

President George Bush Hwy, Plano Texas 75075; 900 Venture Drive, Allen Texas 75013; and 8537 Labelle Road, Beaumont Texas, 77705. Comcast Cable Communications, LLC is registered to conduct business in the state of Texas and has appointed C T Corporation System, located at 1999 Bryan ST., Ste. 900, Dallas, TX 75201 as its agent for service of process.

3. Defendant Comcast Cable Communications Management, LLC is a limited liability company organized and existing under the laws of the State of Delaware that maintains regular and established places of business throughout Texas, for example, at its facilities in this District, such as 135 Houston St., Lewisville Texas, 75057; 1300 Coit Road, Plano Texas 75075; 3033 W. President George Bush Hwy, Plano Texas 75075; 900 Venture Drive, Allen Texas 75013; and 8537 Labelle Road, Beaumont Texas, 77705. Comcast Cable Communications Management, LLC is registered to conduct business in the state of Texas and has appointed Corporation Service Company d/b/a CSC-Lawyers Incorporating Service Company 211 E. 7th Street, Suite 620, Austin, TX 78701 USA as its agent for service of process.

4. Defendant is a multinational information technology company and develops and sells networking equipment and phone services. Defendant sells its products to customers, including customers in this District.

5. Defendant operates and owns the xfinity.com and comcast.com websites, and it markets, offers, distributes, and provides technical support for its networking equipment and phone services throughout the United States including in this District.

6. Defendant develops, designs, manufactures, distributes, markets, offers to sell, and/or sells infringing products and services within the United States, including in this District, and otherwise purposefully directs infringing activities to this District in connection with its aforementioned Texas offices; its aforementioned websites; and its other places of business in

Texas and the rest of the United States. Defendant participates in the design, development, manufacture, sale for importation into the United States, offers for sale for importation into the United States, importation into the United States, sale within the United States after importation, and offers for sale within the United States after importation, of networking equipment and phone services that infringe the Asserted Patents.

7. On information and belief, Defendant is engaged in making, using, selling, offering for sale, and/or importing, and/or inducing its subsidiaries, affiliates, retail partners, and customers in the making, using, selling, offering for sale, and/or importing throughout the United States, including within this District, the products, such as networking equipment, and services, such as phone services, accused of infringement.

8. Prior to the filing of the Complaint, Plaintiff attempted to engage Defendant and/or its agents in good faith licensing discussions related to the Asserted Patents, including by sending them correspondence on September 13, 2024 notifying Defendant of the need to license the Asserted Patents. Defendant's past and continuing sales of its devices i) willfully infringe the Asserted Patents and ii) impermissibly take the significant benefits of Plaintiff's patented technologies without fair compensation to Plaintiff.

9. Through offers to sell, sales, imports, distributions, and other related agreements to transfer ownership of Defendant's electronics, such as networking equipment, and/or Defendant's services, such as phone services, with distributors and customers operating in and maintaining a significant business presence in the U.S. and/or its U.S. subsidiaries Defendant does business in the U.S., the state of Texas, and in this District.

JURISDICTION AND VENUE

10. This action arises under the patent laws of the United States, namely 35 U.S.C. §§ 271, 281, and 284-285, among others.

11. This Court has subject matter jurisdiction pursuant to 28 U.S.C. §§ 1331 and 1338(a).

12. This Court has personal jurisdiction over Defendant in accordance with due process and/or the Texas Long Arm Statute because, in part, Defendant “recruits Texas residents, directly or through an intermediary located in this state, for employment inside or outside this state.” TEX. CIV. PRAC. & REM. CODE § 17.042(3).

13. This Court has personal jurisdiction over Defendant because Defendant has engaged, and continues to engage in continuous, systematic, and substantial activities within this State, including the substantial marketing and sale of products within this State and this District. Furthermore, upon information and belief, this Court has personal jurisdiction over Defendant because Defendant has committed acts giving rise to Plaintiff’s claims for patent infringement within and directed to this District.

14. For example, Defendant is subject to personal jurisdiction in this Court because, *inter alia*, it has regular and established places of business in this District, including offices and data centers located at 135 Houston St., Lewisville Texas, 75057; 1300 Coit Road, Plano Texas 75075; 3033 W. President George Bush Hwy, Plano Texas 75075; 900 Venture Drive, Allen Texas 75013; and 8537 Labelle Road, Beaumont Texas, 77705.

15. Defendant’s offices in the District are regular and established places of business at least because these locations include many members of Defendant’s important teams, including

engineers and sales representatives. Defendant's employees in the District are highly specialized and are important to the operation of Defendant.

16. Defendant, directly and through its agents, regularly conducts, solicits, and transacts business in this District and elsewhere in Texas, including through its xfinity.com and comcast.com websites. For example, Defendant employs sales and marketing employees that regularly sell, offer to sell, or otherwise distribute networking equipment in this District and elsewhere in Texas.

17. Defendant has committed and continues to commit acts of infringement in violation of 35 U.S.C. § 271, and has made, used, marketed, distributed, offered for sale, and sold infringing products in Texas, including in this District, and engaged in infringing conduct within and directed at or from this District. The infringing networking equipment have been and continue to be distributed to and used in this District. Defendant's acts cause injury to Plaintiff, including injury suffered within this District.

18. Moreover, on information and belief, Defendant has previously litigated patent infringement cases before this Court without contesting jurisdiction and venue.

19. Exercising personal jurisdiction over Defendant in this District would not be unreasonable given Defendant's contacts in this District, the interest in this District of resolving disputes related to products sold herein.

20. In addition, Defendant has knowingly induced and continues to knowingly induce infringement within this District by advertising, marketing, offering for sale and/or selling devices pre-loaded with infringing functionality within this District, to consumers, customers, manufacturers, distributors, resellers, partners, and/or end users, and providing instructions, user

manuals, advertising, and/or marketing materials which facilitate, direct or encourage the use of infringing functionality with knowledge thereof.

21. Personal jurisdiction also exists specifically over Defendant because it, directly or through affiliates, subsidiaries, agents, or intermediaries, transacts business in this State or purposefully directed at this State (including, without limitation, retail stores including Best Buy) by making, importing, offering to sell, selling, and/or having sold infringing products within this State and District or purposefully directed at this State or District.

22. Venue is proper in this District under 28 U.S.C. §§ 1391 and 1400(b) because a substantial part of the events or omissions giving rise to the claims occurred in this District, and because Defendant has committed acts of infringement in this District and have a regular and established place of business in this District.

23. On information and belief, Defendant has placed and continues to place infringing products and/or products that practice infringing processes into the stream of commerce via established distribution channels, with the knowledge and/or intent that those products are and/or will be imported, used, offered for sale, sold, and continue to be sold in the United States and Texas, including in this judicial district. As a result, Defendant has, vicariously through and/or in concert with its alter egos, agents, intermediaries, distributors, importers, customers, subsidiaries, and/or consumers, placed the Accused Products (identified in Counts I – IV) into the stream of commerce via established distribution channels with the knowledge and/or intent that those products were sold and continue to be sold in the United States and Texas, including in this judicial district.

COUNT I

(INFRINGEMENT OF U.S. PATENT NO. 7,366,110)

24. Plaintiff incorporates the preceding paragraphs herein by reference.

25. Plaintiff is the assignee of the '110 patent, entitled "Method and apparatus for merging call components during call reconstruction," with ownership of all substantial rights in the '110 patent, including the right to exclude others and to enforce, sue, and recover damages for past and future infringements.

26. The '110 patent is valid, enforceable, and was duly issued in full compliance with Title 35 of the United States Code. The '110 patent issued from U.S. Patent Application No. 11/045,702.

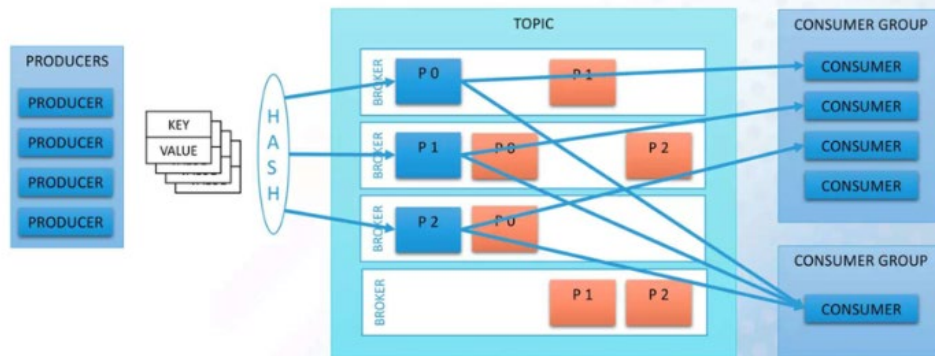
27. Defendant has and continues to directly and/or indirectly infringe (by inducing infringement) one or more claims of the '110 patent in this judicial district and elsewhere in Texas and the United States.

28. Defendant designs, offers for sale, uses, and sells services, such as its broadband infrastructure and ISP services that use Apache Kafka ("the '110 Accused Products"), in a manner that infringes the '110 patent. For example, Defendant uses Apache Kafka to monitor its network infrastructure:

INSIDE KAFKA STREAMS - MONITORING COMCAST'S OUTSIDE PLANT

KAFKA CONSUMERS & PRODUCERS

PARTITION MECHANICS



29. Defendant directly infringes the '110 patent under 35 U.S.C. § 271(a) by using, making, offering for sale, selling, and/or importing the '110 Accused Products, their components and processes, and/or products containing the same that incorporate the fundamental technologies covered by the '110 patent.

30. For example, Defendant infringes claim 1 of the '110 patent via the '110 Accused Products. The '110 Accused Products are resilient to failure of a Kafka broker because they perform a method for migrating Topic partitions from a first broker to a second broker:

4.7 Replication

Kafka replicates the log for each topic's partitions across a configurable number of servers (you can set this replication factor on a topic-by-topic basis). This allows automatic failover to these replicas when a server in the cluster fails so messages remain available in the presence of failures.

Other messaging systems provide some replication-related features, but, in our (totally biased) opinion, this appears to be a tacked-on thing, not heavily used, and with large downsides: replicas are inactive, throughput is heavily impacted, it requires fiddly manual configuration, etc. Kafka is meant to be used with replication by default—in fact we implement un-replicated topics as replicated topics where the replication factor is one.

Source: <https://kafka.apache.org/documentation/#replication>.

31. The '110 Accused Products determine “that at least one communication is to be controlled by a second communication server, wherein the at least one communication was formerly controlled by a first communication server.” For example, when a leader broker fails, the Kafka cluster needs to elect a new leader for the partition that was controlled by the leader broker:

Of course if leaders didn't fail we wouldn't need followers! When the leader does die we need to choose a new leader from among the followers. But followers themselves may fall behind or crash so we must ensure we choose an up-to-date follower. The fundamental guarantee a log replication algorithm must provide is that if we tell the client a message is committed, and the leader fails, the new leader we elect must also have that message. This yields a tradeoff: if the leader waits for more followers to acknowledge a message before declaring it committed then there will be more potentially electable leaders.

Source: <https://kafka.apache.org/documentation/#replication>.

32. The '110 Accused Products receive “from a first communication node, first communication information, wherein the first communication information is associated with the at least one communication and comprises at least one of a first node identifier and a communication identifier, the communication identifier is associated with the at least one

communication, the second communication node comprises second communication information associated with the at least one communication and/or second node, and the first node identifier is associated with second communication information.” For example, the Kafka cluster receives a first message from a first producer. The message includes a topic, a key, and a value. The key is a communication identifier:

Main Concepts and Terminology

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."

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.

Source: https://kafka.apache.org/documentation/#intro_concepts_and_terms. Further, the key included in the message record is associated with a topic because the key determines which partition of a topic receives the message:

send

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

Asynchronously send a record to a topic and invoke the provided callback when the send has been acknowledged.

The send is asynchronous and this method will return immediately once the record has been stored in the buffer of records waiting to be sent. This allows sending many records in parallel without blocking to wait for the response after each one.

The result of the send is a `RecordMetadata` specifying the partition the record was sent to, the offset it was assigned and the timestamp of the record. If the producer is configured with `acks = 0`, the `RecordMetadata` will have `offset = -1` because the producer does not wait for the acknowledgement from the broker. If `CreateTime` is used by the topic, the timestamp will be the user provided timestamp or the record send time if the user did not specify a timestamp for the record. If `LogAppendTime` is used for the topic, the timestamp will be the Kafka broker local time when the message is appended.

Since the send call is asynchronous it returns a `Future` for the `RecordMetadata` that will be assigned to this record. Invoking `get()` on this future will block until the associated request completes and then return the metadata for the record or throw any exception that occurred while sending the record.

If you want to simulate a simple blocking call you can call the `get()` method immediately:

```
byte[] key = "key".getBytes();
byte[] value = "value".getBytes();
ProducerRecord<byte[],byte[]> record = new ProducerRecord<byte[],byte[]>("my-topic", key, value)
producer.send(record).get();
```

Source:

<https://kafka.apache.org/36/javadoc/org/apache/kafka/clients/producer/KafkaProducer.html>.

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.

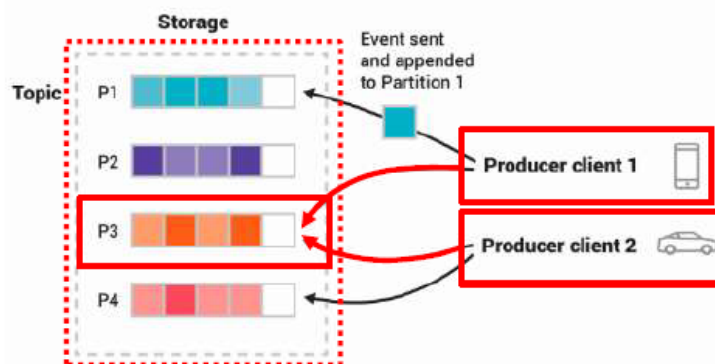


Figure: This example topic has four partitions P1–P4. Two different producer clients are publishing, independently from each other, new events to the topic by writing events over the network to the topic's partitions. Events with the same key (denoted by their color in the figure) are written to the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

Source: https://kafka.apache.org/documentation/#intro_concepts_and_terms. A second producer sends a second message to the Kafka cluster to be written to the same topic. The producer writes the second message, which comprises topic, key, and value information. As explained, the key included in the message record is associated with a topic because the key determines which partition of a topic receives the message. Because the key is associated with the topic, the key need not be associated with the second producer.

33. The '110 Accused Products thereafter receive "from a second communication node, the second communication information." For example, the Kafka cluster receives a second message from a second producer. The message includes a topic, a key, and a value:

Main Concepts and Terminology

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."

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.

Source: https://kafka.apache.org/documentation/#intro_concepts_and_terms. The Kafka cluster receives a second message from a second producer. The message includes a topic, a key, and a value:

send

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

Asynchronously send a record to a topic and invoke the provided callback when the send has been acknowledged.

The send is asynchronous and this method will return immediately once the record has been stored in the buffer of records waiting to be sent. This allows sending many records in parallel without blocking to wait for the response after each one.

The result of the send is a `RecordMetadata` specifying the partition the record was sent to, the offset it was assigned and the timestamp of the record. If the producer is configured with `acks = 0`, the `RecordMetadata` will have `offset = -1` because the producer does not wait for the acknowledgement from the broker. If `CreateTime` is used by the topic, the timestamp will be the user provided timestamp or the record send time if the user did not specify a timestamp for the record. If `LogAppendTime` is used for the topic, the timestamp will be the Kafka broker local time when the message is appended.

Since the send call is asynchronous it returns a `Future` for the `RecordMetadata` that will be assigned to this record. Invoking `get()` on this future will block until the associated request completes and then return the metadata for the record or throw any exception that occurred while sending the record.

If you want to simulate a simple blocking call you can call the `get()` method immediately:

```
byte[] key = "key".getBytes();
byte[] value = "value".getBytes();
ProducerRecord<byte[],byte[]> record = new ProducerRecord<byte[],byte[]>("my-topic", key, value)
producer.send(record).get();
```

Source:

<https://kafka.apache.org/36/javadoc/org/apache/kafka/clients/producer/KafkaProducer.html>.

34. The '110 Accused Products identify "the second communication information based on the at least one of a first node identifier and communication identifier." For example, the second message is identified by the Kafka cluster based on the key. The key determines which topic partition the second message is appended to:

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.

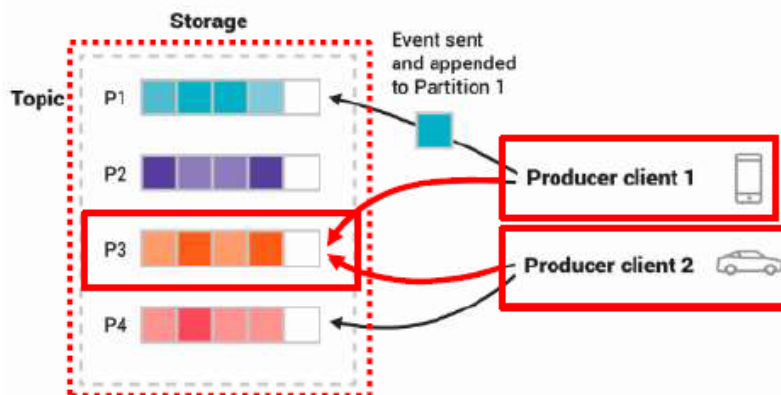


Figure: This example topic has four partitions P1–P4. Two different producer clients are publishing, independently from each other, new events to the topic by writing events over the network to the topic's partitions. Events with the same key (denoted by their color in the figure) are written to the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

Source: https://kafka.apache.org/documentation/#intro_concepts_and_terms.

35. The technology discussion above and the exemplary '110 Accused Products provide context for Plaintiff's infringement allegations.

36. At a minimum, Defendant has known of the '110 patent since at least the service of this complaint on September 26, 2024. Further, Defendant has known of the '110 patent since at least the filing date of the complaint on September 20, 2024. In addition, Defendant has known about the '110 patent since at least September 16, 2024, when it received email correspondence from Plaintiff alerting Defendant to its infringement. Moreover, Defendant has known about the

'110 patent since at least September 13, 2024, when Plaintiff sent correspondence to Defendant via FedEx alerting Defendant to its infringement.

37. On information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has actively induced, under U.S.C. § 271(b), its distributors, customers, subsidiaries, importers, and/or consumers that import, purchase, or sell the '110 Accused Products that include or are made using all of the limitations of one or more claims of the '110 patent to directly infringe one or more claims of the '110 patent (e.g., claim 1, as discussed above) by using, offering for sale, selling, and/or importing the '110 Accused Products. Since at least the notice provided on the above-mentioned date, Defendant does so with knowledge, or with willful blindness of the fact, that the induced acts constitute infringement of the '110 patent. Defendant intends to cause, and has taken affirmative steps to induce infringement by its distributors, importers, customers, subsidiaries, and/or consumers by at least, inter alia, creating advertisements that promote the infringing use of the '110 Accused Products, creating and/or maintaining established distribution channels for the '110 Accused Products into and within the United States, manufacturing the '110 Accused Products in conformity with U.S. laws and regulations, distributing or making available instructions or manuals for these products to purchasers and prospective buyers, and testing the '110 Accused Products, and/or providing technical support, replacement parts, or services for these products to these purchasers in the United States. For example, Defendant configures Apache Kafka in the '110 Accused Products to contain specific instructions, such as source code and configuration files, that cause such products to monitor its network infrastructure. Moreover, in addition to the foregoing, Defendant encourages application developers to build applications in a manner that uses the '110 Accused Products to automatically restart applications on both Defendant's private data centers and public

cloud providers in a manner that necessarily infringes the '110 patent. *See, e.g.,* <https://www.youtube.com/watch?v=oVvoqg-NiKU> (a transcript of which is attached as Exhibit A).

38. In the alternative, on information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has contributorily infringed, under U.S.C. § 271(c), one or more claims of the '110 patent. For example, Defendant contributes to the direct infringement of such claims by distributors, customers, subsidiaries, importers, and/or consumers that use, import, purchase, or sell the '110 Accused Products. To the extent that the '110 Accused Products do not directly infringe one or more claims of the '110 patent, such products contain instructions, such as source code, that are especially adapted to cause the '110 Accused Products to operate in an infringing manner. Such instructions are specifically designed to cause the '110 Accused Products to provide and utilize Apache Kafka in an infringing manner and are a material part of the invention of the '110 patent and are not a staple article of commerce suitable for substantial non-infringing use.

39. On information and belief, despite having knowledge of the '110 patent and knowledge that it is directly and/or indirectly infringing one or more claims of the '110 patent, Defendant has nevertheless continued its infringing conduct and disregarded an objectively high likelihood of infringement. Defendant's infringing activities relative to the '110 patent have been, and continue to be, willful, wanton, malicious, in bad-faith, deliberate, consciously wrongful, flagrant, characteristic of a pirate, and an egregious case of misconduct beyond typical infringement such that Plaintiff is entitled under 35 U.S.C. § 284 to enhanced damages up to three times the amount found or assessed.

40. Plaintiff has been damaged as a result of Defendant's infringing conduct described in this Count. Defendant is, thus, liable to Plaintiff in an amount that adequately compensates Plaintiff for Defendant's infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT II

(INFRINGEMENT OF U.S. PATENT NO. 7,441,141)

41. Plaintiff incorporates the preceding paragraphs herein by reference.

42. Plaintiff is the assignee of the '141 patent, entitled "Backup of network devices," with ownership of all substantial rights in the '141 patent, including the right to exclude others and to enforce, sue, and recover damages for past and future infringements.

43. The '141 patent is valid, enforceable, and was duly issued in full compliance with Title 35 of the United States Code. The '141 patent issued from U.S. Patent Application No. 10/993,519.


44. Defendant has and continues to directly and/or indirectly infringe (by inducing infringement) one or more claims of the '141 patent in this judicial district and elsewhere in Texas and the United States.

45. Defendant designs, offers for sale, uses, and sells services, such as its broadband infrastructure and ISP services that use Apache Kafka ("the '141 Accused Products"), in a manner that infringes the '141 patent. For example, Defendant uses Apache Kafka to monitor its network infrastructure:

INSIDE KAFKA STREAMS - MONITORING COMCAST'S OUTSIDE PLANT

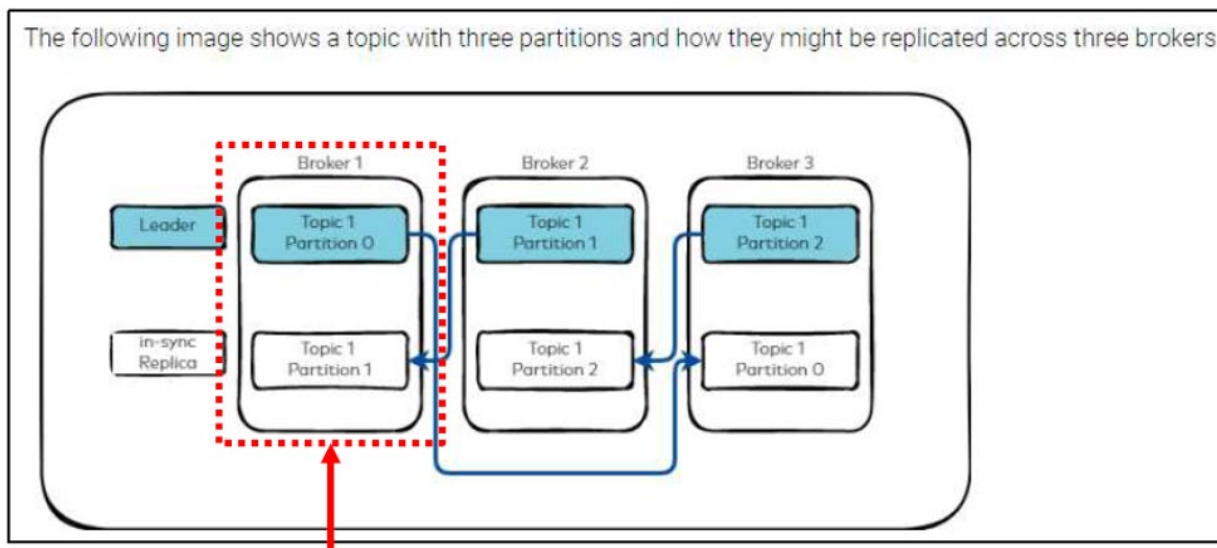
KAFKA CONSUMERS & PRODUCERS

PARTITION MECHANICS



46. Defendant directly infringes the '141 patent under 35 U.S.C. § 271(a) by using, making, offering for sale, selling, and/or importing the '141 Accused Products, their components and processes, and/or products containing the same that incorporate the fundamental technologies covered by the '141 patent.

47. For example, Defendant infringes claim 1 of the '141 patent via the '141 Accused Products. The '141 Accused Products operate at “a first network device of a plurality of network devices each storing device-specific information.” For example, Kafka broker servers are network devices that store device-specific information (i.e., partitions for which they are leaders).



First network device

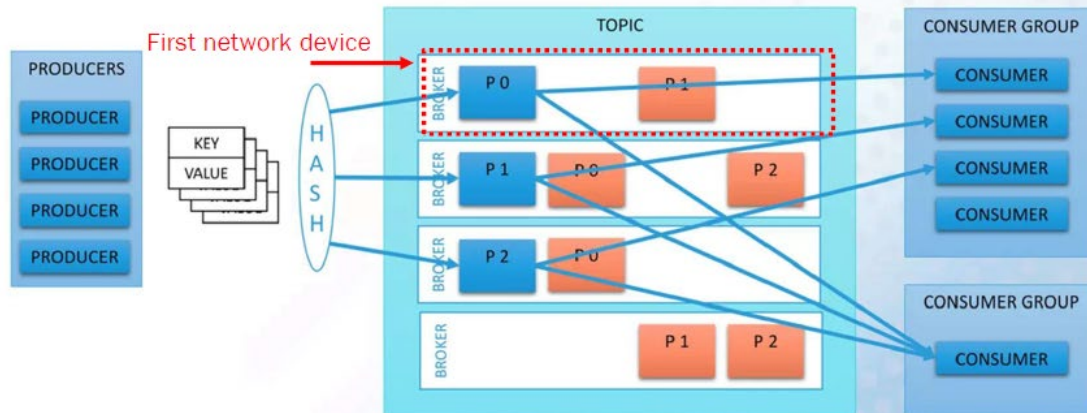
Kafka replicates the log for each topic's partitions across a configurable number of servers (you can set this replication factor on a topic-by-topic basis). This allows automatic failover to these replicas when a server in the cluster fails so messages remain available in the presence of failures.

The unit of replication is the topic partition. Under non-failure conditions, each partition in Kafka has a single leader and zero or more followers. The total number of replicas including the leader constitute the replication factor. All writes go to the leader of the partition, and reads can go to the leader or the followers of the partition. Typically, there are many more partitions than brokers and the leaders are evenly distributed among brokers. The logs on the followers are identical to the leader's log—all have the same offsets and messages in the same order (though, of course, at any given time the leader may have a few as-yet unreplicated messages at the end of its log).

Source: <https://kafka.apache.org/documentation/#replication>.

KAFKA CONSUMERS & PRODUCERS

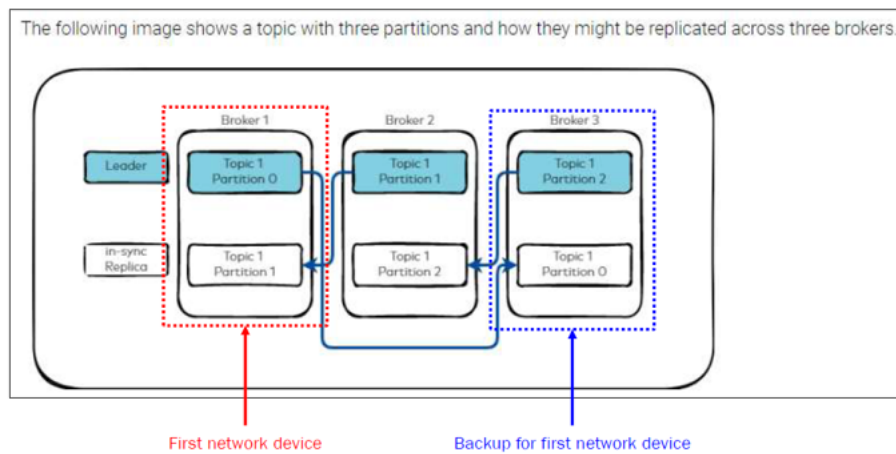
PARTITION MECHANICS



5

COMCAST

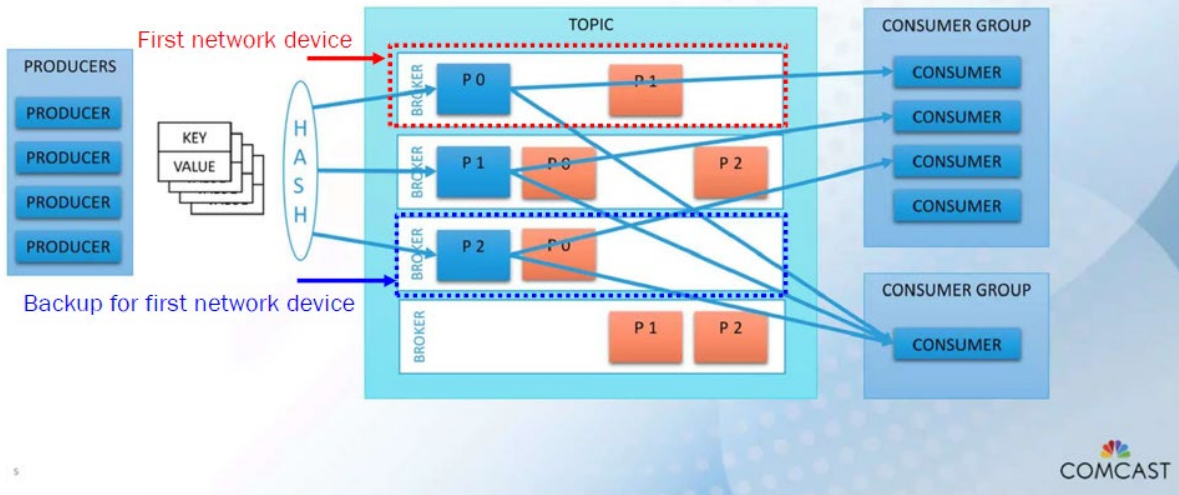
48. The '141 Accused Products select “at least one second network device of said plurality of network devices to act as a backup for said first network device.” For example, a leader broker selects a second broker to act as a backup:



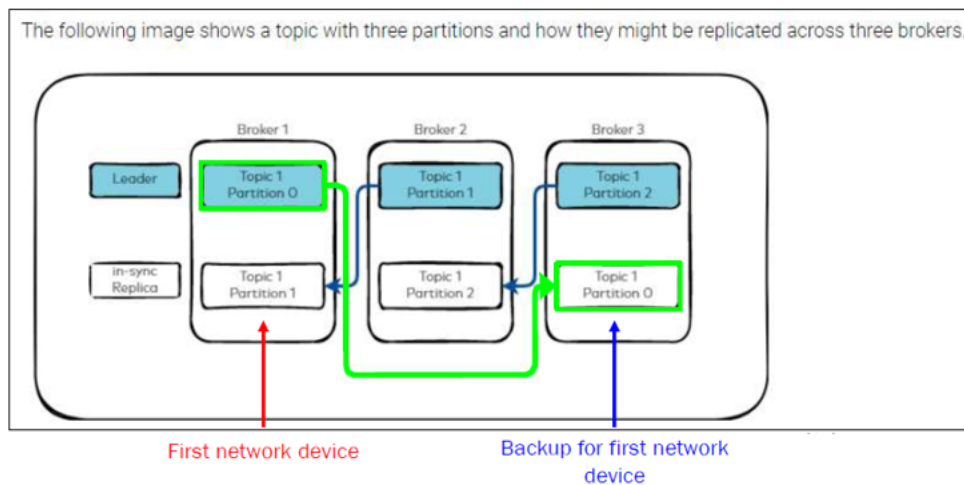
Source: <https://docs.confluent.io/kafka/design/replication.html>.

KAFKA CONSUMERS & PRODUCERS

PARTITION MECHANICS

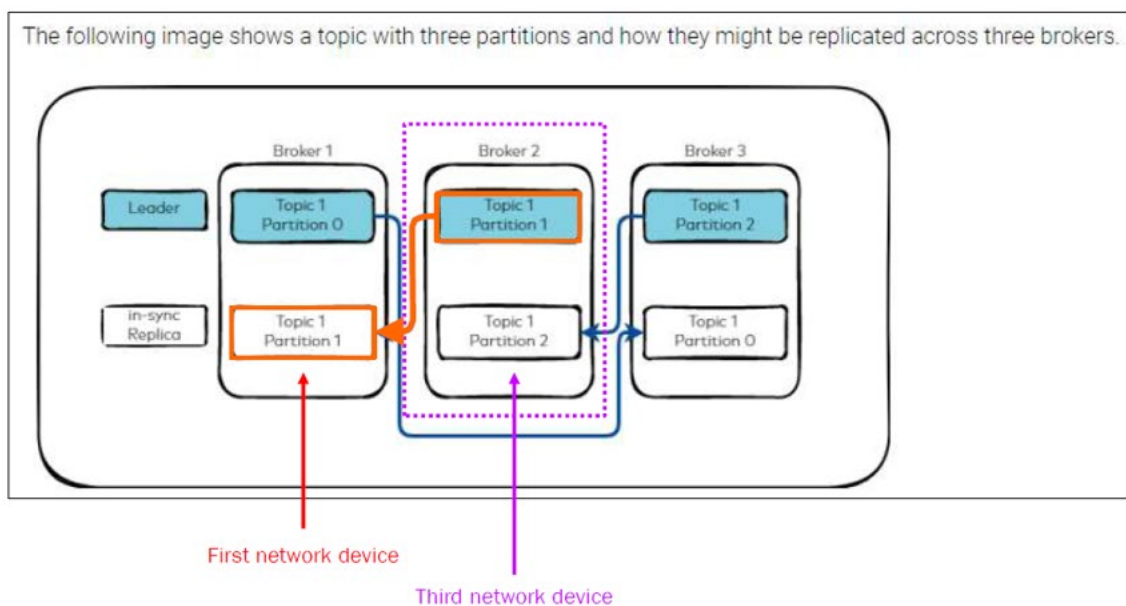


49. The '141 Accused Products communicate “the device-specific information maintained by said first network device to said at least one second network device, said communicated device-specific information for use by said at least one second network device in assuming the role of said first network device upon unavailability of said first network device.” For example, the Kafka leader broker communicates message records (e.g., Partition 0) to the backup broker so that the backup broker can assume the role of leader if the leader broker is unavailable:



Source: <https://docs.confluent.io/kafka/design/replication.html>.

50. The '141 Accused Products thereafter receive “at said first network device device-specific information from at least one third network device for use by said first network device in assuming the role of the third network device upon unavailability of the third network device.” For example, the leader broker may be a follower broker for a different topic partition (e.g., Partition 1):



Source: <https://docs.confluent.io/kafka/design/replication.html>.

51. The '141 Accused Products operate such that “when the device-specific information of said first network device is requested and said first network device is unavailable, communicating the device-specific information of said first network device from one of said at least one second network device.” For example, when messages stored in the topic partition of the leader broker is requested and the leader broker is unavailable, the messages of the topic partition are consumed from the follower broker.

Of course if leaders didn't fail we wouldn't need followers! When the leader does die we need to choose a new leader from among the followers. But followers themselves may fall behind or crash so we must ensure we choose an up-to-date follower. The fundamental guarantee a log replication algorithm must provide is that if we tell the client a message is committed, and the leader fails, the new leader we elect must also have that message. This yields a tradeoff: if the leader waits for more followers to acknowledge a message before declaring it committed then there will be more potentially electable leaders.

Source: https://kafka.apache.org/documentation/#design_replicatedlog.

The unit of replication is the topic partition. Under non-failure conditions, each partition in Kafka has a single leader and zero or more followers. The total number of replicas including the leader constitute the replication factor. All writes go to the leader of the partition, and reads can go to the leader or the followers of the partition. Typically, there are many more partitions than brokers and the leaders are evenly distributed among brokers. The logs on the followers are identical to the leader's log—all have the same offsets and messages in the same order (though, of course, at any given time the leader may have a few as-yet unreplicated messages at the end of its log).

Source: <https://kafka.apache.org/documentation/#replication>.

52. The technology discussion above and the exemplary '141 Accused Products provide context for Plaintiff's infringement allegations.

53. At a minimum, Defendant has known of the '141 patent since at least the service of this complaint on September 26, 2024. Further, Defendant has known of the '141 patent since at least the filing date of the complaint on September 20, 2024. In addition, Defendant has known about the '141 patent since at least September 16, 2024, when it received email correspondence from Plaintiff alerting Defendant to its infringement. Moreover, Defendant has known about the '141 patent since at least September 13, 2024, when Plaintiff sent correspondence to Defendant via FedEx alerting Defendant to its infringement.

54. On information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has actively induced, under U.S.C. § 271(b), its distributors, customers, subsidiaries, importers, and/or consumers that import, purchase, or sell the

'141 Accused Products that include or are made using all of the limitations of one or more claims of the '141 patent to directly infringe one or more claims of the '141 patent (e.g., claim 1, as discussed above) by using, offering for sale, selling, and/or importing the '141 Accused Products. Since at least the notice provided on the above-mentioned date, Defendant does so with knowledge, or with willful blindness of the fact, that the induced acts constitute infringement of the '141 patent. Defendant intends to cause, and has taken affirmative steps to induce infringement by its distributors, importers, customers, subsidiaries, and/or consumers by at least, inter alia, creating advertisements that promote the infringing use of the '141 Accused Products, creating and/or maintaining established distribution channels for the '141 Accused Products into and within the United States, manufacturing the '141 Accused Products in conformity with U.S. laws and regulations, distributing or making available instructions or manuals for these products to purchasers and prospective buyers, and testing the '141 Accused Products, and/or providing technical support, replacement parts, or services for these products to these purchasers in the United States. For example, Defendant configures Apache Kafka in the '141 Accused Products to contain specific instructions, such as source code and configuration files, that cause such products to monitor its network infrastructure. Moreover, in addition to the foregoing, Defendant encourages application developers to build applications in a manner that uses the '141 Accused Products to automatically restart applications on both Defendant's private data centers and public cloud providers in a manner that necessarily infringes the '141 patent. *See, e.g.,* <https://www.youtube.com/watch?v=oVvoqg-NiKU> (a transcript of which is attached as Exhibit A).

55. In the alternative, on information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has contributorily infringed, under

U.S.C. § 271(c), one or more claims of the '141 patent. For example, Defendant contributes to the direct infringement of such claims by distributors, customers, subsidiaries, importers, and/or consumers that use, import, purchase, or sell the '141 Accused Products. To the extent that the '141 Accused Products do not directly infringe one or more claims of the '141 patent, such products contain instructions, such as source code, that are especially adapted to cause the '141 Accused Products to operate in an infringing manner. Such instructions are specifically designed to cause the '141 Accused Products to provide and utilize Apache Kafka in an infringing manner and are a material part of the invention of the '141 patent and are not a staple article of commerce suitable for substantial non-infringing use.

56. On information and belief, despite having knowledge of the '141 patent and knowledge that it is directly and/or indirectly infringing one or more claims of the '141 patent, Defendant has nevertheless continued its infringing conduct and disregarded an objectively high likelihood of infringement. Defendant's infringing activities relative to the '141 patent have been, and continue to be, willful, wanton, malicious, in bad-faith, deliberate, consciously wrongful, flagrant, characteristic of a pirate, and an egregious case of misconduct beyond typical infringement such that Plaintiff is entitled under 35 U.S.C. § 284 to enhanced damages up to three times the amount found or assessed.

57. Plaintiff has been damaged as a result of Defendant's infringing conduct described in this Count. Defendant is, thus, liable to Plaintiff in an amount that adequately compensates Plaintiff for Defendant's infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT III

(INFRINGEMENT OF U.S. PATENT NO. 8,145,945)

58. Plaintiff incorporates the preceding paragraphs herein by reference.

59. Plaintiff is the assignee of the '945 patent, entitled "Packet mirroring between primary and secondary virtualized software images for improved system failover performance," with ownership of all substantial rights in the '945 patent, including the right to exclude others and to enforce, sue, and recover damages for past and future infringements.

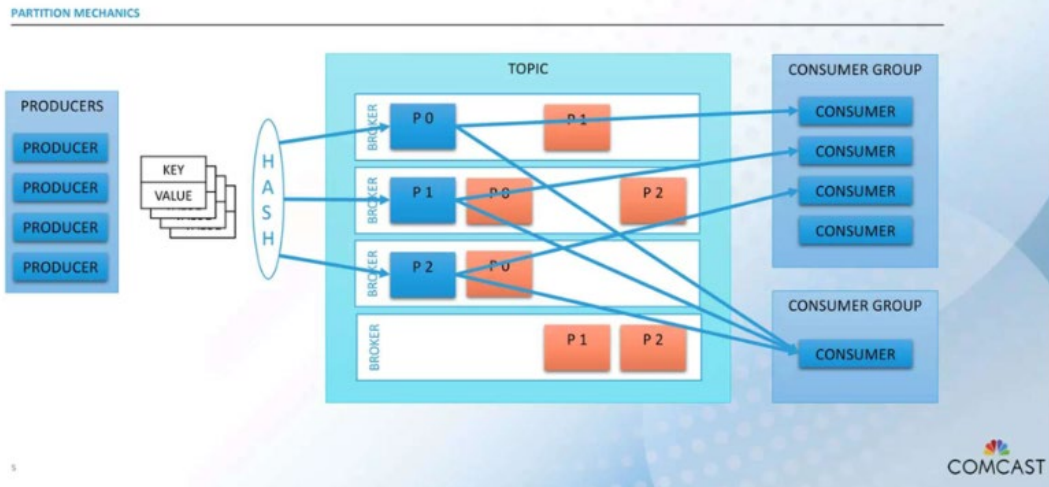
60. The '945 patent is valid, enforceable, and was duly issued in full compliance with Title 35 of the United States Code. The '945 patent issued from U.S. Patent Application No. 12/651,554.

61. Defendant has and continues to directly and/or indirectly infringe (by inducing infringement) one or more claims of the '945 patent in this judicial district and elsewhere in Texas and the United States.

62. Defendant designs, offers for sale, uses, and sells services, such as its broadband infrastructure and ISP services that use Apache Kafka ("the '945 Accused Products"), in a manner that infringes the '945 patent. For example, Defendant uses Apache Kafka to monitor its network infrastructure:

INSIDE KAFKA STREAMS - MONITORING COMCAST'S OUTSIDE PLANT

KAFKA CONSUMERS & PRODUCERS




63. Defendant directly infringes the '945 patent under 35 U.S.C. § 271(a) by using, making, offering for sale, selling, and/or importing the '945 Accused Products, their components and processes, and/or products containing the same that incorporate the fundamental technologies covered by the '945 patent.

64. For example, Defendant infringes claim 1 of the '945 patent via the '945 Accused Products. The '945 Accused Products perform a “method for preserving state and reducing data loss.” For example, Kafka Streams performs a method for preserving state and reducing data loss.

State

Kafka Streams provides so-called *state stores*, which can be used by stream processing applications to store and query data, which is an important capability when implementing *stateful operations*. The *Kafka Streams DSL*, for example, automatically creates and manages such state stores when you are calling stateful operators such as `count()` or `aggregate()`, or when you are *windowing a stream*.

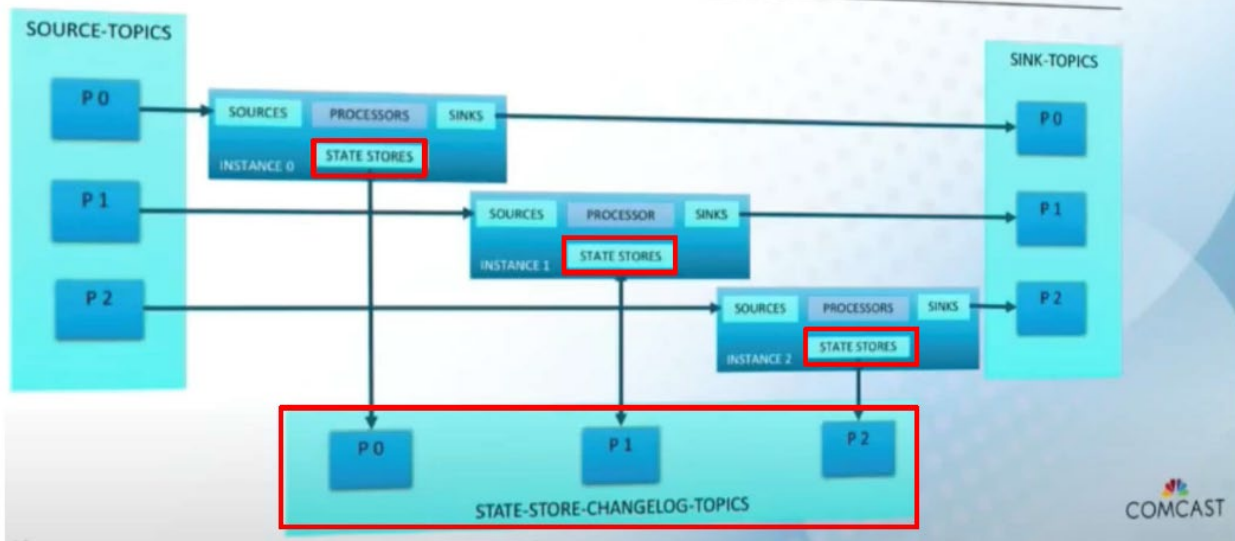
Every stream task in a Kafka Streams application may embed one or more local state stores that can be accessed via APIs to store and query data required for processing. These state stores can either be a [RocksDB](#) database, an in-memory hash map, or another convenient data structure. Kafka Streams offers fault-tolerance and automatic recovery for local state stores.



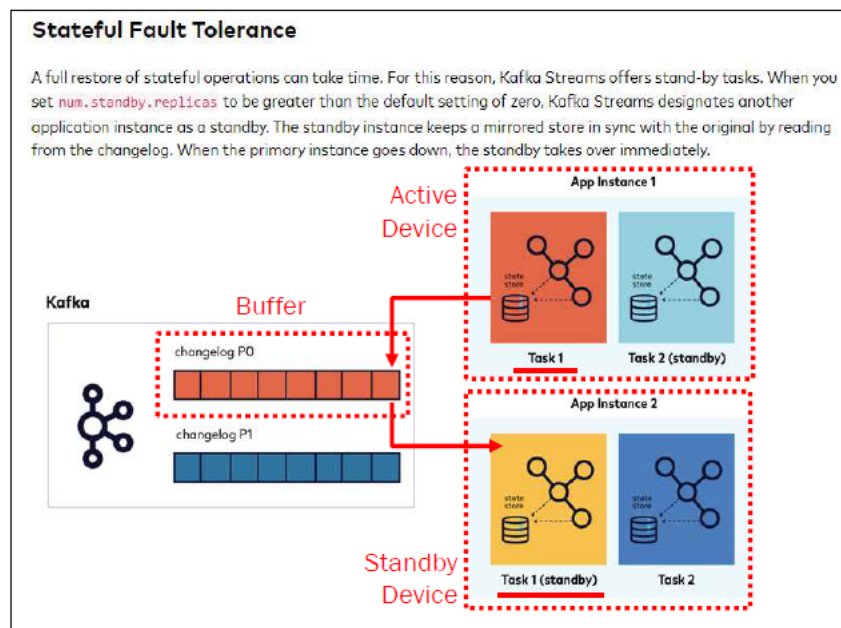
Source: <https://docs.confluent.io/platform/current/streams/architecture.html#state>.

KAFKA STREAMS - PROCESSOR API - PARTITIONING

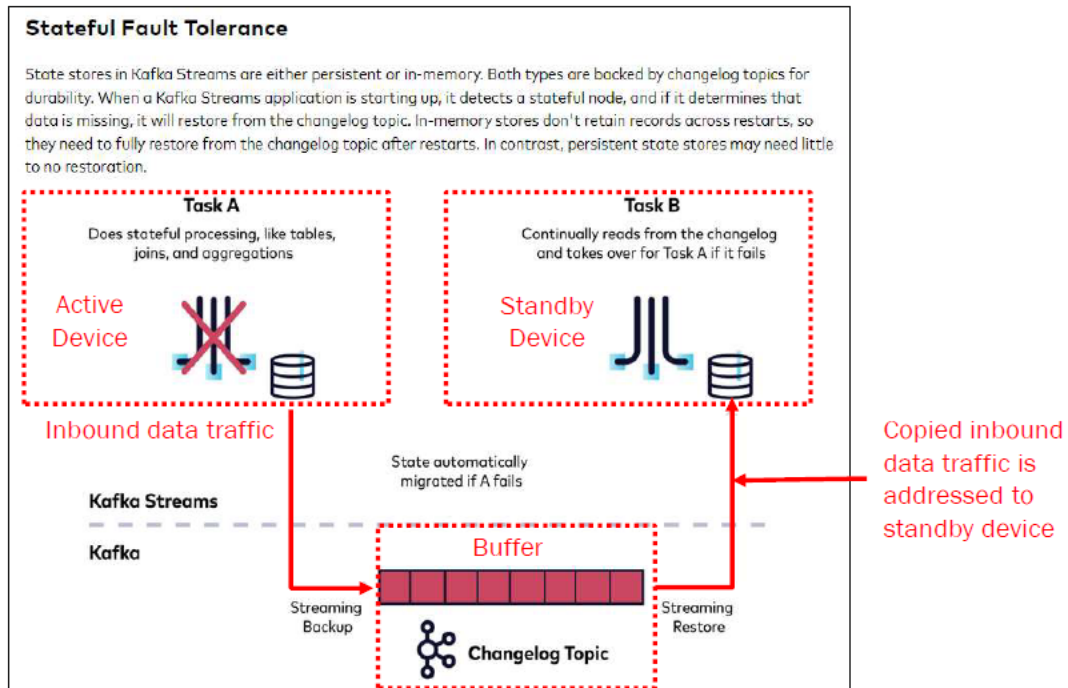
PARTITION MECHANICS IN THE PROCESSOR API



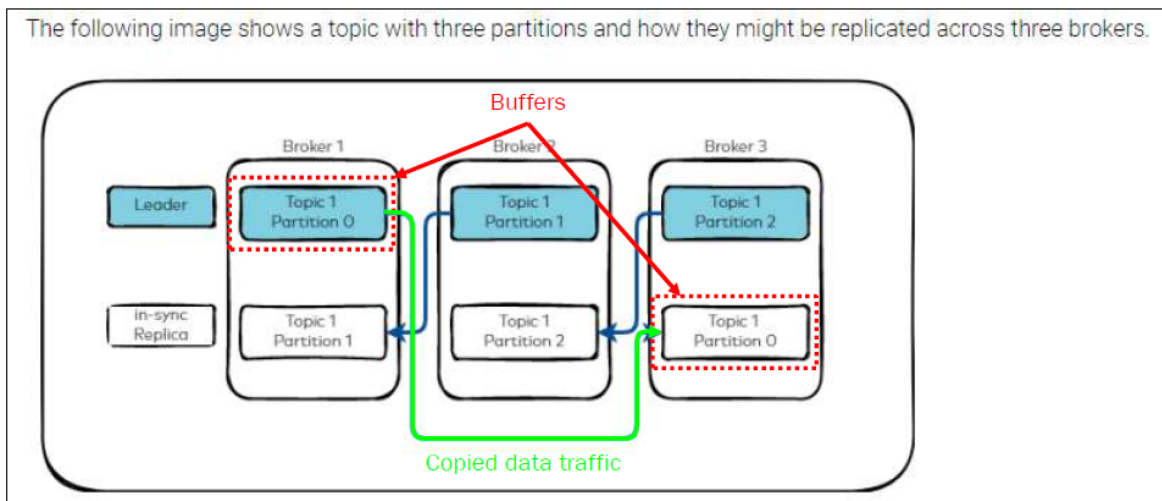
65. The '945 Accused Products operate such that “upon detecting a commit in an active device, continuously copying all inbound data traffic before receipt at the active device to one or more buffers associated with a standby device until a next commit or failure, wherein the copied inbound data traffic has a destination address changed to that of the standby device.” For example, an active instance of a Kafka Streams Application commits state to a local state store that replicated to standby instances through a changelog topic. The data traffic inbound to the changelog is replicated (in alignment with Kafka’s topic replication principles) before the changelog topic is ever received at the active instance. When the state of the active instance is updated, the changelog topic (and its In-Sync Replicas) is updated. The replicated changelog data is addressed to the standby instances:



Source: <https://developer.confluent.io/courses/kafka-streams/stateful-fault-tolerance/>.




Source: <https://developer.confluent.io/courses/kafka-streams/stateful-fault-tolerance/>.

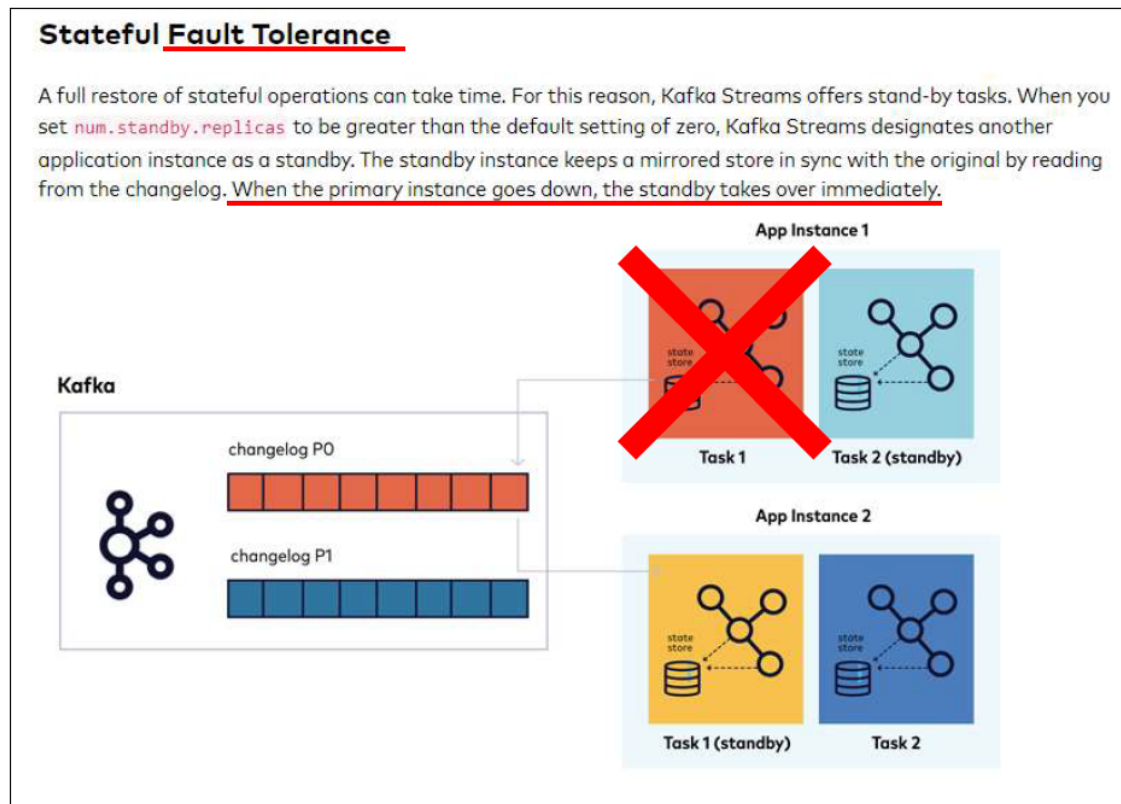


Source: <https://docs.confluent.io/kafka/design/replication.html>.

KAFKA CONSUMERS & PRODUCERS

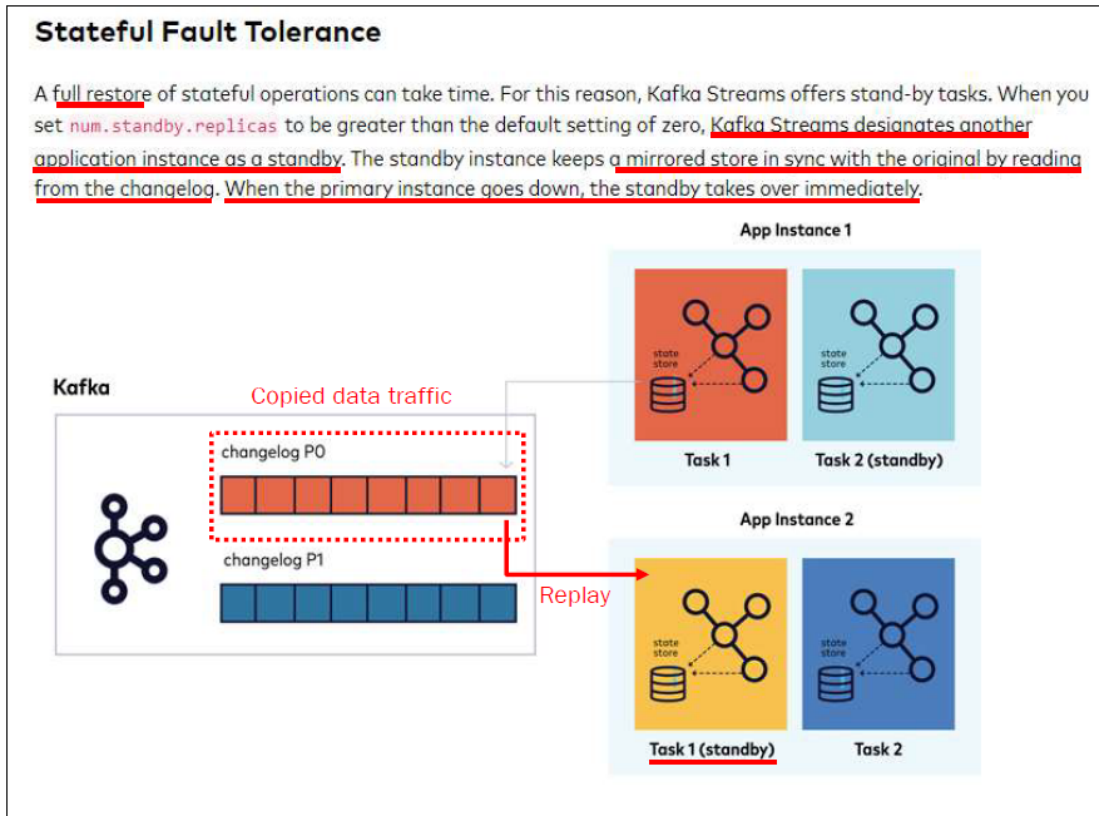


66. The '945 Accused Products detect “a failure.” For example, Kafka Streams is fault-tolerant, so it detects failures:

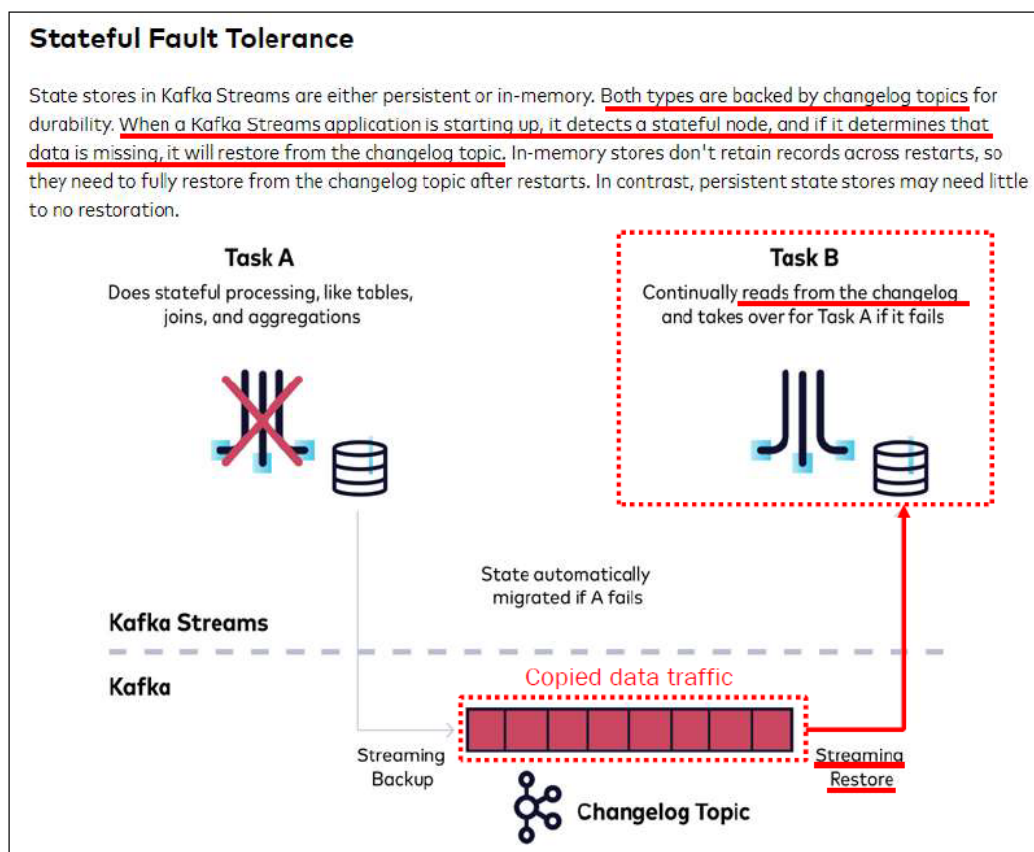


Source: <https://developer.confluent.io/courses/kafka-streams/stateful-fault-tolerance/>.

67. The '945 Accused Products thereafter replay “copied data traffic to restore the standby device to a current state of a failed device.” For example, Kafka replays the changelog to restore standby application instances to the current state of a failed task:



Source: <https://developer.confluent.io/courses/kafka-streams/stateful-fault-tolerance/>.



Source: <https://developer.confluent.io/courses/kafka-streams/stateful-fault-tolerance/>.

68. The technology discussion above and the exemplary '945 Accused Products provide context for Plaintiff's infringement allegations.

69. At a minimum, Defendant has known of the '945 patent since at least the service of this complaint on September 26, 2024. Further, Defendant has known of the '945 patent since at least the filing date of the complaint on September 20, 2024. In addition, Defendant has known about the '945 patent since at least September 16, 2024, when it received email correspondence from Plaintiff alerting Defendant to its infringement. Moreover, Defendant has known about the '945 patent since at least September 13, 2024, when Plaintiff sent correspondence to Defendant via FedEx alerting Defendant to its infringement.

70. On information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has actively induced, under U.S.C. § 271(b), its distributors, customers, subsidiaries, importers, and/or consumers that import, purchase, or sell the '945 Accused Products that include or are made using all of the limitations of one or more claims of the '945 patent to directly infringe one or more claims of the '945 patent (e.g., claim 1, as discussed above) by using, offering for sale, selling, and/or importing the '945 Accused Products. Since at least the notice provided on the above-mentioned date, Defendant does so with knowledge, or with willful blindness of the fact, that the induced acts constitute infringement of the '945 patent. Defendant intends to cause, and has taken affirmative steps to induce infringement by its distributors, importers, customers, subsidiaries, and/or consumers by at least, *inter alia*, creating advertisements that promote the infringing use of the '945 Accused Products, creating and/or maintaining established distribution channels for the '945 Accused Products into and within the United States, manufacturing the '945 Accused Products in conformity with U.S. laws and regulations, distributing or making available instructions or manuals for these products to purchasers and prospective buyers, and testing the '945 Accused Products, and/or providing technical support, replacement parts, or services for these products to these purchasers in the United States. For example, Defendant configures Apache Kafka in the '945 Accused Products to contain specific instructions, such as source code and configuration files, that cause such products to monitor its network infrastructure. Moreover, in addition to the foregoing, Defendant encourages application developers to build applications in a manner that uses the '945 Accused Products to automatically restart applications on both Defendant's private data centers and public cloud providers in a manner that necessarily infringes the '945 patent. *See, e.g.,*

<https://www.youtube.com/watch?v=oVvoqg-NiKU> (a transcript of which is attached as Exhibit A).

71. In the alternative, on information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has contributorily infringed, under U.S.C. § 271(c), one or more claims of the '945 patent. For example, Defendant contributes to the direct infringement of such claims by distributors, customers, subsidiaries, importers, and/or consumers that use, import, purchase, or sell the '945 Accused Products. To the extent that the '945 Accused Products do not directly infringe one or more claims of the '945 patent, such products contain instructions, such as source code, that are especially adapted to cause the '945 Accused Products to operate in an infringing manner. Such instructions are specifically designed to cause the '945 Accused Products to provide and utilize Apache Kafka in an infringing manner and are a material part of the invention of the '945 patent and are not a staple article of commerce suitable for substantial non-infringing use.

72. On information and belief, despite having knowledge of the '945 patent and knowledge that it is directly and/or indirectly infringing one or more claims of the '945 patent, Defendant has nevertheless continued its infringing conduct and disregarded an objectively high likelihood of infringement. Defendant's infringing activities relative to the '945 patent have been, and continue to be, willful, wanton, malicious, in bad-faith, deliberate, consciously wrongful, flagrant, characteristic of a pirate, and an egregious case of misconduct beyond typical infringement such that Plaintiff is entitled under 35 U.S.C. § 284 to enhanced damages up to three times the amount found or assessed.

73. Plaintiff has been damaged as a result of Defendant's infringing conduct described in this Count. Defendant is, thus, liable to Plaintiff in an amount that adequately compensates

Plaintiff for Defendant's infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

COUNT IV

(INFRINGEMENT OF U.S. PATENT NO. 9,026,836)

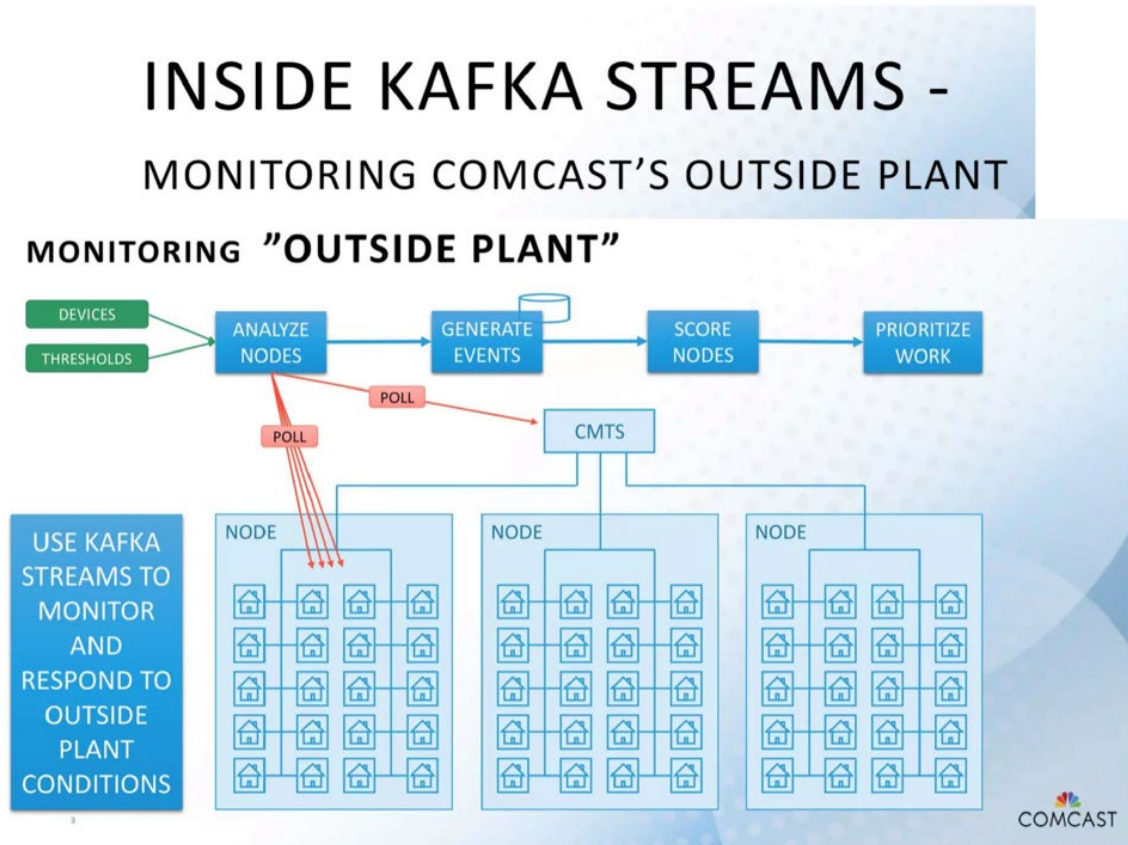
74. Plaintiff incorporates the preceding paragraphs herein by reference.

75. Plaintiff is the assignee of the '836 patent, entitled "Call restoration in response to application failure," with ownership of all substantial rights in the '836 patent, including the right to exclude others and to enforce, sue, and recover damages for past and future infringements.

76. The '836 patent is valid, enforceable, and was duly issued in full compliance with Title 35 of the United States Code. The '836 patent issued from U.S. Patent Application No. 13/476,789.

77. Defendant has and continues to directly and/or indirectly infringe (by inducing infringement) one or more claims of the '836 patent in this judicial district and elsewhere in Texas and the United States.

78. Defendant designs, offers for sale, uses, and sells services, such as its broadband infrastructure and ISP services that use Apache Kafka ("the '836 Accused Products"), in a manner that infringes the '836 patent. For example, Defendant uses Apache Kafka to monitor its network infrastructure:



79. Defendant directly infringes the '836 patent under 35 U.S.C. § 271(a) by using, making, offering for sale, selling, and/or importing the '836 Accused Products, their components and processes, and/or products containing the same that incorporate the fundamental technologies covered by the '836 patent.

80. For example, Defendant infringes claim 1 of the '836 patent via the '836 Accused Products. The '836 Accused Products perform a method of restoration in response to application failure.

Kafka Streams is a client library for processing and analyzing data stored in Kafka. It builds upon important stream processing concepts such as properly distinguishing between event time and processing time, windowing support, and simple yet efficient management and real-time querying of application state.

Kafka Streams has a **low barrier to entry**: You can quickly write and run a small-scale proof-of-concept on a single machine; and you only need to run additional instances of your application on multiple machines to scale up to high-volume production workloads. Kafka Streams transparently handles the load balancing of multiple instances of the same application by leveraging Kafka's parallelism model.

Some highlights of Kafka Streams:

- Designed as a **simple and lightweight client library**, which can be easily embedded in any Java application and integrated with any existing packaging, deployment and operational tools that users have for their streaming applications.
- Has **no external dependencies on systems other than Apache Kafka itself** as the internal messaging layer; notably, it uses Kafka's partitioning model to horizontally scale processing while maintaining strong ordering guarantees.
- Supports **fault-tolerant local state**, which enables very fast and efficient stateful operations like windowed joins and aggregations.
- Supports **exactly-once** processing semantics to guarantee that each record will be processed once and only once even when there is a failure on either Streams clients or Kafka brokers in the middle of processing.
- Employs **one-record-at-a-time processing** to achieve millisecond processing latency, and supports **event-time based windowing operations** with out-of-order arrival of records.
- Offers necessary stream processing primitives, along with a **high-level Streams DSL** and a **low-level Processor API**.

Source: <https://kafka.apache.org/36/documentation/streams/core-concepts>.


81. The '836 Accused Products determine “by a processor, that an application in an application sequence has failed during a communication session that is associated with the application sequence.” For example, Kafka determines that a stream processor in a Kafka Streams application processor topology has failed while processing the stream. The stream is associated with the processor topology:

Stream Processing Topology

- A **stream** is the most important abstraction provided by Kafka Streams: it represents an unbounded, continuously updating data set. A stream is an ordered, replayable, and fault-tolerant sequence of immutable data records, where a **data record** is defined as a key-value pair.
- A **stream processing application** is any program that makes use of the Kafka Streams library. It defines its computational logic through one or more **processor topologies**, where a processor topology is a graph of stream processors (nodes) that are connected by streams (edges).
- A **stream processor** is a node in the processor topology; it represents a processing step to transform data in streams by receiving one input record at a time from its upstream processors in the topology, applying its operation to it, and may subsequently produce one or more output records to its downstream processors.

https://kafka.apache.org/36/documentation/streams/core-concepts#streams_topology

Slightly simplified, the maximum parallelism at which your application may run is bounded by the maximum number of stream tasks, which itself is determined by maximum number of partitions of the input topic(s) the application is reading from. For example, if your input topic has 5 partitions, then you can run up to 5 applications instances. These instances will collaboratively process the topic's data. If you run a larger number of app instances than partitions of the input topic, the "excess" app instances will launch but remain idle; however, if one of the busy instances goes down, one of the idle instances will resume the former's work.



Source: <https://kafka.apache.org/36/documentation/streams/architecture>.

82. The '836 Accused Products send "by the processor, a re-establishment message to a replacement application, the re-establishment message referencing an identifier of at least one of the communication session and its dialog." For example, Kafka sends a re-establishment message to an idle task instance, referencing an identifier of the stream (the stream identifier is referenced so that the idle instance may know which stream to process):

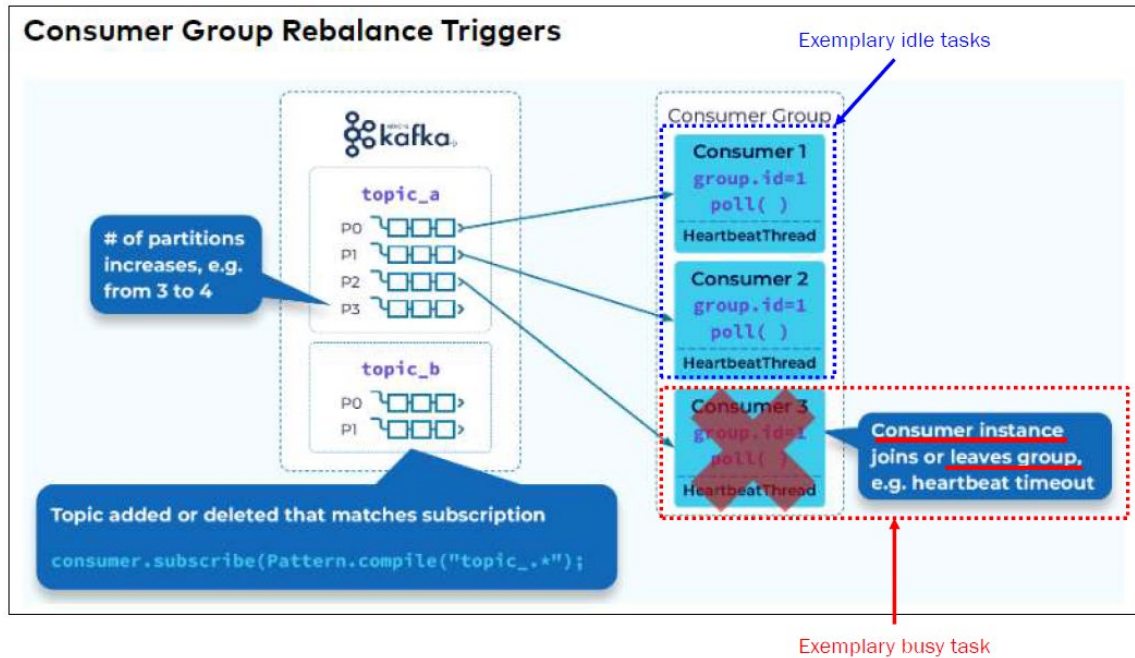
Fault Tolerance

Kafka Streams builds on fault-tolerance capabilities integrated natively within Kafka. Kafka partitions are highly available and replicated; so when stream data is persisted to Kafka it is available even if the application fails and needs to re-process it. Tasks in Kafka Streams leverage the fault-tolerance capability offered by the Kafka consumer client to handle failures. If a task runs on a machine that fails, Kafka Streams automatically restarts the task in one of the remaining running instances of the application.

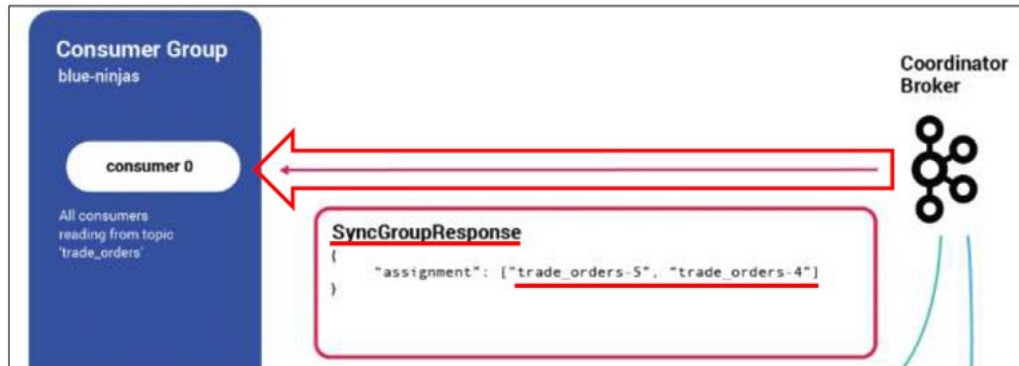

In addition, Kafka Streams makes sure that the local state stores are robust to failures, too. For each state store, it maintains a replicated changelog Kafka topic in which it tracks any state updates. These changelog topics are partitioned as well so that each local state store instance, and hence the task accessing the store, has its own dedicated changelog topic partition. Log compaction is enabled on the changelog topics so that old data can be purged safely to prevent the topics from growing indefinitely. If tasks run on a machine that fails and are restarted on another machine, Kafka Streams guarantees to restore their associated state stores to the content before the failure by replaying the corresponding changelog topics prior to resuming the processing on the newly started tasks. As a result, failure handling is completely transparent to the end user.

Note that the cost of task (re)initialization typically depends primarily on the time for restoring the state by replaying the state stores' associated changelog topics. To minimize this restoration time, users can configure their applications to have **standby replicas** of local states (i.e. fully replicated copies of the state). When a task migration happens, Kafka Streams will assign a task to an application instance where such a standby replica already exists in order to minimize the task (re)initialization cost. See `num.standby.replicas` in the [Kafka Streams Configs](#) section. Starting in 2.6, Kafka Streams will guarantee that a task is only ever assigned to an instance with a fully caught-up local copy of the state, if such an instance exists. Standby tasks will increase the likelihood that a caught-up instance exists in the case of a failure.


Source: <https://kafka.apache.org/36/documentation/streams/architecture>. The consumer client fault tolerance (which is applicable to Kafka Streams Applications) behaves in accordance with Consumer Group Rebalance.



Kafka sends a SyncGroupResponse to the idle task instance. The SyncGroupResponse references the partitions that the failed stream processor was processing.



83. The '836 Accused Products reconstruct “by the processor, the application sequence for the communication session while the communication session is still in progress so that the reconstructed application sequence includes the replacement application.” For example, Kafka reconstructs the processor topology for the stream so that the new instance of the task includes the replacement stream processor and can continue processing forward from the partition state:



Source: <https://kafka.apache.org/36/documentation/streams/architecture>.

84. The technology discussion above and the exemplary '836 Accused Products provide context for Plaintiff's infringement allegations.

85. At a minimum, Defendant has known of the '836 patent since at least the service of this complaint on September 26, 2024. Further, Defendant has known of the '836 patent since at least the filing date of the complaint on September 20, 2024. In addition, Defendant has known about the '836 patent since at least September 16, 2024, when it received email correspondence from Plaintiff alerting Defendant to its infringement. Moreover, Defendant has known about the '836 patent since at least September 13, 2024, when Plaintiff sent correspondence to Defendant via FedEx alerting Defendant to its infringement.

86. On information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has actively induced, under U.S.C. §

271(b), its distributors, customers, subsidiaries, importers, and/or consumers that import, purchase, or sell the '836 Accused Products that include or are made using all of the limitations of one or more claims of the '836 patent to directly infringe one or more claims of the '836 patent (e.g., claim 1, as discussed above) by using, offering for sale, selling, and/or importing the '836 Accused Products. Since at least the notice provided on the above-mentioned date, Defendant does so with knowledge, or with willful blindness of the fact, that the induced acts constitute infringement of the '836 patent. Defendant intends to cause, and has taken affirmative steps to induce infringement by its distributors, importers, customers, subsidiaries, and/or consumers by at least, inter alia, creating advertisements that promote the infringing use of the '836 Accused Products, creating and/or maintaining established distribution channels for the '836 Accused Products into and within the United States, manufacturing the '836 Accused Products in conformity with U.S. laws and regulations, distributing or making available instructions or manuals for these products to purchasers and prospective buyers, and testing the '836 Accused Products, and/or providing technical support, replacement parts, or services for these products to these purchasers in the United States. For example, Defendant configures Apache Kafka in the '836 Accused Products to contain specific instructions, such as source code and configuration files, that cause such products to monitor its network infrastructure. Moreover, in addition to the foregoing, Defendant encourages application developers to build applications in a manner that uses the '836 Accused Products to automatically restart applications on both Defendant's private data centers and public cloud providers in a manner that necessarily infringes the '836 patent. *See, e.g.,* <https://www.youtube.com/watch?v=oVvoqg-NiKU> (a transcript of which is attached as Exhibit A).

87. In the alternative, on information and belief, since at least the above-mentioned date when Defendant was on notice of its infringement, Defendant has contributorily infringed, under U.S.C. § 271(c), one or more claims of the '836 patent. For example, Defendant contributes to the direct infringement of such claims by distributors, customers, subsidiaries, importers, and/or consumers that use, import, purchase, or sell the '836 Accused Products. To the extent that the '836 Accused Products do not directly infringe one or more claims of the '836 patent, such products contain instructions, such as source code, that are especially adapted to cause the '836 Accused Products to operate in an infringing manner. Such instructions are specifically designed to cause the '836 Accused Products to provide and utilize Apache Kafka in an infringing manner and are a material part of the invention of the '836 patent and are not a staple article of commerce suitable for substantial non-infringing use.

88. On information and belief, despite having knowledge of the '836 patent and knowledge that it is directly and/or indirectly infringing one or more claims of the '836 patent, Defendant has nevertheless continued its infringing conduct and disregarded an objectively high likelihood of infringement. Defendant's infringing activities relative to the '836 patent have been, and continue to be, willful, wanton, malicious, in bad-faith, deliberate, consciously wrongful, flagrant, characteristic of a pirate, and an egregious case of misconduct beyond typical infringement such that Plaintiff is entitled under 35 U.S.C. § 284 to enhanced damages up to three times the amount found or assessed.

89. Plaintiff has been damaged as a result of Defendant's infringing conduct described in this Count. Defendant is, thus, liable to Plaintiff in an amount that adequately compensates Plaintiff for Defendant's infringements, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court under 35 U.S.C. § 284.

CONCLUSION

90. Plaintiff is entitled to recover from Defendant the damages sustained by Plaintiff as a result of Defendant's wrongful acts, and willful infringement, in an amount subject to proof at trial, which, by law, cannot be less than a reasonable royalty, together with interest and costs as fixed by this Court.

91. Plaintiff has incurred and will incur attorneys' fees, costs, and expenses in the prosecution of this action. The circumstances of this dispute may give rise to an exceptional case within the meaning of 35 U.S.C. § 285, and Plaintiff is entitled to recover its reasonable and necessary attorneys' fees, costs, and expenses.

JURY DEMAND

92. Plaintiff hereby requests a trial by jury pursuant to Rule 38 of the Federal Rules of Civil Procedure.

PRAYER FOR RELIEF

93. Plaintiff respectfully requests that the Court find in its favor and against Defendant, and that the Court grant Plaintiff the following relief:

1. A judgment that Defendant has infringed the Asserted Patents as alleged herein, directly and/or indirectly by way of inducing infringement of such patents;
2. A judgment for an accounting of all damages sustained by Plaintiff as a result of the acts of infringement by Defendant;
3. A judgment and order requiring Defendant to pay Plaintiff damages under 35 U.S.C. § 284, including up to treble damages as provided by 35 U.S.C. § 284, and any royalties determined to be appropriate;

4. A judgment and order requiring Defendant to pay Plaintiff pre-judgment and post-judgment interest on the damages awarded;
5. A judgment and order finding this to be an exceptional case and requiring Defendant to pay the costs of this action (including all disbursements) and attorneys' fees as provided by 35 U.S.C. § 285; and
6. Such other and further relief as the Court deems just and equitable.

Dated: December 18, 2024

Respectfully submitted,

/s/ Patrick J. Conroy

Patrick J. Conroy (Lead Counsel)

Texas Bar No. 24012448

Justin B. Kimble

Texas Bar No. 24036909

Jon Rastegar

Texas Bar No. 24064043

Nathan L. Levenson

Texas Bar No. 24097992

Nelson Bumgardner Conroy PC

2727 N. Harwood St.

Suite 250

Dallas, TX 75201

Tel: (817) 377-9111

pat@nelbum.com

justin@nelbum.com

jon@nelbum.com

nathan@nelbum.com

Attorneys for Plaintiff

Arlington Technologies LLC

CERTIFICATE OF SERVICE

I hereby certify that a copy of the foregoing document was filed electronically in compliance with Local Rule CV-5(a). Therefore, this document was served on all counsel who are deemed to have consented to electronic service on December 18, 2024.

/s/ Patrick J. Conroy