

**Exhibit 2 to
Intellectual Ventures I LLC and Intellectual Ventures II LLC
Amended Infringement Contentions**

**Amended Infringement Claim Chart for
U.S. Patent No. 8,407,722 (“the ’722 Patent”)**

The Accused Systems and Services include, without limitation, Liberty Mutual’s systems that utilize Kafka; all past, current, and future systems and services that operate in the same or substantially similar manner as the specifically identified systems and services; and all past, current and future Liberty Mutual’s systems and services that have the same or substantially similar features as the specifically identified systems and services (“Liberty Mutual Count I Systems and Services”).¹ Upon information and belief, Liberty Mutual’s systems and services directly infringe the asserted claims of the ’722 Patent through at least using Kafka in its various systems and offerings. For example, as shown in the source code revisions below, Liberty Mutual uses Kafka for various applications. In addition, Liberty Mutual directly infringes the ’722 Patent via testing the various systems and offerings. On information and belief, Liberty Mutual also indirectly infringes the ’722 Patent through inducement. On information and belief, Liberty Mutual also contributes to the direct infringement through its vendors and other third parties’ uses of Liberty Mutual’s Systems and Services that utilize Kafka by providing components relating to Kafka that have no substantial non-infringing uses. On information and belief, Liberty Mutual with knowledge at least at the time of filing the complaint, supplies its customers, employees, and/or vendors and other third parties with access to its systems and services through various client offerings and/or internal services that allows its customers, employees, and/or vendors and other third parties to use the claimed method. On information and belief, Liberty Mutual servers, or servers that are run on behalf of Liberty Mutual, are involved and/or perform the method of the asserted claims of the ’722 Patent. On information and belief, Liberty Mutual provides instructions on how to use its various offerings and the use of which results in infringement of the ’722 Patent claims through performance of the claimed methods.

Pursuant to P.R. 3-1(b), IV identifies Kafka as used in Liberty Mutual’s systems and offerings as the Accused Instrumentality. Information regarding how specifically this Kafka is used for what product and offerings is non-public. IV intends to obtain discovery regarding how the Kafka is used in Liberty products and services and will update its infringement positions once this discovery is completed. For example, IV served a 30(b)(6) deposition notice on Liberty Mutual on June 26, 2024, seeking information relating to

¹ For the avoidance of doubt, Plaintiffs do not accuse the public clouds of Defendants, to the extent those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Defendants’ activities. [REDACTED]

[REDACTED] IV accuses Liberty Mutual private clouds that implement Kafka and non-licensed public clouds that Liberty Mutual uses to support Kafka for its systems and services. To the extent any of these licenses are relevant to Defendants’ activities, Plaintiffs will meet and confer with Defendants about the impact of such license(s).

Liberty Mutual's use and implementation of Kafka, including use and implementation of Kafka, to support services that are offered to Liberty Mutual's customers and/or employees. This information requested by IV includes for example technical documents that show the operation of Liberty Mutual's use and implementation of Kafka, Liberty Mutual source code relating to Kafka, whether third parties have information relating to Liberty Mutual's use and implementation of Kafka, and other technical information. See Plaintiffs' Initial 30(b)(6) Notice to Defendants (Liberty Mutual) at 9-10. IV served an amended 30(b)(6) notice on August 16, 2024, requesting additional relevant information. See *generally* Plaintiffs' Amended 30(b)(6) Notice to Defendants (Liberty Mutual). IV has also served 30(b)(1) deposition notices for several technical Liberty Mutual employees that likely possess knowledge of Liberty Mutual's use and implementation of Kafka. To date, Liberty Mutual has not offered any deposition dates for either potential corporate witnesses or witnesses having personal knowledge of Liberty Mutual's use and implementation of Kafka. These amended contentions are based on information available to IV as of the date of service of these contentions, and are served without the benefit of even a single technical deposition to understand how Liberty Mutual uses and implements Kafka in its systems and services. Further, fact discovery remains open and Plaintiffs intend to serve additional discovery to better understand Liberty Mutual's use and implementation of Kafka. Accordingly, Plaintiffs reserve all rights to amend these contentions consistent with the local rules. Further, IV has filed a motion to compel seeking discovery including the identification of Liberty Mutual software applications that utilize Kafka, and reserve the right to supplement these contentions once the Court has resolved this pending dispute.

U.S. Patent No. 8,407,722	
Claim 14	Liberty Mutual Count II Systems and Services
14[a][i] A method comprising: providing, using a processing device of an input source,	Upon information and belief, Liberty Mutual’s systems perform a method comprising “providing, using a processing device of an input source.” Liberty Mutual uses Kafka in multiple applications <i>for faster processing</i>. Kafka is a system containing servers, or brokers, that allow producers (e.g. processing device of an input source) to provide streams of records or events. Upon information and belief, Liberty Mutual systems leverage Kafka Streams in multiple projects/applications (e.g., accused products). The below provides a non-exhaustive list² of systems that leverage Kafka based on Liberty Mutual Technical Documentation: [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] See LM0002171.

² A more comprehensive list is present in the source code evidence section of this IC

[REDACTED]

Upon information and belief, Liberty Mutual systems leverage Kafka Streams in multiple projects/applications. For example, one such Liberty Mutual project/application shown below is Galileo Analytics. [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED].

[REDACTED]

[REDACTED]

See LM00004346.

[REDACTED]

[REDACTED]

See LM00004348.

Upon information and belief, Liberty Mutual deploys multiple Kafka Topics to facilitate the processing of data for their telematics-risk-service. CMT, a third-party system that records telemetry data on user driving habits (e.g., providing), feeds the data into several ingestion containers such as telematic-data-ingestion and daily-report-ingestion.

[REDACTED]

[REDACTED]

Further, from this trip-data-manager container, trips are sent to the Kafka topic and then sent to the trip-data-scoring container. [REDACTED]

[REDACTED] This

[REDACTED]

See LM0002201

The description below provides information on how Kafka is utilized in the exemplary systems described above for Galileo and CMT.

Upon information and belief, Liberty Mutual systems include Kafka, a distributed streaming system that runs as a cluster of one or more servers. Using the producer client and API (e.g., providing), a first entity/service/application (e.g., processing device) associated with a data server or database sends streams of records or events to the respective Kafka topic(s).

How does Kafka work in a nutshell?

Kafka is a distributed system consisting of **servers** and **clients** that communicate via a high-performance [TCP network protocol](#). It can be deployed on bare-metal hardware, virtual machines, and containers in on-premise as well as cloud environments.

Servers: Kafka is run as a cluster of one or more servers that can span multiple datacenters or cloud regions. Some of these servers form the storage layer, called the brokers. Other servers run [Kafka Connect](#) to continuously import and export data as event streams to integrate Kafka with your existing systems such as relational databases as well as other Kafka clusters. To let you implement mission-critical use cases, a Kafka cluster is highly scalable and fault-tolerant: if any of its servers fails, the other servers will take over their work to ensure continuous operations without any data loss.

Clients: They allow you to write distributed applications and microservices that read, write, and process streams of events in parallel, at scale, and in a fault-tolerant manner even in the case of network problems or machine failures. Kafka ships with some such clients included, which are augmented by [dozens of clients](#) provided by the Kafka community: clients are available for Java and Scala including the higher-level [Kafka Streams](#) library, for Go, Python, C/C++, and many other programming languages as well as REST APIs.

See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).





Figure 2-2. A simple Kafka cluster

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.


Upon information and belief, Liberty Mutual systems comprise GRM, which deploys Kafka to facilitate data flow. Client/Source Application is fed into Kafka to be streamlined into other applications/services and S3 Deployment Buckets.

	 See LM00008865. Upon information and belief, Liberty Mutual systems comprise GRM, which deploys Kafka to facilitate complex event processing. Client/Source Applications (e.g. producer) are fed into Kafka to be streamlined into a cloud service (e.g. consumer/producer), which can also receive data streams from DynamoDB, then back to Kafka.
--	--

Id.


Upon information and belief, Liberty Mutual systems comprise of GRM which deploy a Kafka CI/CD pipeline. YAML configurations from git inputted to Bamboo deployment, is connected to S3 Deployment Bucket where then the Kafka Connect API is utilized for Oauth.

Kafka Ci/Cd



See LM00008866.


Kafka: Topics, Producers, and Consumers



See <http://cloudurable.com/blog/kafka-architecture/index.html> (last accessed March 15, 2024).

Kafka has four core APIs:

- The [Producer API](#) allows an application to publish a stream of records to one or more Kafka topics.
- The [Consumer API](#) allows an application to subscribe to one or more topics and process the stream of records produced to them.
- The [Streams API](#) allows an application to act as a *stream processor*, consuming an input stream from one or more topics and producing an output stream to one or more output topics, effectively transforming the input streams to output streams.
- The [Connector API](#) allows building and running reusable producers or consumers that connect Kafka topics to existing applications or data systems. For example, a connector to a relational database might capture every change to a table.



See <https://kafka.apache.org/20/documentation.html> (last accessed March 15, 2024).


For example, the Producer API allows for a first entity/service/application to send streams of records or events to Kafka.

2.1 Producer API

The Producer API allows applications to send streams of data to topics in the Kafka cluster.


See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).

Producers are those client applications that publish (write) events to Kafka, and **consumers** are those that subscribe to (read and process) these events. In Kafka, producers and consumers are fully decoupled and agnostic of each other, which is a key design element to achieve the high scalability that Kafka is known for. For example, producers never need to wait for consumers. Kafka provides various guarantees such as the ability to process events exactly-once.



See <https://kafka.apache.org/intro> (last accessed on November 13, 2023).

As an example, the class KafkaProducer is provided for writing to the Kafka topic and publishing records to the Kafka cluster. It works with the ProducerRecord class to enable transmitting messages or values (e.g., update message, containing data for updating a property of the live object) within the network with the help of brokers.



See LM00004572.

```
public class KafkaProducer<K,V>
extends Object
implements Producer<K,V>
```

A Kafka client that publishes records to the Kafka cluster.

The producer is *thread safe* and sharing a single producer instance across threads will generally be faster than having multiple instances.

Here is a simple example of using the producer to send records with strings containing sequential numbers as the key/value pairs.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("acks", "all");
props.put("key.serializer", "org.apache.kafka.common.serialization.StringSerializer");
props.put("value.serializer", "org.apache.kafka.common.serialization.StringSerializer");

Producer<String, String> producer = new KafkaProducer<>(props);
for (int i = 0; i < 100; i++)
    producer.send(new ProducerRecord<String, String>("my-topic", Integer.toString(i), Integer.toSt

producer.close();
```

The producer consists of a pool of buffer space that holds records that haven't yet been transmitted to the server as well as a background I/O thread that is responsible for turning these records into requests and transmitting them to the cluster. Failure to close the producer after use will leak these resources.

The `send()` method is asynchronous. When called it adds the record to a buffer of pending record sends and immediately returns. This allows the producer to batch together individual records for efficiency.

See: <https://kafka.apache.org/23/javadoc/index.html?org/apache/kafka/clients/producer/KafkaProducer.html> (last accessed March 11, 2024).

The below figure represents a `ProducerRecord` class with the parameters topic, partition, key, and value.

ProducerRecord

```
public ProducerRecord(String topic,  
                      Integer partition,  
                      K key,  
                      V value)
```

Creates a record to be sent to a specified topic and partition

Parameters:

topic - The topic the record will be appended to
partition - The partition to which the record should be sent
key - The key that will be included in the record
value - The record contents

Class ProducerRecord<K,V>

```
java.lang.Object  
org.apache.kafka.clients.producer.ProducerRecord<K,V>
```

```
public class ProducerRecord<K,V>  
    extends Object
```

A key/value pair to be sent to Kafka. This consists of a topic name to which the record is being sent, an optional partition number, and an optional key and value.

If a valid partition number is specified that partition will be used when sending the record. If no partition is specified but a key is present a partition will be chosen using a hash of the key. If neither key nor partition is present a partition will be assigned in a round-robin fashion. The record also has an associated timestamp. If the user did not provide a timestamp, the producer will stamp the record with its current time. The timestamp eventually used by Kafka depends on the timestamp type configured for the topic.

If the topic is configured to use `CreateTime`, the timestamp in the producer record will be used by the broker.
If the topic is configured to use `LogAppendTime`, the timestamp in the producer record will be overwritten by the broker with the broker local time when it appends the message to its log.
In either of the cases above, the timestamp that has actually been used will be returned to user in `RecordMetadata`.

See <https://kafka.apache.org/0102/javadoc/org/apache/kafka/clients/producer/ProducerRecord.html> (last accessed March 15, 2024).

Liberty Mutual Bank implements Kafka to provide end users with updates in real time.

See <https://github.com/libertymutual/provider-kafka> (last accessed on November 13, 2023).

See <https://www.linkedin.com/in/robert-scibelli/> (last accessed on November 10, 2023).

See <https://searchjobs.libertymutualgroup.com/careers/job/618496325540?microsite=libertymutual.com> (last accessed on March 19, 2024).

[illegible]

A horizontal bar chart with 28 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Very good, 4. Good, 5. Fair, 6. Poor, 7. Very poor, 8. Excellent, 9. Very good, 10. Good, 11. Fair, 12. Poor, 13. Very poor, 14. Excellent, 15. Very good, 16. Good, 17. Fair, 18. Poor, 19. Very poor, 20. Excellent, 21. Very good, 22. Good, 23. Fair, 24. Poor, 25. Very poor, 26. Excellent, 27. Very good, 28. Good. The percentages are: 1. 1%, 2. 2%, 3. 3%, 4. 4%, 5. 5%, 6. 6%, 7. 7%, 8. 8%, 9. 9%, 10. 10%, 11. 11%, 12. 12%, 13. 13%, 14. 14%, 15. 15%, 16. 16%, 17. 17%, 18. 18%, 19. 19%, 20. 20%, 21. 21%, 22. 22%, 23. 23%, 24. 24%, 25. 25%, 26. 26%, 27. 27%, 28. 28%.

Category	Percentage
1. Other	1%
2. Not sure	2%
3. Very good	3%
4. Good	4%
5. Fair	5%
6. Poor	6%
7. Very poor	7%
8. Excellent	8%
9. Very good	9%
10. Good	10%
11. Fair	11%
12. Poor	12%
13. Very poor	13%
14. Excellent	14%
15. Very good	15%
16. Good	16%
17. Fair	17%
18. Poor	18%
19. Very poor	19%
20. Excellent	20%
21. Very good	21%
22. Good	22%
23. Fair	23%
24. Poor	24%
25. Very poor	25%
26. Excellent	26%
27. Very good	27%
28. Good	28%

[illegible]

[illegible]

A horizontal bar chart with 25 categories listed on the left and corresponding percentage values represented by black bars on the right. The categories are: 1. Very good, 2. Good, 3. Fair, 4. Poor, 5. Very poor, 6. Don't know, 7. No answer, 8. Not applicable, 9. Not relevant, 10. Not possible, 11. Not necessary, 12. Not desirable, 13. Not feasible, 14. Not acceptable, 15. Not suitable, 16. Not appropriate, 17. Not effective, 18. Not efficient, 19. Not economical, 20. Not practical, 21. Not realistic, 22. Not sustainable, 23. Not ethical, 24. Not legal, 25. Not moral. The percentages are: 1. 38%, 2. 35%, 3. 32%, 4. 28%, 5. 25%, 6. 22%, 7. 20%, 8. 18%, 9. 15%, 10. 12%, 11. 10%, 12. 8%, 13. 7%, 14. 6%, 15. 5%, 16. 4%, 17. 3%, 18. 2%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%.

Category	Percentage
1. Very good	38%
2. Good	35%
3. Fair	32%
4. Poor	28%
5. Very poor	25%
6. Don't know	22%
7. No answer	20%
8. Not applicable	18%
9. Not relevant	15%
10. Not possible	12%
11. Not necessary	10%
12. Not desirable	8%
13. Not feasible	7%
14. Not acceptable	6%
15. Not suitable	5%
16. Not appropriate	4%
17. Not effective	3%
18. Not efficient	2%
19. Not economical	1%
20. Not practical	1%
21. Not realistic	1%
22. Not sustainable	1%
23. Not ethical	1%
24. Not legal	1%
25. Not moral	1%

[illegible]

[illegible]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages shown as black bars on the right. The percentages are as follows:

Category	Percentage
1. Very high	35
2. High	25
3. Medium-high	40
4. Medium	20
5. Medium-low	25
6. Low	30
7. Very low	35
8. Not at all	30
9. Don't know	25
10. No answer	20
11. Very high	15
12. High	20
13. Medium-high	15
14. Medium	25
15. Medium-low	15
16. Low	10
17. Very low	15
18. Not at all	10
19. Don't know	15
20. No answer	10
21. Very high	20
22. High	25
23. Medium-high	30
24. Medium	35
25. Medium-low	30
26. Low	25
27. Very low	20
28. Not at all	15
29. Don't know	10
30. No answer	5

[illegible]

[illegible]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. Very good, 2. Good, 3. Fair, 4. Poor, 5. Very poor, 6. Don't know, 7. No answer, 8. Not applicable, 9. Not relevant, 10. Not necessary, 11. Not possible, 12. Not desirable, 13. Not feasible, 14. Not realistic, 15. Not acceptable, 16. Not suitable, 17. Not appropriate, 18. Not effective, 19. Not efficient, 20. Not economical, 21. Not practical, 22. Not convenient, 23. Not comfortable, 24. Not safe, 25. Not healthy. The percentages are: 1. 45%, 2. 40%, 3. 35%, 4. 30%, 5. 25%, 6. 20%, 7. 15%, 8. 10%, 9. 5%, 10. 4%, 11. 3%, 12. 2%, 13. 1%, 14. 1%, 15. 1%, 16. 1%, 17. 1%, 18. 1%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%.

Category	Percentage
1. Very good	45%
2. Good	40%
3. Fair	35%
4. Poor	30%
5. Very poor	25%
6. Don't know	20%
7. No answer	15%
8. Not applicable	10%
9. Not relevant	5%
10. Not necessary	4%
11. Not possible	3%
12. Not desirable	2%
13. Not feasible	1%
14. Not realistic	1%
15. Not acceptable	1%
16. Not suitable	1%
17. Not appropriate	1%
18. Not effective	1%
19. Not efficient	1%
20. Not economical	1%
21. Not practical	1%
22. Not convenient	1%
23. Not comfortable	1%
24. Not safe	1%
25. Not healthy	1%

[illegible]

[illegible]

[REDACTED]

.		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]

	• [REDACTED]
14[a][ii] a data representation to a client device, different from the input source, coupled to a routing network,	Upon information and belief, Liberty Mutual’s systems include “a data representation to a client device, different from the input source, coupled to a routing network.” Kafka is a distributed streaming system that is run as a cluster of one or more servers. Using the producer client and API, a first entity/service/application sends streams of records or events (e.g., data representation) to respective Kafka topics. These records or events, as an example, include dynamic objects such as transactions. They are routed to a destination entity/service/application (e.g., client device) as a consumer client with the consumer API via brokers in the network (e.g., coupled to a routing network).

How does Kafka work in a nutshell?

Kafka is a distributed system consisting of **servers** and **clients** that communicate via a high-performance [TCP network protocol](#). It can be deployed on bare-metal hardware, virtual machines, and containers in on-premise as well as cloud environments.

Servers: Kafka is run as a cluster of one or more servers that can span multiple datacenters or cloud regions. Some of these servers form the storage layer, called the brokers. Other servers run [Kafka Connect](#) to continuously import and export data as event streams to integrate Kafka with your existing systems such as relational databases as well as other Kafka clusters. To let you implement mission-critical use cases, a Kafka cluster is highly scalable and fault-tolerant: if any of its servers fails, the other servers will take over their work to ensure continuous operations without any data loss.

Clients: They allow you to write distributed applications and microservices that read, write, and process streams of events in parallel, at scale, and in a fault-tolerant manner even in the case of network problems or machine failures. Kafka ships with some such clients included, which are augmented by [dozens of clients](#) provided by the Kafka community: clients are available for Java and Scala including the higher-level [Kafka Streams](#) library, for Go, Python, C/C++, and many other programming languages as well as REST APIs.

See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).

Producers are those client applications that publish (write) events to Kafka, and **consumers** are those that subscribe to (read and process) these events. In Kafka, producers and consumers are fully decoupled and agnostic of each other, which is a key design element to achieve the high scalability that Kafka is known for. For example, producers never need to wait for consumers. Kafka provides various [guarantees](#) such as the ability to process events exactly-once.

See <https://kafka.apache.org/intro> (last accessed on November 13, 2023).

[REDACTED]

2.2 Consumer API

The Consumer API allows applications to read streams of data from topics in the Kafka cluster.


Examples showing how to use the consumer are given in the [javadocs](#).

To use the consumer, you can use the following maven dependency:

See <https://kafka.apache.org/documentation.html> (last accessed March 15, 2024).

Upon information and belief, Liberty Mutual systems leverage Kafka Streams in the following projects/applications. For example, one such Liberty Mutual project/application is Galileo Analytics, which is used to process user interactions on the browser [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]



See LM00004348.

For example, many destination entities/services/applications (e.g. client device) can interact with a Kafka consumer instance to connect to the Kafka cluster. For all consumer instances reading from the same topic, they are grouped with distinctive IDs. An even distribution assignment of partitions of the topics will happen across each consumer group. The consumer API offers the functionality of offset committing to provide an accurate consumption of the messages in a topic stream.

[Redacted]

Id.

[Redacted]

[Redacted]

Id.

For example, Kafka producer and consumer client applications communicate in a cluster through the Kafka brokers via a network (e.g. routing network) to read/write events to respective topics.

A Kafka client communicates with the Kafka brokers via the network for writing (or reading) events. Once received, the brokers will store the events in a durable and fault-tolerant manner for as long as you need—even forever.

See <https://kafka.apache.org/quickstart> (last accessed July 26, 2024).

Source Code Evidence

[Redacted]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 'Other', 'No answer', 'Don't know', 'Yes, definitely', 'Yes, probably', 'Yes, maybe', 'Yes, not sure', 'Yes, not at all', 'No, definitely', 'No, probably', 'No, maybe', 'No, not sure', 'No, not at all', 'Don't know', 'No answer', 'Other', 'Yes, definitely', 'Yes, probably', 'Yes, maybe', 'Yes, not sure', 'Yes, not at all', 'No, definitely', 'No, probably', 'No, maybe', 'No, not sure', 'No, not at all'. The percentages are: 10%, 15%, 20%, 25%, 30%, 35%, 30%, 25%, 20%, 15%, 10%, 5%, 5%, 5%, 5%, 5%, 10%, 15%, 20%, 25%, 30%, 35%, 30%, 25%, 20%, 15%, 10%.

Category	Percentage
Other	10%
No answer	15%
Don't know	20%
Yes, definitely	25%
Yes, probably	30%
Yes, maybe	35%
Yes, not sure	30%
Yes, not at all	25%
No, definitely	20%
No, probably	15%
No, maybe	10%
No, not sure	5%
No, not at all	5%
Don't know	5%
No answer	5%
Other	10%
Yes, definitely	15%
Yes, probably	20%
Yes, maybe	25%
Yes, not sure	30%
Yes, not at all	35%
No, definitely	30%
No, probably	25%
No, maybe	20%
No, not sure	15%
No, not at all	10%

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. (10%), 2. (15%), 3. (20%), 4. (25%), 5. (30%), 6. (35%), 7. (40%), 8. (45%), 9. (10%), 10. (15%), 11. (20%), 12. (25%), 13. (30%), 14. (35%), 15. (40%), 16. (45%), 17. (10%), 18. (15%), 19. (20%), 20. (25%), 21. (30%), 22. (35%), 23. (40%), 24. (45%), 25. (10%).

Category	Percentage
1.	10%
2.	15%
3.	20%
4.	25%
5.	30%
6.	35%
7.	40%
8.	45%
9.	10%
10.	15%
11.	20%
12.	25%
13.	30%
14.	35%
15.	40%
16.	45%
17.	10%
18.	15%
19.	20%
20.	25%
21.	30%
22.	35%
23.	40%
24.	45%
25.	10%

A horizontal bar chart consisting of 30 black bars of varying lengths. The bars are arranged in a single column, with the longest bar in the middle and the shortest at the top and bottom. The bars represent data points, likely percentages, as indicated by the 'Percentage' label in the adjacent column.

[illegible]

[illegible]

[illegible]

[illegible]

A horizontal bar chart with 28 categories listed on the left and their corresponding percentages shown as black bars on the right. The percentages are as follows:

Category	Percentage
1. Not at all	1%
2. Not very much	1%
3. Somewhat	1%
4. A great deal	1%
5. Not at all	1%
6. Not very much	1%
7. Somewhat	1%
8. A great deal	1%
9. Not at all	1%
10. Not very much	1%
11. Somewhat	1%
12. A great deal	1%
13. Not at all	1%
14. Not very much	1%
15. Somewhat	1%
16. A great deal	1%
17. Not at all	1%
18. Not very much	1%
19. Somewhat	1%
20. A great deal	1%
21. Not at all	1%
22. Not very much	1%
23. Somewhat	1%
24. A great deal	1%
25. Not at all	1%
26. Not very much	1%
27. Somewhat	1%
28. A great deal	1%

	• [REDACTED]
14[a][iii] wherein the data representation includes at least one live object recognizable by the client device,	Upon information and belief, Liberty Mutual’s systems are configured such that “wherein the data representation includes at least one live object recognizable by the client device.” Kafka is a distributed streaming system that is run as a cluster of one or more servers. Using the producer client and API, a first entity/service/application associated with a data server sends streams of records or events (e.g., data representation) to respective Kafka topics. When a producer reads or writes to Kafka, they are writing an event(s) that consists of something that has occurred (e.g. at least one live object). Each new

event/record updates a previous event/record within a topic. Consumer entity/service/application (e.g., client device) that subscribes to the said topic will receive the updated event occurrence.

Kafka is a distributed *event streaming platform* that lets you read, write, store, and process *events* (also called *records* or *messages* in the documentation) across many machines.

Example events are payment transactions, geolocation updates from mobile phones, shipping orders, sensor measurements from IoT devices or medical equipment, and much more. These events are organized and stored in *topics*. Very simplified, a topic is similar to a folder in a filesystem, and the events are the files in that folder.

See <https://kafka.apache.org/quickstart> (last accessed on November 13, 2023).

An **event** records the fact that "something happened" in the world or in your business. It is also called record or message in the documentation. When you read or write data to Kafka, you do this in the form of events. Conceptually, an event has a key, value, timestamp, and optional metadata headers. Here's an example event:

- Event key: "Alice"
- Event value: "Made a payment of \$200 to Bob"
- Event timestamp: "Jun. 25, 2020 at 2:06 p.m."

Topics are **partitioned**, meaning a topic is spread over a number of "buckets" located on different Kafka brokers. This distributed placement of your data is very important for scalability because it allows client applications to both read and write the data from/to many brokers at the same time. When a new event is published to a topic, it is actually appended to one of the topic's partitions. Events with the same event key (e.g., a customer or vehicle ID) are written to the same partition, and Kafka guarantees that any consumer of a given topic-partition will always read that partition's events in exactly the same order as they were written.





Figure: This example topic has four partitions P1–P4. Two different producer clients are publishing, independently from each other, new events to the topic by writing events over the network to the topic's partitions. Events with the same key (denoted by their color in the figure) are written to the same partition. Note that both producers can write to the same partition if appropriate.

See <https://kafka.apache.org/intro> (last accessed on November 13, 2023)

For example, destination entities/services/applications (e.g. client device) can initialize a Kafka consumer instance to connect to the Kafka cluster. For all consumer instances reading from the same topic, they are grouped with distinctive IDs. An even assignment of partitions of the topics will be distributed across each consumer group. The consumer API offers the functionality of offset committing to provide an accurate consumption of the messages in a topic stream.



See LM00004571.

Each consumer in a group can dynamically set the list of topics it wants to subscribe to through one of the `subscribe` APIs. Kafka will deliver each message in the subscribed topics to one process in each consumer group. This is achieved by balancing the partitions between all members in the consumer group so that each partition is assigned to exactly one consumer in the group. So if there is a topic with four partitions, and a consumer group with two processes, each process would consume from two partitions.

See <https://kafka.apache.org/10/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html> (last accessed March 15, 2024).

Usage Examples

The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them.

Automatic Offset Committing

This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("group.id", "test");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props);
consumer.subscribe(Arrays.asList("foo", "bar"));
while (true) {
    ConsumerRecords<String, String> records = consumer.poll(100);
    for (ConsumerRecord<String, String> record : records)
        System.out.printf("offset = %d, key = %s, value = %s\n", record.offset(), record.key(), record.value());
}
```

The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration `>bootstrap.servers`. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting).

Setting `enable.auto.commit` means that offsets are committed automatically with a frequency controlled by the config `auto.commit.interval.ms`.

In this example the consumer is subscribing to the topics `foo` and `bar` as part of a group of consumers called `test` as configured with `group.id`.

See <https://kafka.apache.org/documentation.html> (last accessed March 15, 2024).

Source Code Evidence

[REDACTED]

	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]

[illegible]

[illegible]

[illegible]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Yes, 4. No, 5. Yes, 6. No, 7. Yes, 8. No, 9. Yes, 10. No, 11. Yes, 12. No, 13. Yes, 14. No, 15. Yes, 16. No, 17. Yes, 18. No, 19. Yes, 20. No, 21. Yes, 22. No, 23. Yes, 24. No, 25. Yes. The percentages are: 1. 10%, 2. 15%, 3. 45%, 4. 12%, 5. 18%, 6. 10%, 7. 15%, 8. 12%, 9. 18%, 10. 10%, 11. 15%, 12. 12%, 13. 18%, 14. 10%, 15. 15%, 16. 12%, 17. 18%, 18. 10%, 19. 15%, 20. 12%, 21. 18%, 22. 10%, 23. 15%, 24. 12%, 25. 18%.

Category	Percentage
1. Other	10%
2. Not sure	15%
3. Yes	45%
4. No	12%
5. Yes	18%
6. No	10%
7. Yes	15%
8. No	12%
9. Yes	18%
10. No	10%
11. Yes	15%
12. No	12%
13. Yes	18%
14. No	10%
15. Yes	15%
16. No	12%
17. Yes	18%
18. No	10%
19. Yes	15%
20. No	12%
21. Yes	18%
22. No	10%
23. Yes	15%
24. No	12%
25. Yes	18%

[illegible]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. Very good, 2. Good, 3. Fair, 4. Poor, 5. Very poor, 6. Don't know, 7. No answer, 8. Not applicable, 9. Not relevant, 10. Not necessary, 11. Not possible, 12. Not desirable, 13. Not feasible, 14. Not realistic, 15. Not acceptable, 16. Not suitable, 17. Not appropriate, 18. Not effective, 19. Not efficient, 20. Not economical, 21. Not practical, 22. Not convenient, 23. Not comfortable, 24. Not safe, 25. Not healthy. The percentages are: 1. 35%, 2. 30%, 3. 25%, 4. 20%, 5. 15%, 6. 10%, 7. 5%, 8. 3%, 9. 2%, 10. 1%, 11. 1%, 12. 1%, 13. 1%, 14. 1%, 15. 1%, 16. 1%, 17. 1%, 18. 1%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%.

Category	Percentage
1. Very good	35%
2. Good	30%
3. Fair	25%
4. Poor	20%
5. Very poor	15%
6. Don't know	10%
7. No answer	5%
8. Not applicable	3%
9. Not relevant	2%
10. Not necessary	1%
11. Not possible	1%
12. Not desirable	1%
13. Not feasible	1%
14. Not realistic	1%
15. Not acceptable	1%
16. Not suitable	1%
17. Not appropriate	1%
18. Not effective	1%
19. Not efficient	1%
20. Not economical	1%
21. Not practical	1%
22. Not convenient	1%
23. Not comfortable	1%
24. Not safe	1%
25. Not healthy	1%

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages shown as black bars on the right. The categories are: 1. Very good, 2. Good, 3. Fair, 4. Poor, 5. Very poor, 6. Don't know, 7. No answer, 8. Not applicable, 9. Not relevant, 10. Not possible, 11. Not necessary, 12. Not useful, 13. Not effective, 14. Not efficient, 15. Not safe, 16. Not healthy, 17. Not comfortable, 18. Not convenient, 19. Not affordable, 20. Not accessible, 21. Not reliable, 22. Not secure, 23. Not legal, 24. Not ethical, 25. Not sustainable. The percentages are: 1. 35%, 2. 25%, 3. 20%, 4. 15%, 5. 10%, 6. 5%, 7. 3%, 8. 2%, 9. 1%, 10. 1%, 11. 1%, 12. 1%, 13. 1%, 14. 1%, 15. 1%, 16. 1%, 17. 1%, 18. 1%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%.


Category	Percentage
1. Very good	35%
2. Good	25%
3. Fair	20%
4. Poor	15%
5. Very poor	10%
6. Don't know	5%
7. No answer	3%
8. Not applicable	2%
9. Not relevant	1%
10. Not possible	1%
11. Not necessary	1%
12. Not useful	1%
13. Not effective	1%
14. Not efficient	1%
15. Not safe	1%
16. Not healthy	1%
17. Not comfortable	1%
18. Not convenient	1%
19. Not affordable	1%
20. Not accessible	1%
21. Not reliable	1%
22. Not secure	1%
23. Not legal	1%
24. Not ethical	1%
25. Not sustainable	1%

	• [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
14[a][iv] and causing the client device to respond to the live object of the data representation by determining an object identifier (ID) of the live object	Upon information and belief, Liberty Mutual’s systems include “causing the client device to respond to the live object of the data representation by determining an object identifier (ID) of the live object.” When a producer reads or writes an event(s) to Kafka, they are written to topics. A topic is a timeline of events, where each contains an event key or an identifier (e.g., object identifier (ID) of the live object) and is divided into several partitions. Kafka facilitates the receipt of events where destination entity/service/applications (e.g. client device(s)) subscribed to that topic can read through the stream and consume the events (e.g., respond to the live object). An event records the fact that "something happened" in the world or in your business. It is also called record or message in the documentation. When you read or write data to Kafka, you do this in the form of events. Conceptually, an event has a key, value, timestamp, and optional metadata headers. Here's an example event: • Event key: "Alice" • Event value: "Made a payment of \$200 to Bob" • Event timestamp: "Jun. 25, 2020 at 2:06 p.m."

See <https://kafka.apache.org/intro> (last accessed on November 13, 2023)

Partitions

Kafka's topics are divided into several partitions. While the topic is a logical concept in Kafka, a partition is the smallest storage unit that holds a subset of records owned by a topic. Each partition is a single log file where records are written to it in an append-only fashion.



When talking about the content inside a partition, I will use the terms **record** and **message** interchangeably.

Offsets and the ordering of messages

The records in the partitions are each assigned a sequential identifier called the offset, which is unique for each record within the partition.

The offset is an incremental and immutable number, maintained by Kafka. When a record is written to a partition, it is appended to the end of the log, assigning the next sequential offset. Offsets are particularly useful for consumers when reading records from a partition. We'll come to that at a later point.

See <https://medium.com/event-driven-utopia/understanding-kafka-topic-partitions-ae40f80552e8> (last accessed March 11, 2024).

For example, Kafka destination entities/services/applications (e.g. client device) initializing a consumer instance subscribing to topics are noted as subscribers. Consumers subscribing to topics will be grouped to

keep track, and the consumer API offers the functionality of offset assignments to provide recovery in case of any disconnects or failures.

org.apache.kafka.clients.consumer

Class KafkaConsumer<K,V>

java.lang.Object

org.apache.kafka.clients.consumer.KafkaConsumer<K,V>

See <https://kafka.apache.org/22/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html> (last accessed March 15, 2024).

Each consumer in a group can dynamically set the list of topics it wants to subscribe to through one of the `subscribe` APIs. Kafka will deliver each message in the subscribed topics to one process in each consumer group. This is achieved by balancing the partitions between all members in the consumer group so that each partition is assigned to exactly one consumer in the group. So if there is a topic with four partitions, and a consumer group with two processes, each process would consume from two partitions.

Id.

called subscribers or readers. The consumer subscribes to one or more topics and reads the messages in the order in which they were produced to each partition. The consumer keeps track of which messages it has already consumed by keeping track of the offset of messages. The *offset*—an integer value that continually increases—is another piece of metadata that Kafka adds to each message as it is produced. Each message in a given partition has a unique offset, and the following message has a greater offset (though not necessarily monotonically greater). By storing the next possible offset for each partition, typically in Kafka itself, a consumer can stop and restart without losing its place.

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

Source Code Evidence

[REDACTED]

	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

A horizontal bar chart with 28 categories listed on the left and their corresponding percentages shown as black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Don't know, 4. No answer, 5. No opinion, 6. No response, 7. No comment, 8. No information, 9. No data, 10. No results, 11. No findings, 12. No conclusions, 13. No recommendations, 14. No suggestions, 15. No advice, 16. No help, 17. No support, 18. No assistance, 19. No guidance, 20. No direction, 21. No motivation, 22. No inspiration, 23. No encouragement, 24. No support, 25. No assistance, 26. No guidance, 27. No direction, 28. No motivation. The percentages are: 1. 1%, 2. 1%, 3. 1%, 4. 1%, 5. 1%, 6. 1%, 7. 1%, 8. 1%, 9. 1%, 10. 1%, 11. 1%, 12. 1%, 13. 1%, 14. 1%, 15. 1%, 16. 1%, 17. 1%, 18. 1%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%, 26. 1%, 27. 1%, 28. 1%.

Category	Percentage
1. Other	1%
2. Not sure	1%
3. Don't know	1%
4. No answer	1%
5. No opinion	1%
6. No response	1%
7. No comment	1%
8. No information	1%
9. No data	1%
10. No results	1%
11. No findings	1%
12. No conclusions	1%
13. No recommendations	1%
14. No suggestions	1%
15. No advice	1%
16. No help	1%
17. No support	1%
18. No assistance	1%
19. No guidance	1%
20. No direction	1%
21. No motivation	1%
22. No inspiration	1%
23. No encouragement	1%
24. No support	1%
25. No assistance	1%
26. No guidance	1%
27. No direction	1%
28. No motivation	1%

[illegible]

[illegible]

	• [REDACTED]
14[a][v] and to register for updates of the live object with the routing network, such that registering the client device with the routing network provides client connection information to a node in the routing network; and	Upon information and belief, Liberty Mutual’s systems include a “register for updates of the live object with the routing network, such that registering the client device with the routing network provides client connection information to a node in the routing network.” A consumer can subscribe to a topic, or topics, (e.g. register for updates of the live object) and Kafka brokers (e.g. node in the routing network) facilitate the receipt of those events to consumer entity/service/application user devices as they read through the topic stream. Kafka uses a network protocol such as TCP to facilitate message requests and delivery through the Kafka brokers. A Kafka client communicates with the Kafka brokers via the network for writing (or reading) events. Once received, the brokers will store the events in a durable and fault-tolerant manner for as long as you need—even forever. See https://kafka.apache.org/quickstart (last accessed July 26, 2024). Each consumer in a group can dynamically set the list of topics it wants to subscribe to through one of the <code>subscribe</code> APIs. Kafka will deliver each message in the subscribed topics to one process in each consumer group. This is achieved by balancing the partitions between all members in the consumer group so that each partition is assigned to exactly one consumer in the group. So if there is a topic with four partitions, and a consumer group with two processes, each process would consume from two partitions.

See <https://kafka.apache.org/10/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html> (last accessed March 15, 2024).

Kafka uses a binary protocol over TCP. The protocol defines all apis as request response message pairs. All messages are size delimited and are made up of the following primitive types.

The client initiates a socket connection and then writes a sequence of request messages and reads back the corresponding response message. No handshake is required on connection or disconnection. TCP is happier if you maintain persistent connections used for many requests to amortize the cost of the TCP handshake, but beyond this penalty connecting is pretty cheap.

The client will likely need to maintain a connection to multiple brokers, as data is partitioned and the clients will need to talk to the server that has their data. However it should not generally be necessary to maintain multiple connections to a single broker from a single client instance (i.e. connection pooling).

The server guarantees that on a single TCP connection, requests will be processed in the order they are sent and responses will return in that order as well. The broker's request processing allows only a single in-flight request per connection in order to guarantee this ordering. Note that clients can (and ideally should) use non-blocking IO to implement request pipelining and achieve higher throughput. i.e., clients can send requests even while awaiting responses for preceding requests since the outstanding requests will be buffered in the underlying OS socket buffer. All requests are initiated by the client, and result in a corresponding response message from the server except where noted.

The server has a configurable maximum limit on request size and any request that exceeds this limit will result in the socket being disconnected.

See https://kafka.apache.org/0100/protocol.html#protocol_network (last accessed July 29, 2024).

As an example, Kafka Streams applications read through a changelog stream or table of records with unique keys. When a new event is processed, the state of the table updates as well with a new column addition.

KTable is an abstraction of a *changelog stream* from a primary-keyed table. Each record in this changelog stream is an update on the primary-keyed table with the record key as the primary key.

See <https://kafka.apache.org/23/javadoc/org/apache/kafka/streams/kstream/KTable.html> (last accessed on November 10, 2023).

Let's illustrate this with an example. Imagine a table that tracks the total number of pageviews by user (first column of diagram below). Over time, whenever a new pageview event is processed, the state of the table is updated accordingly. Here, the state changes between different points in time – and different revisions of the table – can be represented as a changelog stream (second column).

See <https://docs.confluent.io/platform/current/streams/concepts.html> (last accessed on November 10, 2023)

As an example, Kafka can use DNS to facilitate connectivity of any producer or consumer entities/services/applications (e.g., client device). On information and belief, for example, a first connection made by a consumer will connect to a bootstrap server, where the broker will provide metadata regarding other details of the Kafka cluster. DNS resolution provides registering the client device with the routing network on first connection (e.g., registering the client) which provides client connection information to a node in the routing network.

DNS Resolution for Kafka clients

DNS resolution is an important aspect of Kafka client connectivity as it enables the clients to connect to the correct Kafka brokers. When a Kafka client is configured to connect to a Kafka cluster, it typically specifies a list of broker hostnames or IP addresses in its configuration.

When a Kafka client is started, it first reads its configuration file to determine the initial list of bootstrap brokers. These brokers are specified as a comma-separated list of hostname and port pairs (e.g., `broker1.example.com:9092,broker2.example.com:9092`).

The Kafka client then performs DNS resolution to obtain the IP addresses of the specified hostnames. By default, the client uses the operating system's default DNS resolution behavior, which typically involves sending DNS queries to a DNS resolver provided by the operating system.

See <https://levelup.gitconnected.com/dns-resolution-for-kafka-clients-2e76b28e4717> (last accessed March 15, 2024).

Source Code Evidence

[REDACTED]

[illegible]

[Redacted]

	■	[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
	■	[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]
		[Redacted]

[Redacted]

		[Redacted] [Redacted] [Redacted] [Redacted] [Redacted] [Redacted] [Redacted] [Redacted] [Redacted]
	■	[Redacted] [Redacted]
	■	[Redacted] [Redacted] [Redacted] [Redacted] [Redacted] [Redacted]
	■	[Redacted] [Redacted] [Redacted]
	■	[Redacted] [Redacted] [Redacted]
	■	[Redacted] [Redacted] [Redacted]
	■	[Redacted] [Redacted]

	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
14[b] sending, using the processing device of the input source, an update message to the routing network, wherein the update message identifies the live object and contains update data that updates a property of the live object;	Upon information and belief, Liberty Mutual’s systems perform “sending, using the processing device of the input source, an update message to the routing network, wherein the update message identifies the live object and contains update data that updates a property of the live object.” Each new event/record that is written to a Kafka topic from a first entity/service/application (e.g., sending, using the processing device of an input source) comprises content (e.g. update data) that includes an event key (e.g. update message identifies the live object). The ProducerRecord function includes several parameters for the topic, partition, key, and value to which the record belongs to. ProducerRecord <pre>public ProducerRecord(String topic, Integer partition, K key, V value)</pre> Creates a record to be sent to a specified topic and partition Parameters: - topic - The topic the record will be appended to - partition - The partition to which the record should be sent - key - The key that will be included in the record - value - The record contents See https://kafka.apache.org/0102/javadoc/org/apache/kafka/clients/producer/ProducerRecord.html (last accessed March 15, 2024).





Figure 2-2. A simple Kafka cluster

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

Kafka: Topics, Producers, and Consumers



See <http://cloudurable.com/blog/kafka-architecture/index.html> (last accessed March 15, 2024).

For example, using the *send* function, a record or event is asynchronously sent to its corresponding topic. The key for each event is used to determine the corresponding partition/topic it should be written to.

send

```
public Future<RecordMetadata> send(ProducerRecord<K,V> record)
```

Asynchronously send a record to a topic. Equivalent to `send(record, null)`. See `send(ProducerRecord, Callback)` for details.

Specified by:

`send` in interface `Producer<K,V>`

Parameters:

`record` - The record to send

Returns:

A future which will eventually contain the response information

See <https://kafka.apache.org/0100/javadoc/org/apache/kafka/clients/producer/KafkaProducer.html> (last accessed on November 10, 2023).

For example, records or events include dynamic objects, such as transactions, and make up respective topic streams that can be partitioned across multiple Kafka brokers as they are written by a first entity/service/application (e.g. sending, using the processing device of the input source). For consumer clients subscribed to and reading these topic streams, the API provides continuous coordination and rebalancing on how these topics are partitioned.

An **event** records the fact that "something happened" in the world or in your business. It is also called record or message in the documentation. When you read or write data to Kafka, you do this in the form of events. Conceptually, an event has a key, value, timestamp, and optional metadata headers. Here's an example event:

- Event key: "Alice"
- Event value: "Made a payment of \$200 to Bob"
- Event timestamp: "Jun. 25, 2020 at 2:06 p.m."

Topics are **partitioned**, meaning a topic is spread over a number of "buckets" located on different Kafka brokers. This distributed placement of your data is very important for scalability because it allows client applications to both read and write the data from/to many brokers at the same time. When a new event is published to a topic, it is actually appended to one of the topic's partitions. Events with the same event key (e.g., a customer or vehicle ID) are written to the same partition, and Kafka guarantees that any consumer of a given topic-partition will always read that partition's events in exactly the same order as they were written.




Figure: This example topic has four partitions P1–P4. Two different producer clients are publishing, independently from each other, new events to the topic by writing events over the network to the topic's partitions. Events with the same key (denoted by their color in the figure) are written to the same partition. Note that both producers can write to the same partition if appropriate.

See <https://kafka.apache.org/> (last accessed on November 13, 2023).

At the heart of the consumer API is a simple loop for polling the server for more data. Once the consumer subscribes to topics, the poll loop handles all details of coordination, partition rebalances, heartbeats, and data fetching, leaving the developer with a clean API that simply returns available data from the assigned partitions. The main body of a consumer will look as follows:

```
try {
    while (true) { ❶
        ConsumerRecords<String, String> records = consumer.poll(100); ❷
        for (ConsumerRecord<String, String> record : records) ❸
        {
            log.debug("topic = %s, partition = %d, offset = %d,"
                + "customer = %s, country = %s\n",
                record.topic(), record.partition(), record.offset(),
                record.key(), record.value());

            int updatedCount = 1;
            if (custCountryMap.containsKey(record.value())) {
                updatedCount = custCountryMap.get(record.value()) + 1;
            }
            custCountryMap.put(record.value(), updatedCount)

            JSONObject json = new JSONObject(custCountryMap);
            System.out.println(json.toString(4)) ❹
        }
    }
}
```

See <https://www.oreilly.com/library/view/kafka-the-definitive/9781491936153/ch04.html> (last accessed March 15, 2024).

The Kafka Streams API in a Nutshell

The Streams API of Kafka, available through a Java library, can be used to **build highly scalable, elastic, fault-tolerant, distributed applications, and microservices**. First and foremost, the Kafka Streams API allows you **to create real-time applications that power your core business**. It is the easiest yet the most powerful technology to process data stored in Kafka. It builds upon important **concepts** for **stream processing** such as efficient management of application state, fast and efficient aggregations and joins, properly distinguishing between event-time and processing-time, and seamless handling of **out-of-order data**.

A unique feature of the Kafka Streams API is that the applications you build with it are **normal Java applications**. These applications can be packaged, deployed, and monitored like any other Java application – there is **no need to install separate processing clusters or similar special-purpose and expensive infrastructure!**

See <https://docs.confluent.io/platform/current/streams/introduction.html> (last accessed August 6, 2024).

A **KTable** is an abstraction of a **changelog stream**, where each data record represents an update. More precisely, the value in a data record is interpreted as an “UPDATE” of the last value for the same record key, if any (if a corresponding key doesn’t exist yet, the update will be considered an INSERT). Using the table

See <https://docs.confluent.io/platform/current/streams/concepts.html> (last accessed March 15, 2024).

Source Code Evidence

[REDACTED]

[REDACTED]

		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]

[illegible]

[REDACTED]

	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	

[REDACTED]

		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]

²The list is non exhaustive; Intellectual Ventures reserve the right to supplement the source code analysis.

[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

5. [REDACTED]
[REDACTED]

[illegible]

[Redacted]

9. [Redacted]
[Redacted]
[Redacted]

[Redacted]
[Redacted]

[Redacted]
[Redacted]

[Redacted]
[Redacted]

[Redacted]
[Redacted]

[Redacted]
[Redacted]

[Redacted]
[Redacted]
[Redacted]

[Redacted]
[Redacted]
[Redacted]
[Redacted]

[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]

████████████████████

████████████████████

████████████████████

██████████

[illegible]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

Age Group	Percentage of Respondents
18-29	90%
30-49	85%
50-64	75%
65+	65%

31. [REDACTED]
[REDACTED]
[REDACTED]

████████████████████

[REDACTED]

34. [REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]

████████████████████

[REDACTED]

	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	

[REDACTED]

48. [REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]
	[REDACTED]	[REDACTED]

© 2006 The Authors

© 2006 The Authors

[REDACTED]

[illegible]

62. [REDACTED]
[REDACTED]

- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]
- [REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[illegible]

[REDACTED]

	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	
	[REDACTED]	

[Redacted]

	[Redacted]	
	[Redacted]	
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	
	[Redacted]	
	[Redacted]	
	[Redacted]	
	[Redacted]	
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]
	[Redacted]	[Redacted]

[REDACTED]

	[REDACTED]	
	[REDACTED]	
	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
	[REDACTED]	[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

[REDACTED]

		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]
		[REDACTED]

[REDACTED]

	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]

	[REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]

[REDACTED]

95. [REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED]
[REDACTED]
[REDACTED]
[REDACTED]

[REDACTED] [REDACTED]
[REDACTED]

[REDACTED] [REDACTED]
[REDACTED]

[REDACTED] [REDACTED]
[REDACTED]
[REDACTED]

	[REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
	[REDACTED] [REDACTED]

[REDACTED]

	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]
	[REDACTED]

	[REDACTED]
[REDACTED]	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
[REDACTED]	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
[REDACTED]	[REDACTED] [REDACTED]
[REDACTED]	[REDACTED] [REDACTED]
	[REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]

14[c] wherein a gateway device at the routing network is configured to identify a category of the update message based on the input source, to determine a node type to which the identified category maps, and to route the update message to the node, having the node type, at the routing network,

Upon information and belief, Liberty Mutual’s systems include “wherein a gateway device at the routing network is configured to identify a category of the update message based on the input source, to determine a node type to which the identified category maps, and to route the update message to the node, having the node type, at the routing network.” Liberty Mutual’s systems include Kafka Clusters that comprise multiple Kafka brokers (e.g. nodes), with the initial broker or broker leader (e.g. gateway device) monitoring the activity of the other brokers. It handles new client connections, as well as events written and read by the event key (e.g. identify a category of the update message) to the broker responsible. A broker is selected based on the topic/partition to which the event belongs (e.g. determine a node type to which the identified category maps).




Figure 1-7. Replication of partitions in a cluster

A single Kafka server is called a *broker*. The broker receives messages from producers, assigns offsets to them, and writes the messages to storage on disk. It also services consumers, responding to fetch requests for partitions and responding with the messages that have been published. Depending on the specific hardware and its performance characteristics, a single broker can easily handle thousands of partitions and millions of messages per second.

Kafka brokers are designed to operate as part of a *cluster*. Within a cluster of brokers, one broker will also function as the cluster *controller* (elected automatically from the live members of the cluster). The controller is responsible for administrative operations, including assigning partitions to brokers and monitoring for broker failures. A partition is owned by a single broker in the cluster, and that broker is called the *leader* of the partition. A replicated partition (as seen in **Figure 1-7**) is assigned to additional brokers, called *followers* of the partition. Replication provides redundancy of messages in the partition, such that one of the followers can take over leadership if there is a broker failure. All producers must connect to the leader in order to publish messages, but consumers may fetch from either the leader or one of the followers. Cluster operations, including partition replication, are covered in detail in **Chapter 7**.

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

Servers: Kafka is run as a cluster of one or more servers that can span multiple datacenters or cloud regions. Some of these servers form the storage layer, called the brokers. Other servers run [Kafka Connect](#) to continuously import and export data as event streams to integrate Kafka with your existing systems such as relational databases as well as other Kafka clusters. To let you implement mission-critical use cases, a Kafka cluster is highly scalable and fault-tolerant: if any of its servers fails, the other servers will take over their work to ensure continuous operations without any data loss.

For example, records or events include dynamic objects, such as transactions, and make up respective topic streams that can be partitioned across multiple Kafka brokers. As they are written by a first entity/service/application, any new updates also be routed to the topic stream/partition it belongs to at respective Kafka brokers (e.g. route the update message to the node, have the node type).

An **event** records the fact that "something happened" in the world or in your business. It is also called record or message in the documentation. When you read or write data to Kafka, you do this in the form of events. Conceptually, an event has a key, value, timestamp, and optional metadata headers. Here's an example event:


- Event key: "Alice"
- Event value: "Made a payment of \$200 to Bob"
- Event timestamp: "Jun. 25, 2020 at 2:06 p.m."

Topics are **partitioned**, meaning a topic is spread over a number of "buckets" located on different Kafka brokers. This distributed placement of your data is very important for scalability because it allows client applications to both read and write the data from/to many brokers at the same time. When a new event is published to a topic, it is actually appended to one of the topic's partitions. Events with the same event key (e.g., a customer or vehicle ID) are written to the same partition, and Kafka guarantees that any consumer of a given topic-partition will always read that partition's events in exactly the same order as they were written.

See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).

A Kafka Cluster comprises of multiple Kafka brokers ("nodes"). A broker includes multiple topics (and partitions). A broker is selected based upon the topic/partition to which the event belongs. The broker stores

the events inside the partitions of the topics, where topics can be spread over several partitions to which a group of consumers can access.



A two server Kafka cluster hosting four partitions (P0-P3) with two consumer groups. Consumer group A has two consumer instances and group B has four.




Figure: This example topic has four partitions P1–P4. Two different producer clients are publishing, independently from each other, new events to the topic by writing events over the network to the topic's partitions. Events with the same key (denoted by their color in the figure) are written to the same partition. Note that both producers can write to the same partition if appropriate.

See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).

For example, when message writes are happening to a topic that is partitioned, it adds to the end of that specified partition to keep the order which these writes are happening. Furthermore, when a consumer group containing multiple consumer clients are reading from this topic, it will read from the offset/index it previously had.




Figure 1-5. Representation of a topic with multiple partitions




Figure 1-6. A consumer group reading from a topic





Figure 3-1. High-level overview of Kafka producer components


See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

	We start producing messages to Kafka by creating a <code>ProducerRecord</code>, which must include the topic we want to send the record to and a value. Optionally, we can also specify a key and/or a partition. Once we send the <code>ProducerRecord</code>, the first thing the producer will do is serialize the key and value objects to <code>ByteArrays</code> so they can be sent over the network. Next, the data is sent to a partitioner. If we specified a partition in the <code>ProducerRecord</code>, the partitioner doesn't do anything and simply returns the partition we specified. If we didn't, the partitioner will choose a partition for us, usually based on the <code>ProducerRecord</code> key. Once a partition is selected, the producer knows which topic and partition the record will go to. It then adds the record to a batch of records that will also be sent to the same topic and partition. A separate thread is responsible for sending those batches of records to the appropriate Kafka brokers. See https://www.confluent.io/resources/kafka-the-definitive-guide/ at p. 43 (last accessed March 11, 2024). Unable to load the picture See Kafka Architecture & Internal. This talks about the Apache Kafka... by Narayan Kumar Medium (last accessed on March 19, 2023). Unable to load the picture See https://www.cloudkarafka.com/blog/2016-11-30-part1-kafka-for-beginners-what-is-apache-kafka.html (last accessed March 15, 2024).
--	---



See <http://cloudurable.com/blog/kafka-architecture/index.html> (last accessed March 11, 2024).

Anatomy of a Topic



See <https://rafael-natali.medium.com/automatically-reassigning-partitions-in-apache-kafka-cluster-cd2d920fd5c9> (last accessed March 15, 2024).





Figure 3-1. High-level overview of Kafka producer components



See <https://www.cloudkarafka.com/blog/2016-11-30-part1-kafka-for-beginners-what-is-apache-kafka.html> (last accessed March 15, 2024).

As an example, (on information and belief) see the following sample showing configured to identify a category of the update message based on the input source, available at <https://kafka.apache.org/23/javadoc/index.html?org/apache/kafka/clients/producer/KafkaProducer.html> (last accessed March 15, 2024).

Method Detail

partition

```
int partition(String topic,
              Object key,
              byte[] keyBytes,
              Object value,
              byte[] valueBytes,
              Cluster cluster)
```

Compute the partition for the given record.

Parameters:

topic - The topic name

Id.

	org.apache.kafka.clients.producer Class ProducerRecord<K,V> java.lang.Object org.apache.kafka.clients.producer.ProducerRecord<K,V> <i>Id.</i> Constructor Summary Constructors Constructor and Description ProducerRecord(String topic, Integer partition, K key, V value) Creates a record to be sent to a specified topic and partition ProducerRecord(String topic, Integer partition, K key, V value, Iterable<Header> headers) Creates a record to be sent to a specified topic and partition ProducerRecord(String topic, Integer partition, Long timestamp, K key, V value) Creates a record with a specified timestamp to be sent to a specified topic and partition ProducerRecord(String topic, Integer partition, Long timestamp, K key, V value, Iterable<Header> headers) Creates a record with a specified timestamp to be sent to a specified topic and partition ProducerRecord(String topic, K key, V value) Create a record to be sent to Kafka ProducerRecord(String topic, V value) Create a record with no key <i>Id.</i>
--	--

14[d] wherein the node is configured to identify the client device as a registered device and to route the update message to the client device, and	Upon information and belief, Liberty Mutual’s systems include “wherein the node is configured to identify the client device as a registered device and to route the update message to the client device.” Consumer entity/service/application user devices (e.g. client device) subscribe (e.g. registered device) to a topic or multiple topics. The broker pushes the events of a topic (e.g. routes the update message) to the consumer (e.g., client device) in real-time that has subscribed to the topic. Producers are those client applications that publish (write) events to Kafka, and consumers are those that subscribe to (read and process) these events. In Kafka, producers and consumers are fully decoupled and agnostic of each other, which is a key design element to achieve the high scalability that Kafka is known for. For example, producers never need to wait for consumers. Kafka provides various guarantees such as the ability to process events exactly-once. Events are organized and durably stored in topics. Very simplified, a topic is similar to a folder in a filesystem, and the events are the files in that folder. An example topic name could be "payments". Topics in Kafka are always multi-producer and multi-subscriber: a topic can have zero, one, or many producers that write events to it, as well as zero, one, or many consumers that subscribe to these events. Events in a topic can be read as often as needed—unlike traditional messaging systems, events are not deleted after consumption. Instead, you define for how long Kafka should retain your events through a per-topic configuration setting, after which old events will be discarded. Kafka's performance is effectively constant with respect to data size, so storing data for a long time is perfectly fine. • The Producer API to publish (write) a stream of events to one or more Kafka topics. • The Consumer API to subscribe to (read) one or more topics and to process the stream of events produced to them. See https://kafka.apache.org/intro (last accessed on November 13, 2023).
--	---

2.4 Connect API

The Connect API allows implementing connectors that continually pull from some source data system into Kafka or push from Kafka into some sink data system.

Many users of Connect won't need to use this API directly, though, they can use pre-built connectors without needing to write any code. Additional information on using Connect is available [here](#).

See <https://kafka.apache.org/documentation/> (last accessed on November 10, 2023).

As an example (on information and belief), the following usage example provides a sample of the claim limitation. For all consumer instances reading from the same topic, they are grouped with distinctive IDs. An even distribution assignment of partitions of the topics will happen across each consumer group. The consumer API offers the functionality of offset committing to provide an accurate consumption of the messages in a topic stream.

Each consumer in a group can dynamically set the list of topics it wants to subscribe to through one of the `subscribe` APIs. Kafka will deliver each message in the subscribed topics to one process in each consumer group. This is achieved by balancing the partitions between all members in the consumer group so that each partition is assigned to exactly one consumer in the group. So if there is a topic with four partitions, and a consumer group with two processes, each process would consume from two partitions.

See <https://kafka.apache.org/10/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html>.

Usage Examples

The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them.

Automatic Offset Committing

This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("group.id", "test");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props);
consumer.subscribe(Arrays.asList("foo", "bar"));
while (true) {
    ConsumerRecords<String, String> records = consumer.poll(100);
    for (ConsumerRecord<String, String> record : records)
        System.out.printf("offset = %d, key = %s, value = %s\n", record.offset(), record.key(), record.value());
}
```

The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration `bootstrap.servers`. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting).


Setting `enable.auto.commit` means that offsets are committed automatically with a frequency controlled by the config `auto.commit.interval.ms`.

In this example the consumer is subscribing to the topics `foo` and `bar` as part of a group of consumers called `test` as configured with `group.id`.

Id.; see also <https://kafka.apache.org/20/documentation.html> (last accessed March 15, 2024) (detailing Consumer API).

- The [Consumer API](#) allows an application to subscribe to one or more topics and process the stream of records produced to them.

Id.



See <https://www.confluent.io/blog/a-comprehensive-rest-proxy-for-kafka/> (last accessed March 15, 2024).

Source Code Evidence

[illegible]

A horizontal bar chart with 25 categories listed on the left and corresponding percentage values represented by black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Yes, 4. No, 5. Yes, 6. No, 7. Yes, 8. No, 9. Yes, 10. No, 11. Yes, 12. No, 13. Yes, 14. No, 15. Yes, 16. No, 17. Yes, 18. No, 19. Yes, 20. No, 21. Yes, 22. No, 23. Yes, 24. No, 25. Yes. The percentages are: 1. 10%, 2. 15%, 3. 20%, 4. 25%, 5. 30%, 6. 35%, 7. 40%, 8. 35%, 9. 30%, 10. 25%, 11. 20%, 12. 15%, 13. 10%, 14. 15%, 15. 20%, 16. 25%, 17. 30%, 18. 35%, 19. 40%, 20. 35%, 21. 30%, 22. 25%, 23. 20%, 24. 15%, 25. 10%.

Category	Percentage
1. Other	10%
2. Not sure	15%
3. Yes	20%
4. No	25%
5. Yes	30%
6. No	35%
7. Yes	40%
8. No	35%
9. Yes	30%
10. No	25%
11. Yes	20%
12. No	15%
13. Yes	10%
14. No	15%
15. Yes	20%
16. No	25%
17. Yes	30%
18. No	35%
19. Yes	40%
20. No	35%
21. Yes	30%
22. No	25%
23. Yes	20%
24. No	15%
25. Yes	10%

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages shown as black bars on the right. The categories are: 1. (1%), 2. (35%), 3. (32%), 4. (30%), 5. (25%), 6. (28%), 7. (35%), 8. (32%), 9. (20%), 10. (28%), 11. (25%), 12. (20%), 13. (28%), 14. (25%), 15. (28%), 16. (25%), 17. (20%), 18. (25%), 19. (20%), 20. (25%), 21. (30%), 22. (32%), 23. (28%), 24. (25%), 25. (35%).

Category	Percentage
1.	1%
2.	35%
3.	32%
4.	30%
5.	25%
6.	28%
7.	35%
8.	32%
9.	20%
10.	28%
11.	25%
12.	20%
13.	28%
14.	25%
15.	28%
16.	25%
17.	20%
18.	25%
19.	20%
20.	25%
21.	30%
22.	32%
23.	28%
24.	25%
25.	35%

[illegible]

[illegible]

[illegible]

[illegible]

A horizontal bar chart with 30 categories listed on the left and their corresponding percentages represented by black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Very good, 4. Good, 5. Fair, 6. Poor, 7. Very poor, 8. Excellent, 9. Very good, 10. Good, 11. Fair, 12. Poor, 13. Very poor, 14. Excellent, 15. Very good, 16. Good, 17. Fair, 18. Poor, 19. Very poor, 20. Excellent, 21. Very good, 22. Good, 23. Fair, 24. Poor, 25. Very poor, 26. Excellent, 27. Very good, 28. Good, 29. Fair, 30. Poor, 31. Very poor, 32. Excellent, 33. Very good, 34. Good, 35. Fair, 36. Poor, 37. Very poor, 38. Excellent, 39. Very good, 40. Good, 41. Fair, 42. Poor, 43. Very poor, 44. Excellent, 45. Very good, 46. Good, 47. Fair, 48. Poor, 49. Very poor, 50. Excellent, 51. Very good, 52. Good, 53. Fair, 54. Poor, 55. Very poor, 56. Excellent, 57. Very good, 58. Good, 59. Fair, 60. Poor, 61. Very poor, 62. Excellent, 63. Very good, 64. Good, 65. Fair, 66. Poor, 67. Very poor, 68. Excellent, 69. Very good, 70. Good, 71. Fair, 72. Poor, 73. Very poor, 74. Excellent, 75. Very good, 76. Good, 77. Fair, 78. Poor, 79. Very poor, 80. Excellent, 81. Very good, 82. Good, 83. Fair, 84. Poor, 85. Very poor, 86. Excellent, 87. Very good, 88. Good, 89. Fair, 90. Poor, 91. Very poor, 92. Excellent, 93. Very good, 94. Good, 95. Fair, 96. Poor, 97. Very poor, 98. Excellent, 99. Very good, 100. Good, 101. Fair, 102. Poor, 103. Very poor, 104. Excellent, 105. Very good, 106. Good, 107. Fair, 108. Poor, 109. Very poor, 110. Excellent, 111. Very good, 112. Good, 113. Fair, 114. Poor, 115. Very poor, 116. Excellent, 117. Very good, 118. Good, 119. Fair, 120. Poor, 121. Very poor, 122. Excellent, 123. Very good, 124. Good, 125. Fair, 126. Poor, 127. Very poor, 128. Excellent, 129. Very good, 130. Good, 131. Fair, 132. Poor, 133. Very poor, 134. Excellent, 135. Very good, 136. Good, 137. Fair, 138. Poor, 139. Very poor, 140. Excellent, 141. Very good, 142. Good, 143. Fair, 144. Poor, 145. Very poor, 146. Excellent, 147. Very good, 148. Good, 149. Fair, 150. Poor, 151. Very poor, 152. Excellent, 153. Very good, 154. Good, 155. Fair, 156. Poor, 157. Very poor, 158. Excellent, 159. Very good, 160. Good, 161. Fair, 162. Poor, 163. Very poor, 164. Excellent, 165. Very good, 166. Good, 167. Fair, 168. Poor, 169. Very poor, 170. Excellent, 171. Very good, 172. Good, 173. Fair, 174. Poor, 175. Very poor, 176. Excellent, 177. Very good, 178. Good, 179. Fair, 180. Poor, 181. Very poor, 182. Excellent, 183. Very good, 184. Good, 185. Fair, 186. Poor, 187. Very poor, 188. Excellent, 189. Very good, 190. Good, 191. Fair, 192. Poor, 193. Very poor, 194. Excellent, 195. Very good, 196. Good, 197. Fair, 198. Poor, 199. Very poor, 200. Excellent, 201. Very good, 202. Good, 203. Fair, 204. Poor, 205. Very poor, 206. Excellent, 207. Very good, 208. Good, 209. Fair, 210. Poor, 211. Very poor, 212. Excellent, 213. Very good, 214. Good, 215. Fair, 216. Poor, 217. Very poor, 218. Excellent, 219. Very good, 220. Good, 221. Fair, 222. Poor, 223. Very poor, 224. Excellent, 225. Very good, 226. Good, 227. Fair, 228. Poor, 229. Very poor, 230. Excellent, 231. Very good, 232. Good, 233. Fair, 234. Poor, 235. Very poor, 236. Excellent, 237. Very good, 238. Good, 239. Fair, 240. Poor, 241. Very poor, 242. Excellent, 243. Very good, 244. Good, 245. Fair, 246. Poor, 247. Very poor, 248. Excellent, 249. Very good, 250. Good, 251. Fair, 252. Poor, 253. Very poor, 254. Excellent, 255. Very good, 256. Good, 257. Fair, 258. Poor, 259. Very poor, 260. Excellent, 261. Very good, 262. Good, 263. Fair, 264. Poor, 265. Very poor, 266. Excellent, 267. Very good, 268. Good, 269. Fair, 270. Poor, 271. Very poor, 272. Excellent, 273. Very good, 274. Good, 275. Fair, 276. Poor, 277. Very poor, 278. Excellent, 279. Very good, 280. Good, 281. Fair, 282. Poor, 283. Very poor, 284. Excellent, 285. Very good, 286. Good, 287. Fair, 288. Poor, 289. Very poor, 290. Excellent, 291. Very good, 292. Good, 293. Fair, 294. Poor, 295. Very poor, 296. Excellent, 297. Very good, 298. Good, 299. Fair, 300. Poor, 301. Very poor, 302. Excellent, 303. Very good, 304. Good, 305. Fair, 306. Poor, 307. Very poor, 308. Excellent, 309. Very good, 310. Good, 311. Fair, 312. Poor, 313. Very poor, 314. Excellent, 315. Very good, 316. Good, 317. Fair, 318. Poor, 319. Very poor, 320. Excellent, 321. Very good, 322. Good, 323. Fair, 324. Poor, 325. Very poor, 326. Excellent, 327. Very good, 328. Good, 329. Fair, 330. Poor, 331. Very poor, 332. Excellent, 333. Very good, 334. Good, 335. Fair, 336. Poor, 337. Very poor, 338. Excellent, 339. Very good, 340. Good, 341. Fair, 342. Poor, 343. Very poor, 344. Excellent, 345. Very good, 346. Good, 347. Fair, 348. Poor, 349. Very poor, 350. Excellent, 351. Very good, 352. Good, 353. Fair, 354. Poor, 355. Very poor, 356. Excellent, 357. Very good, 358. Good, 359. Fair, 360. Poor, 361. Very poor, 362. Excellent, 363. Very good, 364. Good, 365. Fair, 366. Poor, 367. Very poor, 368. Excellent, 369. Very good, 370. Good, 371. Fair, 372. Poor, 373. Very poor, 374. Excellent, 375. Very good, 376. Good, 377. Fair, 378. Poor, 379. Very poor, 380. Excellent, 381. Very good, 382. Good, 383. Fair, 384. Poor, 385. Very poor, 386. Excellent, 387. Very good, 388. Good, 389. Fair, 390. Poor, 391. Very poor, 392. Excellent, 393. Very good, 394. Good, 395. Fair, 396. Poor, 397. Very poor, 398. Excellent, 399. Very good, 400. Good, 401. Fair, 402. Poor, 403. Very poor, 404. Excellent, 405. Very good, 406. Good, 407. Fair, 408. Poor, 409. Very poor, 410. Excellent, 411. Very good, 412. Good, 413. Fair, 414. Poor, 415. Very poor, 416. Excellent, 417. Very good, 418. Good, 419. Fair, 420. Poor, 421. Very poor, 422. Excellent, 423. Very good, 424. Good, 425. Fair, 426. Poor, 427. Very poor, 428. Excellent, 429. Very good, 430. Good, 431. Fair, 432. Poor, 433. Very poor, 434. Excellent, 435. Very good, 436. Good, 437. Fair, 438. Poor, 439. Very poor, 440. Excellent, 441. Very good, 442. Good, 443. Fair, 444. Poor, 445. Very poor, 446. Excellent, 447. Very good, 448. Good, 449. Fair, 450. Poor, 451. Very poor, 452. Excellent, 453. Very good, 454. Good, 455. Fair, 456. Poor, 457. Very poor, 458. Excellent, 459. Very good, 460. Good, 461. Fair, 462. Poor, 463. Very poor, 464. Excellent, 465. Very good, 466. Good, 467. Fair, 468. Poor, 469. Very poor, 470. Excellent, 471. Very good, 472. Good, 473. Fair, 474. Poor, 475. Very poor, 476. Excellent, 477. Very good, 478. Good, 479. Fair, 480. Poor, 481. Very poor, 482. Excellent, 483. Very good, 484. Good, 485. Fair, 486. Poor, 487. Very poor, 488. Excellent, 489. Very good, 490. Good, 491. Fair, 492. Poor, 493. Very poor, 494. Excellent, 495. Very good, 496. Good, 497. Fair, 498. Poor, 499. Very poor, 500. Excellent, 501. Very good, 502. Good, 503. Fair, 504. Poor, 505. Very poor, 506. Excellent, 507. Very good, 508. Good, 509. Fair, 510. Poor, 511. Very poor, 512. Excellent, 513. Very good, 514. Good, 515. Fair, 516. Poor, 517. Very poor, 518. Excellent, 519. Very good, 520. Good, 521. Fair, 522. Poor, 523. Very poor, 524. Excellent, 525. Very good, 526. Good, 527. Fair, 528. Poor, 529. Very poor, 530. Excellent, 531. Very good, 532. Good, 533. Fair, 534. Poor, 535. Very poor, 536. Excellent, 537. Very good, 538. Good, 539. Fair, 540. Poor, 541. Very poor, 542. Excellent, 543. Very good, 544. Good, 545. Fair, 546. Poor, 547. Very poor, 548. Excellent, 549. Very good, 550. Good, 551. Fair, 552. Poor, 553. Very poor, 554. Excellent, 555. Very good, 556. Good, 557. Fair, 558. Poor, 559.

A horizontal bar chart with 25 rows. Each row consists of a small black dot followed by a black horizontal bar. The bars vary in length, representing different percentages. The longest bar is in the 19th row, and the shortest is in the 10th row.

Category	Percentage
1	10%
2	15%
3	12%
4	8%
5	14%
6	13%
7	16%
8	18%
9	11%
10	7%
11	14%
12	13%
13	17%
14	12%
15	15%
16	14%
17	13%
18	12%
19	20%
20	18%
21	16%
22	11%
23	10%
24	14%
25	13%

14[e] wherein the client device processes the update message upon receipt to update the property of the live object at the client device.

Upon information and belief, Liberty Mutual's systems include "wherein the client device processes the update message upon receipt to update the property of the live object at the client device." Upon receiving the events from one or more topics, the consumer entity/service/application processes the events (e.g., client device processes the update message). These consumed update messages are aggregated, enriched, or transformed into a new topic (e.g., update the property of the live object).

Kafka APIs

In addition to command line tooling for management and administration tasks, Kafka has five core APIs for Java and Scala:

- The [Admin API](#) to manage and inspect topics, brokers, and other Kafka objects.
- The [Producer API](#) to publish (write) a stream of events to one or more Kafka topics.
- The [Consumer API](#) to subscribe to (read) one or more topics and to process the stream of events produced to them.

Stream Processing

Many users of Kafka process data in processing pipelines consisting of multiple stages, where raw input data is consumed from Kafka topics and then aggregated, enriched, or otherwise transformed into new topics for further consumption or follow-up processing. For example, a processing pipeline for recommending news articles might crawl article content from RSS feeds and publish it to an "articles" topic; further processing might normalize or deduplicate this content and publish the cleansed article content to a new topic; a final processing stage might attempt to recommend this content to users. Such processing pipelines create graphs of real-time data flows based on the individual topics. Starting in 0.10.0.0, a light-weight but powerful stream processing library called [Kafka Streams](#) is available in Apache Kafka to perform such data processing as described above. Apart from Kafka Streams, alternative open source stream processing tools include [Apache Storm](#) and [Apache Samza](#).

See <https://kafka.apache.org/uses> (last accessed March 21, 2024).

As an example (on information and belief), the following usage example provides a sample of the claim limitation. For all consumer instances reading from the same topic, they are grouped with distinctive IDs. An even distribution assignment of partitions of the topics will happen across each consumer group. The consumer API offers the functionality of offset committing to provide an accurate consumption of the messages in a topic stream.

Each consumer in a group can dynamically set the list of topics it wants to subscribe to through one of the `subscribe` APIs. Kafka will deliver each message in the subscribed topics to one process in each consumer group. This is achieved by balancing the partitions between all members in the consumer group so that each partition is assigned to exactly one consumer in the group. So if there is a topic with four partitions, and a consumer group with two processes, each process would consume from two partitions.

Usage Examples

The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them.

Automatic Offset Committing

This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("group.id", "test");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props);
consumer.subscribe(Arrays.asList("foo", "bar"));
while (true) {
    ConsumerRecords<String, String> records = consumer.poll(100);
    for (ConsumerRecord<String, String> record : records)
        System.out.printf("offset = %d, key = %s, value = %s\n", record.offset(), record.key(), record.value());
}
```

The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration `>bootstrap.servers`. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting).

Setting `enable.auto.commit` means that offsets are committed automatically with a frequency controlled by the config `auto.commit.interval.ms`.

In this example the consumer is subscribing to the topics `foo` and `bar` as part of a group of consumers called `test` as configured with `group.id`.

See <https://kafka.apache.org/10/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html>.

Once we create a consumer, the next step is to subscribe to one or more topics. The `subscribe()` method takes a list of topics as a parameter, so it's pretty simple to use:

```
consumer.subscribe(Collections.singletonList("customerCountries")); ❶
```

❶ Here we simply create a list with a single element: the topic name `customerCountries`.

It is also possible to call `subscribe` with a regular expression. The expression can match multiple topic names, and if someone creates a new topic with a name that matches, a rebalance will happen almost immediately and the consumers will start consuming from the new topic. This is useful for applications that need to consume from multiple topics and can handle the different types of data the topics will contain. Subscribing to multiple topics using a regular expression is most commonly used in applications that replicate data between Kafka and another system or streams processing applications.

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

As another example, the below guidebook provides another example showing this limitation.

Chapter 4. Kafka Consumers: Reading Data from Kafka

Applications that need to read data from Kafka use a `KafkaConsumer` to subscribe to Kafka topics and receive messages from these topics. Reading data from Kafka is a bit different than reading data from other messaging systems, and there are few unique concepts and ideas involved. It is difficult to understand how to use the consumer API without understanding these concepts first. We'll start by explaining some of the important concepts, and then we'll go through some examples that show the different ways consumer APIs can be used to implement applications with varying requirements.

Commits and Offsets

Whenever we call `poll()`, it returns records written to Kafka that consumers in our group have not read yet. This means that we have a way of tracking which records were read by a consumer of the group. As discussed before, one of Kafka's unique characteristics is that it does not track acknowledgments from consumers the way many JMS queues do. Instead, it allows consumers to use Kafka to track their position (offset) in each partition.

We call the action of updating the current position in the partition a **commit**.

How does a consumer commit an offset? It produces a message to Kafka, to a special `__consumer_offsets` topic, with the committed offset for each partition. As long as all your consumers are up, running, and churning away, this will have no impact. However, if a consumer crashes or a new consumer joins the consumer group, this will *trigger a rebalance*. After a rebalance, each consumer may be assigned a new set of partitions than the one it processed before. In order to know where to pick up the work, the consumer will read the latest committed offset of each partition and continue from there.

If the committed offset is smaller than the offset of the last message the client processed, the messages between the last processed offset and the committed offset will be processed twice. See Figure 4-6.




Figure 4-6. Re-processed messages

If the committed offset is larger than the offset of the last message the client actually processed, all messages between the last processed offset and the committed offset will be missed by the consumer group. See Figure 4-7.

See <https://www.oreilly.com/library/view/kafka-the-definitive/9781491936153/ch04.html> (last accessed March 11, 2024).

Source Code Evidence

[REDACTED]

[REDACTED]

[REDACTED]

1	100%
2	95%
3	65%
4	85%
5	55%
6	80%
7	82%
8	58%
9	60%
10	62%
11	65%
12	60%
13	60%
14	60%
15	85%
16	50%
17	60%
18	65%
19	65%
20	85%
21	65%
22	68%
23	70%
24	60%
25	60%
26	65%
27	50%
28	55%
29	55%
30	60%
31	45%

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages shown as black bars on the right. The categories are: 1. Very much, 2. Somewhat, 3. Not much, 4. Not at all, 5. Don't know, 6. Very much, 7. Somewhat, 8. Not much, 9. Not at all, 10. Don't know, 11. Very much, 12. Somewhat, 13. Not much, 14. Not at all, 15. Don't know, 16. Very much, 17. Somewhat, 18. Not much, 19. Not at all, 20. Don't know, 21. Very much, 22. Somewhat, 23. Not much, 24. Not at all, 25. Don't know. The percentages are: 1. 1%, 2. 28%, 3. 27%, 4. 26%, 5. 25%, 6. 31%, 7. 24%, 8. 23%, 9. 22%, 10. 21%, 11. 20%, 12. 19%, 13. 18%, 14. 17%, 15. 16%, 16. 15%, 17. 14%, 18. 13%, 19. 12%, 20. 11%, 21. 10%, 22. 9%, 23. 8%, 24. 7%, 25. 6%.

Category	Percentage
1. Very much	1%
2. Somewhat	28%
3. Not much	27%
4. Not at all	26%
5. Don't know	25%
6. Very much	31%
7. Somewhat	24%
8. Not much	23%
9. Not at all	22%
10. Don't know	21%
11. Very much	20%
12. Somewhat	19%
13. Not much	18%
14. Not at all	17%
15. Don't know	16%
16. Very much	15%
17. Somewhat	14%
18. Not much	13%
19. Not at all	12%
20. Don't know	11%
21. Very much	10%
22. Somewhat	9%
23. Not much	8%
24. Not at all	7%
25. Don't know	6%

[illegible]

[illegible]

A horizontal bar chart consisting of 30 black bars of varying lengths. The bars are arranged in a single column, with the longest bar in the 10th position from the top and the shortest in the 1st position. The bars represent data points, with the longest bar reaching approximately 75% of the chart's width and the shortest bar reaching approximately 5%.

[illegible]

A horizontal bar chart with 25 categories listed on the left and their corresponding percentages shown as black bars on the right. The categories are: 1. Other, 2. Not sure, 3. Yes, 4. No, 5. Yes, 6. No, 7. Yes, 8. No, 9. Yes, 10. No, 11. Yes, 12. No, 13. Yes, 14. No, 15. Yes, 16. No, 17. Yes, 18. No, 19. Yes, 20. No, 21. Yes, 22. No, 23. Yes, 24. No, 25. Yes. The percentages are: 1. 1%, 2. 1%, 3. 31%, 4. 1%, 5. 1%, 6. 1%, 7. 1%, 8. 1%, 9. 1%, 10. 1%, 11. 1%, 12. 1%, 13. 1%, 14. 1%, 15. 1%, 16. 1%, 17. 1%, 18. 1%, 19. 1%, 20. 1%, 21. 1%, 22. 1%, 23. 1%, 24. 1%, 25. 1%.

Category	Percentage
1. Other	1%
2. Not sure	1%
3. Yes	31%
4. No	1%
5. Yes	1%
6. No	1%
7. Yes	1%
8. No	1%
9. Yes	1%
10. No	1%
11. Yes	1%
12. No	1%
13. Yes	1%
14. No	1%
15. Yes	1%
16. No	1%
17. Yes	1%
18. No	1%
19. Yes	1%
20. No	1%
21. Yes	1%
22. No	1%
23. Yes	1%
24. No	1%
25. Yes	1%

Category	Percentage
Category 1	~10%
Category 2	~95%
Category 3	~90%
Category 4	~85%
Category 5	~80%
Category 6	~75%
Category 7	~70%
Category 8	~65%
Category 9	~60%
Category 10	~55%
Category 11	~50%
Category 12	~45%
Category 13	~40%
Category 14	~35%
Category 15	~30%
Category 16	~25%
Category 17	~20%
Category 18	~15%
Category 19	~10%
Category 20	~5%

	• [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED] [REDACTED]
Claim 16	Liberty Mutual Count II Systems and Services



16. The method of claim 14, wherein providing the data representation to the client device includes providing the live object that causes the client device to register with the node of the routing network.	Upon information and belief, Liberty Mutual’s systems perform the “method of claim 14, wherein providing the data representation to the client device includes providing the live object that causes the client device to register with the node of the routing network.” On information and belief, Kafka facilitate the performance of this claim by Kafka clients performing DNS resolution. <i>See Claim 14, supra.</i> <i>See also</i> LM00004591. Kafka has four core APIs: • The Producer API allows an application to publish a stream of records to one or more Kafka topics. • The Consumer API allows an application to subscribe to one or more topics and process the stream of records produced to them. <i>See</i> https://kafka.apache.org/20/documentation.html (last accessed March 11, 2024).
---	--

	Usage Examples The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them. Automatic Offset Committing This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing. <pre>Properties props = new Properties(); props.put("bootstrap.servers", "localhost:9092"); props.put("group.id", "test"); props.put("enable.auto.commit", "true"); props.put("auto.commit.interval.ms", "1000"); props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer"); props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer"); KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props); consumer.subscribe(Arrays.asList("foo", "bar")); while (true) { ConsumerRecords<String, String> records = consumer.poll(100); for (ConsumerRecord<String, String> record : records) System.out.printf("offset = %d, key = %s, value = %s\n", record.offset(), record.key(), record.value()); }</pre> The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration <code>>bootstrap.servers</code>. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting). Setting <code>enable.auto.commit</code> means that offsets are committed automatically with a frequency controlled by the config <code>auto.commit.interval.ms</code>. In this example the consumer is subscribing to the topics <code>foo</code> and <code>bar</code> as part of a group of consumers called <code>test</code> as configured with <code>group.id</code>. See https://kafka.apache.org/10/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html (last accessed March 11, 2024). As an example (on information and belief), Kafka validates, resolves, and creates connections based on the hostname. Kafka clients perform DNS resolution to help establish the connection to a broker. To enable an easier connection configuration, a single DNS alias to address all brokers in a cluster may be created. By default, Kafka validates, resolves, and creates connections based on the hostname provided in the bootstrap server configuration (and later in the names returned by the brokers as specified in the <code>advertised.listeners</code> configuration). This simple model works most of the time but fails to cover two important use cases: the use of DNS aliases, especially in a bootstrap configuration, and the use of a single DNS that maps to multiple IP addresses. These sound similar but are slightly different. Let's look at each of these mutually exclusive scenarios in a bit more detail.
--	---

Suppose you have multiple brokers with the following naming convention: `broker1.hostname.com`, `broker2.hostname.com`, etc. Rather than specifying all of them in a bootstrap server configuration, which can easily become challenging to maintain, you may want to create a single DNS alias that will map to all of them.

You'll use `all-brokers.hostname.com` for bootstrapping, since you don't actually care which broker gets the initial connection from clients. This is all very convenient, except if you use SASL to authenticate. If you use SASL, the client will try to authenticate `all-brokers.hostname.com`, but the server principal will be `broker2.hostname.com`. If the names don't match, SASL will refuse to authenticate (the broker certificate could be a man-in-the-middle attack), and the connection will fail.

In this scenario, you'll want to use `client.dns.lookup=resolve_canonical_bootstrap_servers_only`. With this configuration, the client will "expend" the DNS alias, and the result will be the same as if you included all the broker names the DNS alias connects to as brokers in the original bootstrap list.

See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. *Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale*. O'Reilly Media, Inc.

As a further example, (on information and belief) Kafka clients perform DNS resolution to help establish connection to a broker in the Kafka cluster.

DNS Resolution for Kafka clients



	DNS resolution is an important aspect of Kafka client connectivity as it enables the clients to connect to the correct Kafka brokers. When a Kafka client is configured to connect to a Kafka cluster, it typically specifies a list of broker hostnames or IP addresses in its configuration. When a Kafka client is started, it first reads its configuration file to determine the initial list of bootstrap brokers. These brokers are specified as a comma-separated list of hostname and port pairs (e.g., <code>broker1.example.com:9092,broker2.example.com:9092</code>). The Kafka client then performs DNS resolution to obtain the IP addresses of the specified hostnames. By default, the client uses the operating system's default DNS resolution behavior, which typically involves sending DNS queries to a DNS resolver provided by the operating system. See https://levelup.gitconnected.com/dns-resolution-for-kafka-clients-2e76b28e4717 (last accessed March 11, 2024). Source Code Evidence See Claim 14, <i>supra</i>.
Claim 17	Liberty Mutual Count II Systems and Services

17. The method of claim 14, wherein providing the data representation to the client device includes providing an activation module that is executed by the client device and that registers the live object with the routing network.	Upon information and belief, Liberty Mutual’s systems perform the “method of claim 14, wherein providing the data representation to the client device includes providing an activation module that is executed by the client device and that registers the live object with the routing network.” On information and belief, Kafka facilitates the performance of this claim by handling validation, resolving, and creating of connections. <i>See Claim 14, supra.</i> As an example, the Consumer API allows an application to process the stream of records produced to them. https://kafka.apache.org/20/documentation.html (last accessed March 20, 2024). Kafka has four core APIs: • The Producer API allows an application to publish a stream of records to one or more Kafka topics. • The Consumer API allows an application to subscribe to one or more topics and process the stream of records produced to them. <i>Id.</i> As a further example (on information and belief) the below usage example shows how this limitation is met:
---	--

Usage Examples

The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them.

Automatic Offset Committing

This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("group.id", "test");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props);
consumer.subscribe(Arrays.asList("foo", "bar"));
while (true) {
    ConsumerRecords<String, String> records = consumer.poll(100);
    for (ConsumerRecord<String, String> record : records)
        System.out.printf("offset = %d, key = %s, value = %s\n", record.offset(), record.key(), record.value());
}
```

See <https://kafka.apache.org/24/javadoc/index.html?org/apache/kafka/clients/consumer/KafkaConsumer.html#:~:text=Consumer%20Groups%20and%20Topic%20Subscriptions&text=Each%20consumer%20in%20a%20group,process%20in%20each%20consumer%20group> (last accessed March 11, 2024).

The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration `>bootstrap.servers`. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting).

Setting `enable.auto.commit` means that offsets are committed automatically with a frequency controlled by the config `auto.commit.interval.ms`.

In this example the consumer is subscribing to the topics `foo` and `bar` as part of a group of consumers called `test` as configured with `group.id`.

Id.

See also <https://www.conduktor.io/kafka/complete-kafka-consumer-with-java/> (last accessed March 15, 2024) (describing an example using the feature that subscribes the consumer to a specific topic):

	Subscribe the consumer to a specific topic To read the messages from a topic, we need to connect the consumer to the specified topic. Here, we use <code>Arrays.asList()</code>, as it allows our consumer to subscribe to multiple topics. Below code shows the implementation of subscription of the consumer to one topic: <pre>1 // subscribe consumer to our topic(s) 2 consumer.subscribe(Arrays.asList(topic));</pre> <i>Id.</i> As a further example (on information and belief), Kafka destination entities/services/applications (e.g. client device) initializing a consumer instance subscribing to topics are also noted as subscribers.
--	--



Once we create a consumer, the next step is to subscribe to one or more topics. The `subscribe()` method takes a list of topics as a parameter, so it's pretty simple to use:

```
consumer.subscribe(Collections.singletonList("customerCountries")); ❶
```

- ❶ Here we simply create a list with a single element: the topic name `customerCountries`.

It is also possible to call `subscribe` with a regular expression. The expression can match multiple topic names, and if someone creates a new topic with a name that matches, a rebalance will happen almost immediately and the consumers will start consuming from the new topic. This is useful for applications that need to consume from multiple topics and can handle the different types of data the topics will contain. Subscribing to multiple topics using a regular expression is most commonly used in applications that replicate data between Kafka and another system or streams processing applications.

For example, Kafka validates, resolves, and creates connections based on the hostname. Kafka clients perform DNS resolution to help establish the connection to a broker. To enable an easier connection configuration, a single DNS alias to address all brokers in a cluster may be created.

By default, Kafka validates, resolves, and creates connections based on the hostname provided in the bootstrap server configuration (and later in the names returned by the brokers as specified in the `advertised.listeners` configuration). This simple model works most of the time but fails to cover two important use cases: the use of DNS aliases, especially in a bootstrap configuration, and the use of a single DNS that maps to multiple IP addresses. These sound similar but are slightly different. Let's look at each of these mutually exclusive scenarios in a bit more detail.

Suppose you have multiple brokers with the following naming convention: `broker1.hostname.com`, `broker2.hostname.com`, etc. Rather than specifying all of them in a bootstrap server configuration, which can easily become challenging to maintain, you may want to create a single DNS alias that will map to all of them.



	You'll use <code>all-brokers.hostname.com</code> for bootstrapping, since you don't actually care which broker gets the initial connection from clients. This is all very convenient, except if you use SASL to authenticate. If you use SASL, the client will try to authenticate <code>all-brokers.hostname.com</code>, but the server principal will be <code>broker2.hostname.com</code>. If the names don't match, SASL will refuse to authenticate (the broker certificate could be a man-in-the-middle attack), and the connection will fail. In this scenario, you'll want to use <code>client.dns.lookup=resolve_canonical_bootstrap_servers_only</code>. With this configuration, the client will "expend" the DNS alias, and the result will be the same as if you included all the broker names the DNS alias connects to as brokers in the original bootstrap list. See Neha Narkhede, Gwen Shapira, Todd Palino. 2017. <i>Kafka: The Definitive Guide: Real-Time Data and Stream Processing at Scale</i>. O'Reilly Media, Inc. Source Code Evidence See Claim 14, <i>supra</i>.
Claim 18	Liberty Mutual Count II Systems and Services

18. The method of claim 17, wherein providing the activation module includes providing the activation module that is configured to determine a node type that handles registration and that causes the client device to register with the node corresponding to the node type.	Upon information and belief, Liberty Mutual’s systems perform the “method of claim 17, wherein providing the activation module includes providing the activation module that is configured to determine a node type that handles registration and that causes the client device to register with the node corresponding to the node type.” On information and belief, Kafka facilitates the performance of this claim by handling validation, resolving, and creating of connections. <i>See Claim 17, supra.</i> As an example (on information and belief), Kafka provides offsets including automatic offsets, which determine a node type that handles registration and that causes the client device to register with the node corresponding to the node type. We start producing messages to Kafka by creating a <code>ProducerRecord</code>, which must include the topic we want to send the record to and a value. Optionally, we can also specify a key and/or a partition. Once we send the <code>ProducerRecord</code>, the first thing the producer will do is serialize the key and value objects to <code>ByteArrays</code> so they can be sent over the network. Next, the data is sent to a partitioner. If we specified a partition in the <code>ProducerRecord</code>, the partitioner doesn’t do anything and simply returns the partition we specified. If we didn’t, the partitioner will choose a partition for us, usually based on the <code>ProducerRecord</code> key. Once a partition is selected, the producer knows which topic and partition the record will go to. It then adds the record to a batch of records that will also be sent to the same topic and partition. A separate thread is responsible for sending those batches of records to the appropriate Kafka brokers. <i>See https://www.confluent.io/resources/kafka-the-definitive-guide/ (last accessed March 11, 2024)</i>
---	---

Usage Examples

The consumer APIs offer flexibility to cover a variety of consumption use cases. Here are some examples to demonstrate how to use them.

Automatic Offset Committing

This example demonstrates a simple usage of Kafka's consumer api that relying on automatic offset committing.

```
Properties props = new Properties();
props.put("bootstrap.servers", "localhost:9092");
props.put("group.id", "test");
props.put("enable.auto.commit", "true");
props.put("auto.commit.interval.ms", "1000");
props.put("key.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
props.put("value.deserializer", "org.apache.kafka.common.serialization.StringDeserializer");
KafkaConsumer<String, String> consumer = new KafkaConsumer<>(props);
consumer.subscribe(Arrays.asList("foo", "bar"));
while (true) {
    ConsumerRecords<String, String> records = consumer.poll(100);
    for (ConsumerRecord<String, String> record : records)
        System.out.printf("offset = %d, key = %s, value = %s%n", record.offset(), record.key(), record.value());
}
```

See <https://kafka.apache.org/24/javadoc/index.html?org/apache/kafka/clients/consumer/KafkaConsumer.html>
(last accessed March 11, 2024)

The connection to the cluster is bootstrapped by specifying a list of one or more brokers to contact using the configuration >bootstrap.servers. This list is just used to discover the rest of the brokers in the cluster and need not be an exhaustive list of servers in the cluster (though you may want to specify more than one in case there are servers down when the client is connecting).

Setting enable.auto.commit means that offsets are committed automatically with a frequency controlled by the config auto.commit.interval.ms.

In this example the consumer is subscribing to the topics *foo* and *bar* as part of a group of consumers called *test* as configured with group.id.

Id.

As a further example (on information and belief), Kafka destination entities/services/applications (e.g. client device) initializing a consumer instance subscribing to topics are also noted as subscribers. The KafkaConsumer class with the consumer API maintains an offset committing to provide an accurate consumption of the messages in a topic stream.

	org.apache.kafka.clients.consumer Class KafkaConsumer<K,V> java.lang.Object org.apache.kafka.clients.consumer.KafkaConsumer<K,V> See https://kafka.apache.org/22/javadoc/org/apache/kafka/clients/consumer/KafkaConsumer.html (last accessed March 11, 2024) <i>Offsets and Consumer Position</i> Kafka maintains a numerical offset for each record in a partition. This offset acts as a unique identifier of a record within that partition, and also denotes the position of the consumer in the partition. For example, a consumer which is at position 5 has consumed records with offsets 0 through 4 and will next receive the record with offset 5. There are actually two notions of position relevant to the user of the consumer: <i>Id.</i> Source Code Evidence See Claim 14, <i>supra</i>.
--	--