

Please note

IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.

Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.

The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Notices and disclaimers

© 2018 International Business Machines Corporation. No part of this document may be reproduced or transmitted in any form without written permission from IBM.

U.S. Government Users Restricted Rights — use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM.

Information in these presentations (including information relating to products that have not yet been announced by IBM) has been reviewed for accuracy as of the date of initial publication and could include unintentional technical or typographical errors. IBM shall have no responsibility to update this information. **This document is distributed “as is” without any warranty, either express or implied. In no event, shall IBM be liable for any damage arising from the use of this information, including but not limited to, loss of data, business interruption, loss of profit or loss of opportunity.** IBM products and services are warranted per the terms and conditions of the agreements under which they are provided.

IBM products are manufactured from new parts or new and used parts. In some cases, a product may not be new and may have been previously installed. Regardless, our warranty terms apply.”

Any statements regarding IBM's future direction, intent or product plans are subject to change or withdrawal without notice.

Performance data contained herein was generally obtained in a controlled, isolated environments. Customer examples are presented as illustrations of how those customers have used IBM products and the results they may have achieved. Actual performance, cost, savings or other results in other operating environments may vary.

References in this document to IBM products, programs, or services does not imply that IBM intends to make such products, programs or services available in all countries in which IBM operates or does business.

Workshops, sessions and associated materials may have been prepared by independent session speakers, and do not necessarily reflect the views of IBM. All materials and discussions are provided for informational purposes only, and are neither intended to, nor shall constitute legal or other guidance or advice to any individual participant or their specific situation.

It is the customer's responsibility to insure its own compliance with legal requirements and to obtain advice of competent legal counsel as to the identification and interpretation of any relevant laws and regulatory requirements that may affect the customer's business and any actions the customer may need to take to comply with such laws. IBM does not provide legal advice or represent or warrant that its services or products will ensure that the customer follows any law.

Notices and disclaimers continued

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products about this publication and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products. IBM does not warrant the quality of any third-party products, or the ability of any such third-party products to interoperate with IBM's products. **IBM expressly disclaims all warranties, expressed or implied, including but not limited to, the implied warranties of merchantability and fitness for a purpose.**

The provision of the information contained herein is not intended to, and does not, grant any right or license under any IBM patents, copyrights, trademarks or other intellectual property right.

IBM, the IBM logo, ibm.com and [names of other referenced IBM products and services used in the presentation] are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at “Copyright and trademark information” at: www.ibm.com/legal/copytrade.shtml.

Introducing IBM Event Streams

Making Apache Kafka a piece of cake

think 2019

—
Emma Humber
IBM Event Streams



Event Streaming and Messaging



Apache Kafka



IBM Event Streams

What is Event Streaming

Event-Driven in Action

Getting data to where it's needed, before it's needed



**Respond to events before
the moment passes**




**Responsive & personalised
customer experiences**




**Bring real time intelligence
to your apps**

Properties of the Event Backbone




Stream history




Scalable
consumption



Immutable data



Scalable



Highly available

Isn't This “Just Messaging”?

Message Queuing & Event Streaming focus on different aspects of Messaging

Operations

Messages that represent some **current or future processing**. For instance: request and response messages.

Message Queuing

Informational

Messages that represent the **state of the system**. For instance: logging, measurements and notifications.


Event Streaming

Event Streaming & Message Queuing Need Different Capabilities

Event Streaming



Stream History



Scalable
consumption



Immutable data

Message Queuing



Transient Data
Persistence



Request / Reply



Assured Delivery

What is Apache Kafka?

Apache Kafka is an **open source, distributed streaming platform**

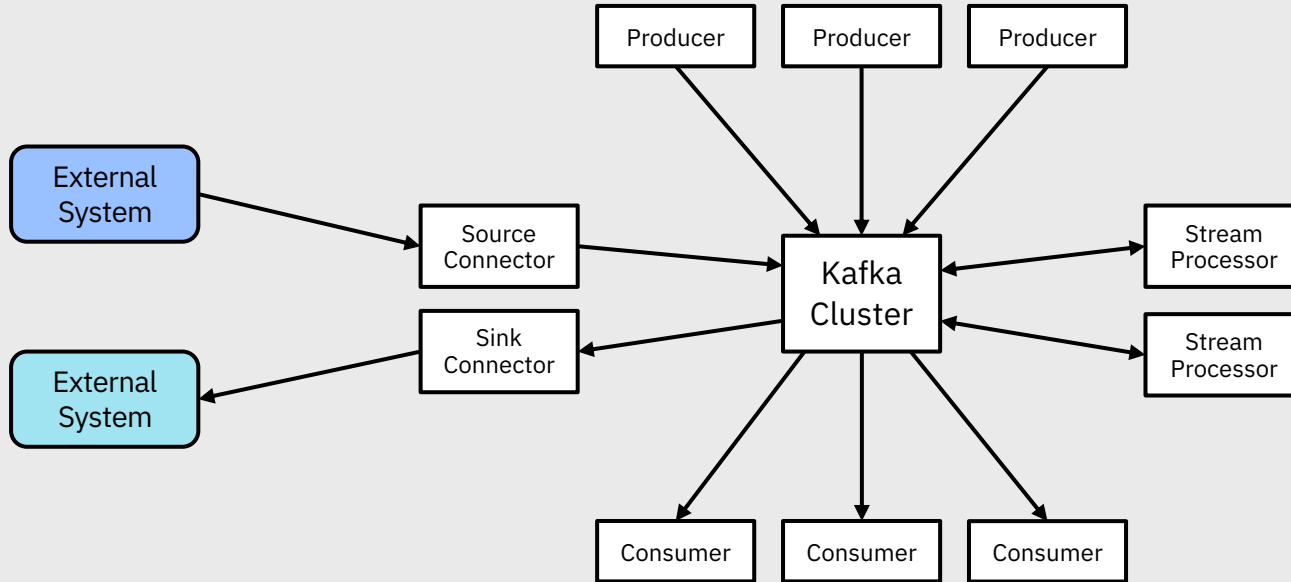


Publish and subscribe to streams of events

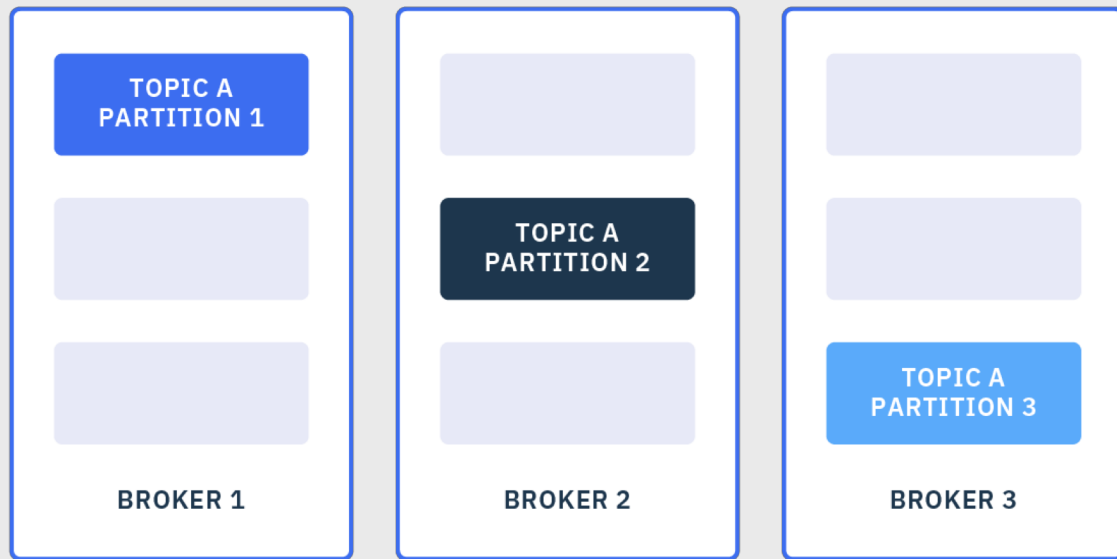
Store events in durable way

Process streams of events as they occur

Apache Kafka is an Open-Source Streaming Platform

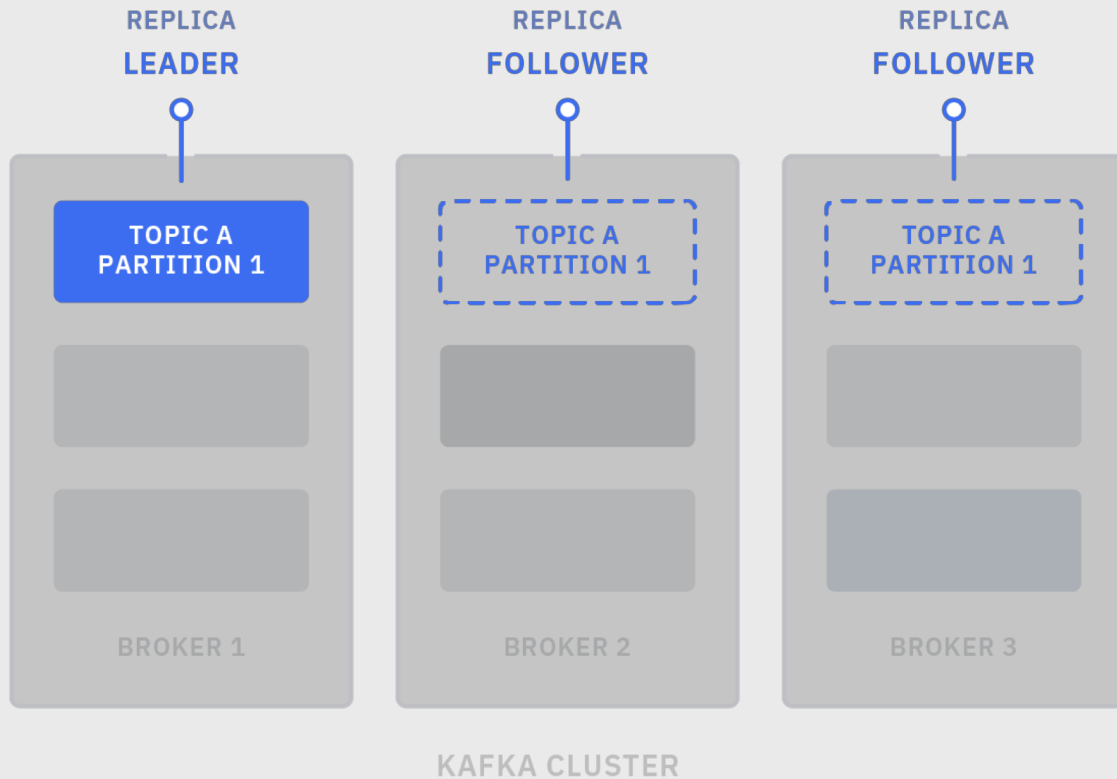


Kafka cluster




KAFKA CLUSTER

Replication



Topics




Producers

TOPIC


PARTITION 0



PARTITION 1



PARTITION 2



Producer can choose acknowledgement level:

0

Fire-and-forget
Fast, but risky

1

Waits for 1 broker to acknowledge

ALL

Waits for all replica brokers to acknowledge

Producer can choose whether to retry:

0

Do not retry
Loses messages on error

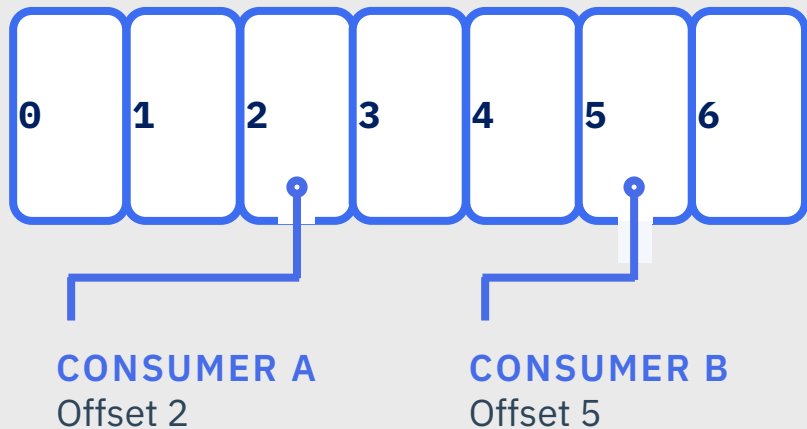
>0

Retry
Retry, might result in duplicates on error

Producer can also choose idempotence

Can retry without risking duplicates

Consumers



Consumer can choose how to commit offsets:

Automatic

Commits might go faster than processing

Manual,
asynchronous

Fairly safe, but could re-process messages

Manual,
synchronous

Safe, but slows down processing

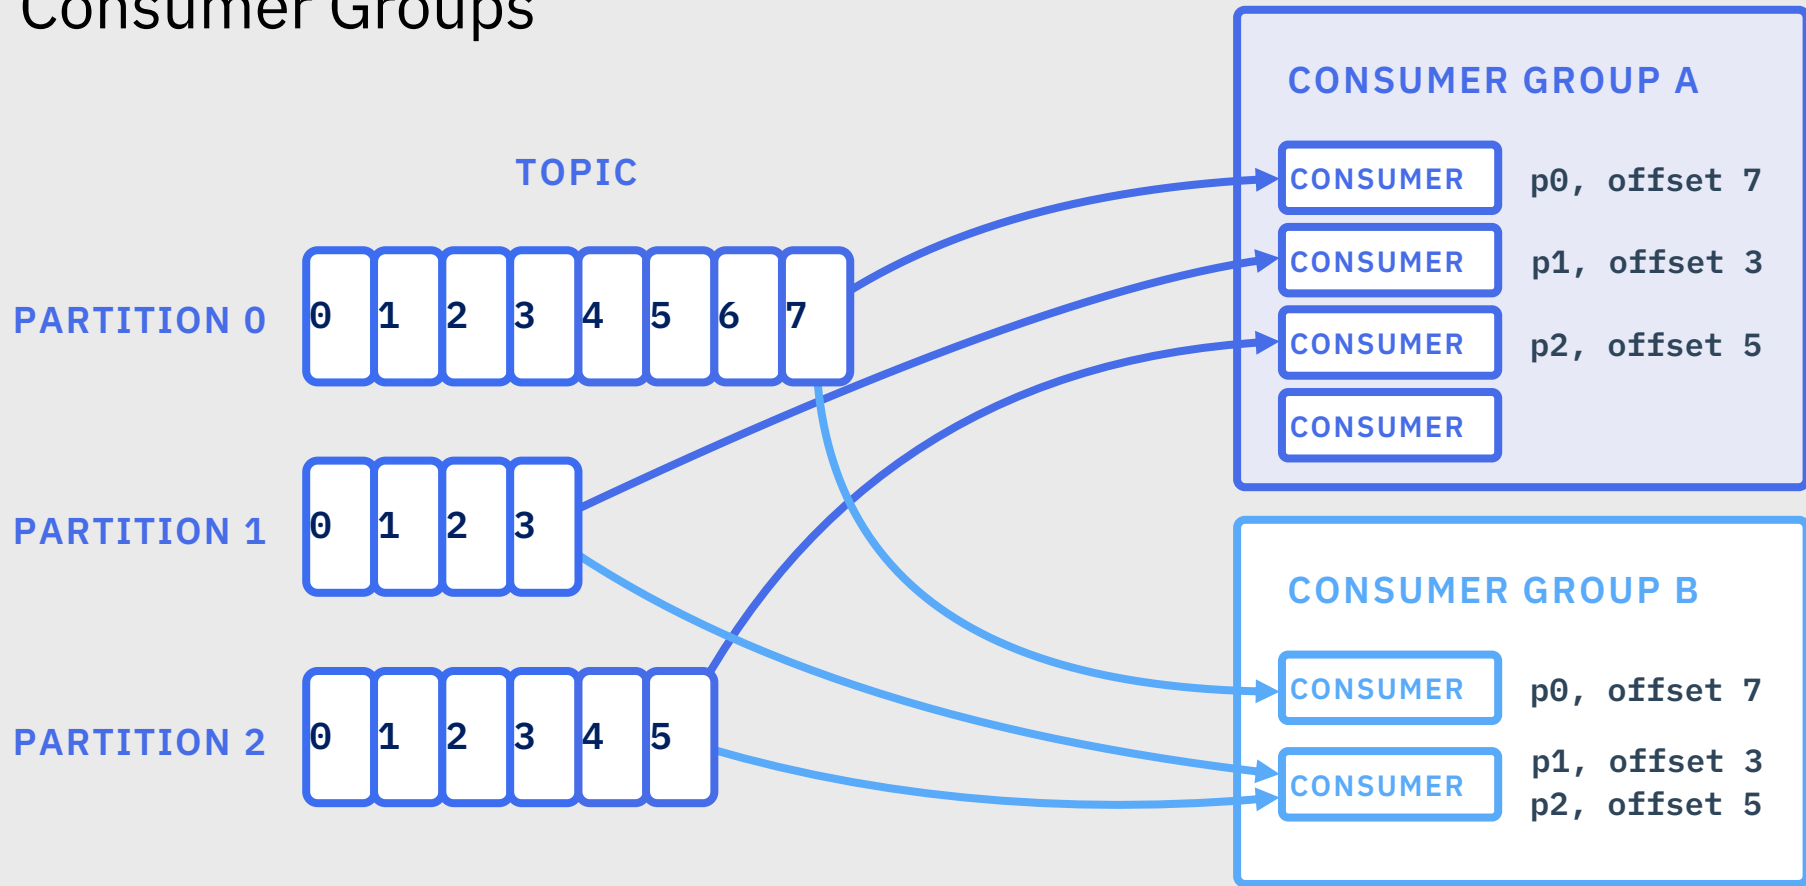
A common pattern is to commit offsets on a timer

Exactly once semantics

Can group sending messages and committing offsets into transactions

Primarily aimed at stream processing applications

Consumer Groups



What is IBM Event Streams?



IBM Event Streams

Fully supported
Apache Kafka®
with value-add
capabilities



Differentiated Value

- **IBM has years of operational expertise** running Apache Kafka for Enterprises
 - This experience has been embedded in the DNA of Event Streams
- Event Streams makes Kafka easy to run, manage & consume, **reducing skill requirements** and increasing speed of deployment for **faster time to value**
- IBM Cloud security integration **simplifies Kafka access control** using roles and policies
- **IBM's experience in Enterprise-critical software** has shaped features like geo-replication for disaster recovery & integration with IBM MQ, to give confidence deploying **mission-critical workloads**
- **Support you can Trust** – IBM has decades of experience supporting the World's toughest environments

IBM offers **Event Streams** in 3 form factors:

**Container-native
Software**

IBM Private Cloud
x86 and zLinux

**Isolated
aaS**

**Multi-tenant
aaS**

IBM Public Cloud

Making Apache Kafka intuitive and easy

IBM Cloud Private

Create resource Catalog Docs Support

← View All

Configure ibm-eventstreams-dev V 0.1.1

Kafka broker configuration

| | |
|---------------------------------|------------------------------------|
| CPU limit for Kafka brokers * | Memory limit for Kafka brokers * |
| 1000m | 1Gi |
| CPU request for Kafka brokers * | Memory request for Kafka brokers * |
| 1000m | 1Gi |

Cancel Install

- Many distinct components to deploy, configure and coordinate secure connectivity
- Container placement critical to ensure production level availability
- Secured network traffic ingress
- Ensuring consistent and repeatable deployment

think-2019

Getting started

Topics

Consumer groups

Monitor

Toolbox

[Connect to this cluster](#) 

Welcome to IBM Event Streams, let's get you up and running...



Use a simulated topic

Start exploring what IBM Event Streams has to offer with our simulated topic. You can do this even if your brokers aren't ready



Generate a starter application

Download and install our starter Kafka application and view data flowing to and from IBM Event Streams in just a few minutes

Learn more...[FAQs](#) [GitHub](#) [Documentation](#)

Kafka basics

Learn the basics of Apache Kafka, the heart of Event Streams. **System is healthy**



Topics

Airline_delays

Messages

Consumer groups

Connection information



All partitions




Showing 871 message(s) across all partitions

Find message



19 September 2018, 15:42:18



View live data



Select timeframe of data to display

Hours



Select start date of data

19/09/2018

| SEPTEMBER 2018 | | | | | | |
|----------------|----|----|----|----|----|----|
| S | M | T | W | Th | F | S |
| 26 | 27 | 28 | 29 | 30 | 31 | 1 |
| 2 | 3 | 4 | 5 | 6 | 7 | 8 |
| 9 | 10 | 11 | 12 | 13 | 14 | 15 |
| 16 | 17 | 18 | 19 | 20 | 21 | 22 |
| 23 | 24 | 25 | 26 | 27 | 28 | 29 |

Select start time of data

14:46:42

Timestamp

Partition

Offset

19/09/2018, 15:04:52

0

0

19/09/2018, 15:04:56

0

1

19/09/2018, 15:04:57

0

2

19/09/2018, 15:05:01

0

3

19/09/2018, 15:05:01

0

4

19/09/2018, 15:05:03

0

5

19/09/2018, 15:05:05

0

6

19/09/2018, 15:05:07

0

7

19/09/2018, 15:05:11

0

8

19/09/2018, 15:05:12

0

9



System is healthy


[← Topics / MY.FIRST.TOPIC](#)[Messages](#)[Consumer groups](#)Topic: MY.FIRST.TOPIC [< Previous offset](#)[Next offset >](#)

Partition

0

Offset

2

[View live data](#) 

Select timeframe of data to display

Hours 

Select start date of data

 09/02/2019

<

FEBRUARY 2019

| S | M | T | W | Th | F | S |
|----|----|----|----|----|----|----|
| 27 | 28 | 29 | 30 | 31 | 1 | 2 |
| 3 | 4 | 5 | 6 | 7 | 8 | 9 |
| 10 | 11 | 12 | 13 | 14 | 15 | 16 |
| 17 | 18 | 19 | 20 | 21 | 22 | 23 |
| 24 | 25 | 26 | 27 | 28 | 1 | 2 |
| 3 | 4 | 5 | 6 | 7 | 8 | 9 |

Select start time of data

 21:47:30Indexed timestamp 

Partition

09/02/2019, 22:46:56

0

09/02/2019, 22:46:58

0

09/02/2019, 22:47:00

0

09/02/2019, 22:47:02

0

09/02/2019, 22:47:04

0

09/02/2019, 22:47:06

0

09/02/2019, 22:47:08

0

09/02/2019, 22:47:10

0

09/02/2019, 22:47:12

0

09/02/2019, 22:47:14

0

 Message size
4 B Kafka timestamp 
09/02/2019, 22:47:00 Key
-[Raw payload](#)demo 

[← Topics](#)

MY.FIRST.TOPIC

[Messages](#)[Consumer groups](#)[All partitions](#)

Showing 34 messages

2 February 2019, 22:48:50

2:00:00 AM

2/5/2019

[View live data](#)

Select timeframe of data to display

[Hours](#)

Select start date of data

 09/02/2019

FEBRUARY 2019

S M T W Th F S

Topic connection

[Connect a client](#)[Sample code](#)[Geo-replication](#)

Sample connection code

Use this snippet of code to set the properties in your Kafka client to connect securely. Replace the values in <brackets>.

[Java](#)

```
import java.util.Properties;

import org.apache.kafka.clients.CommonClientConfigs;
import org.apache.kafka.common.config.SaslConfigs;
import org.apache.kafka.common.config.SslConfigs;


Properties properties = new Properties();
properties.put(CommonClientConfigs.BOOTSTRAP_SERVERS_CONFIG, "9.20.192.113:31934");
properties.put(CommonClientConfigs.SECURITY_PROTOCOL_CONFIG, "SASL_SSL");
properties.put(SslConfigs.SSL_PROTOCOL_CONFIG, "TLSv1.2");
properties.put(SslConfigs.SSL_TRUSTSTORE_LOCATION_CONFIG, "<certs.jks_file_location>");
properties.put(SslConfigs.SSL_TRUSTSTORE_PASSWORD_CONFIG, "<truststore_password>");
```

[Show more](#)



Sample configuration properties

Use this snippet to create a properties file for use by Kafka tools to connect securely. Replace the values in <brackets>.

Producers sending messages
1





Type to filter producers



| Kafka producer ID | Min data point | Max data point | Data range | Warnings |
|---|----------------|----------------|------------|----------|
| <input checked="" type="checkbox"/> producer-1 ServiceID=49268369-8e59-4e1c-a23e-08cdx19a2a5 | 0 B | 507.26 kB | 507.26 kB | None |

26445696 messages over the last 2 hours

Less than 60 seconds since last message produced



Effortless geo-replication

1

Geo-replication

Geo-replication is a way to duplicate your topics to different instances of Event Streams. These are usually located in different locations to minimize the risk of data loss in the event of a cluster failure.

Origin locations

Want to replicate topics to this cluster?

To add this cluster as a destination location on an origin cluster, you will need to generate connection information.

Generate connection information for this cluster

Destination locations

Add destination cluster

| | | | |
|----------|----------|-----------|---|
| rc6-repl | TOPICS 1 | WORKERS 2 | → |
|----------|----------|-----------|---|

2

← Back to geo-replication

rc6-repl TOPICS 0 WORKERS 2

Geo-replicated topics

| Name | Replication health |
|----------|--------------------|
| timtopic | Awaiting creation |

You are about to create a replicator for 1 topic.

☐ Add prefix to destination topic names ☒ Include message history

Cancel Create

3

← Back to geo-replication


rc6-repl TOPICS 0 WORKERS 2


Geo-replicated topics

| Name | Replication health |
|----------|--------------------|
| timtopic | CREATING |


Topics

🔍 Type to search topics

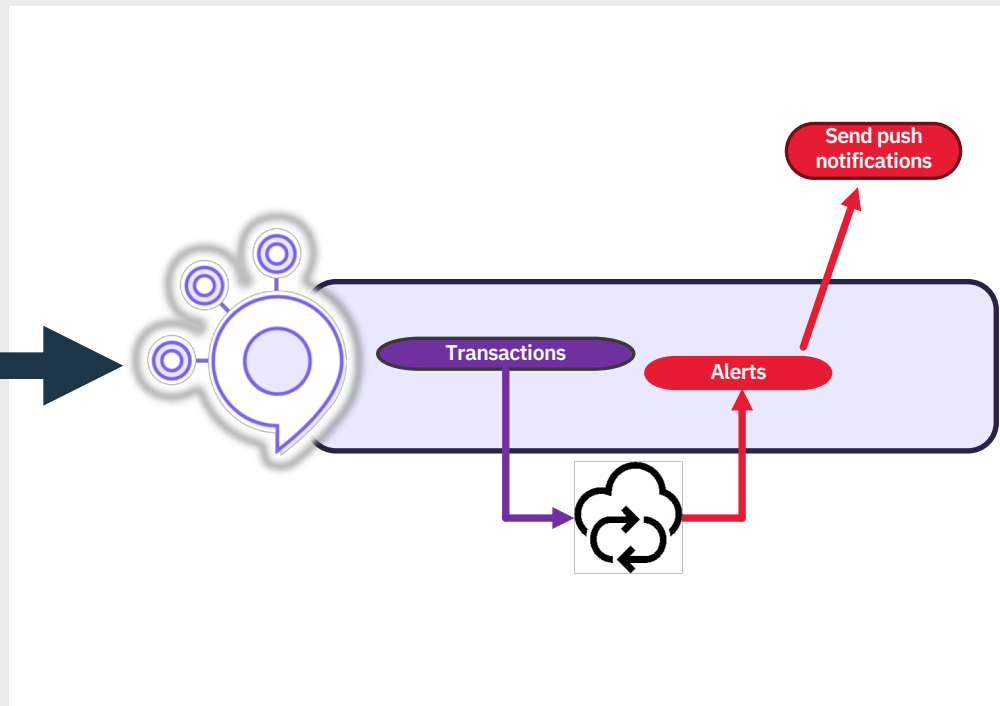
Geo-replication  Create topic +

| Name | Replicas | Partitions | |
|----------|----------|------------|---|
| testfest | 3 | 3 | ⋮ |
| timtopic | 2 | 1 |  (1) ⋮ |

Integrates with IBM MQ



IBM MQ connects mission-critical Systems of Record, requiring **transactional, once-only delivery** E.g. payment transactions



IBM Event Streams distributes and processes streams of events in real-time to intelligently engage with customers E.g. alerts on spending patterns

Find out more

- Try out Event Streams
 - <https://ibm.github.io/event-streams/installing/trying-out/> (On premise)
 - <https://console.bluemix.net/catalog/services/event-streams> (IBM hosted)
- Get in touch
 - <https://ibm.github.io/event-streams/support/>
- Apache Kafka
 - <https://kafka.apache.org/>
- Find out more
 - <https://www.ibm.com/cloud/event-streams>

Thank you

Emma Humber
IBM Event Streams

—

emma.humber@uk.ibm.com