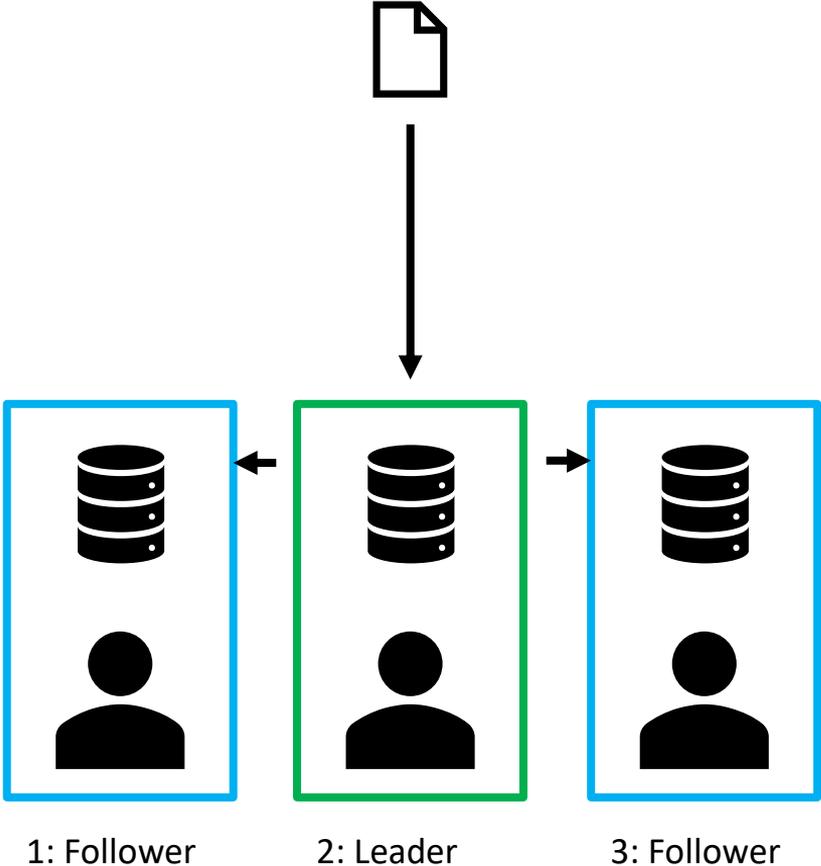
Consensus continued



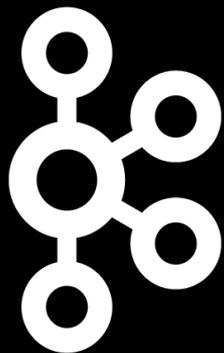
Consensus continued



Consensus continued



Kafka and Zones



Align Kafka with Kubernetes

Spread topics and partitions

Keep your data

Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



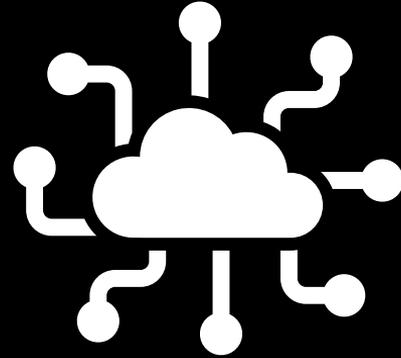
Questions



CICS Integration

CICS events = Kafka messages

Secure by design



Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



Questions



Thank you for listening

Documentation:

<https://ibm.github.io/event-streams/schemas/overview/>

<https://ibm.github.io/event-streams/installing/preparing-multizone/>

https://www.ibm.com/support/knowledgecenter/en/SSGMCP_5.2.0/com.ibm.cics.ts.event_processing.doc/concepts/dfhep_event_processing_adapters.html

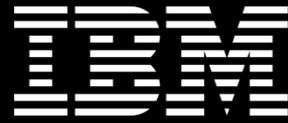
<https://ibm.github.io/event-streams/connecting/rest-api/>

Ask more questions on slack:

<https://ibm.github.io/event-streams/support/>

IBM Cloud Technology - #ibm-eventstreams





Apache Avro schemas

What is a schema

Why use schemas

How Event Streams
improves the experience



Multizone

Why multizone

Kafka and zones



CICS integration



Questions

