

KIP-766: fetch/findSessions queries with open endpoints for SessionStore/WindowStore

- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
- [Proposed Changes](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Rejected Alternatives](#)
 - [Adding new method methods for unbounded range queries](#)
 - [Modify existing range interface](#)

Status

Current state: *"Accept"*

Discussion thread: [here](#)

JIRA: [here](#)

Motivation

This is a follow-up KIP for [KIP-763: Range queries with open endpoints](#). In KIP-763, we focused on **ReadOnlyKeyValueStore**, in this KIP, we'll focus on **ReadOnlySessionStore** and **ReadOnlyWindowStore**.

Currently, **ReadOnlySessionStore/ReadOnlyWindowStore** in Kafka streams only support bounded ranged queries. More precisely, the `fetch/backwardFetch`, and `findSessions/backwardFindSessions` interfaces.

In the **ReadOnlyWindowStore** looks as follows:

```

1  /**
2
3  * Get all the key-value pairs in the given key range and time range from all the existing windows.
4  * <p>
5  * This iterator must be closed after use.
6  *
7  * @param keyFrom      the first key in the range
8  * @param keyTo        the last key in the range
9  * @param timeFrom     time range start (inclusive), where iteration starts.
10 * @param timeTo      time range end (inclusive), where iteration ends.
11 * @return an iterator over windowed key-value pairs {@code <Windowed<K>, value>}, from beginning
12 to end of time.
13 * @throws InvalidStateStoreException if the store is not initialized
14 * @throws NullPointerException      if {@code null} is used for any key.
15 * @throws IllegalArgumentException  if duration is negative or can't be represented as {@code lon
16 g milliseconds}
17 */
18 KeyValueIterator<Windowed<K>, V> fetch(K keyFrom, K keyTo, Instant timeFrom, Instant timeTo)
19     throws IllegalArgumentException;
20
21 /**
22 * Get all the key-value pairs in the given key range and time range from all the existing windows
23 * in backward order with respect to time (from end to beginning of time).
24 * <p>
25 * This iterator must be closed after use.
26 *
27 * @param keyFrom      the first key in the range
28 * @param keyTo        the last key in the range
29 * @param timeFrom     time range start (inclusive), where iteration ends.
30 * @param timeTo      time range end (inclusive), where iteration starts.
31 * @return an iterator over windowed key-value pairs {@code <Windowed<K>, value>}, from end to
32 beginning of time.
33 * @throws InvalidStateStoreException if the store is not initialized
34 * @throws NullPointerException      if {@code null} is used for any key.
35 * @throws IllegalArgumentException  if duration is negative or can't be represented as {@code lon
36 g milliseconds}
37 */
38 default KeyValueIterator<Windowed<K>, V> backwardFetch(K keyFrom, K keyTo, Instant timeFrom,
39 Instant timeTo)
40     throws IllegalArgumentException {
41     throw new UnsupportedOperationException();
42 }

```

In the `ReadOnlySessionStore` looks as follows:

```

1  /**
2  * Fetch any sessions in the given range of keys and the sessions end is &ge;
3  * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
4  * earliest to latest.
5  * <p>
6  * This iterator must be closed after use.
7  *
8  * @param keyFrom      The first key that could be in the range
9  * @param keyTo        The last key that could be in the range
10 * @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
11 * iteration starts.
12 * @param latestSessionStartTime the end timestamp of the latest session to search for, where
13 * iteration ends.
14 * @return iterator of sessions with the matching keys and aggregated values, from earliest to
15 * latest session time.
16 * @throws NullPointerException If null is used for any key.
17 */
18 default KeyValueIterator<Windowed<K>, AGG> findSessions(final K keyFrom,
19 final K keyTo,
20 final long earliestSessionEndTime,
21 final long latestSessionStartTime) {
22     throw new UnsupportedOperationException(
23         "This API is not supported by this implementation of ReadOnlySessionStore.");
24 }
25 /**
26 * Fetch any sessions in the given range of keys and the sessions end is &ge;
27 * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
28 * earliest to latest.
29 * <p>
30 * This iterator must be closed after use.
31 *

```

```

* @param keyFrom          The first key that could be in the range
* @param keyTo           The last key that could be in the range
* @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
*                          iteration starts.
* @param latestSessionStartTime the end timestamp of the latest session to search for, where
*                          iteration ends.
* @return iterator of sessions with the matching keys and aggregated values, from earliest to
* latest session time.
* @throws NullPointerException If null is used for any key.
*/
default KeyValueTypeIterator<Windowed<K>, AGG> findSessions(final K keyFrom,
                                                           final K keyTo,
                                                           final Instant earliestSessionEndTime,
                                                           final Instant latestSessionStartTime) {
    throw new UnsupportedOperationException(
        "This API is not supported by this implementation of ReadOnlySessionStore.");
}

/**
 * Fetch any sessions in the given range of keys and the sessions end is &ge;
 * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
 * latest to earliest.
 * <p>
 * This iterator must be closed after use.
 *
 * @param keyFrom          The first key that could be in the range
 * @param keyTo           The last key that could be in the range
 * @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
 *                          iteration ends.
 * @param latestSessionStartTime the end timestamp of the latest session to search for, where
 *                          iteration starts.
 * @return backward iterator of sessions with the matching keys and aggregated values, from
 * latest to earliest session time.
 * @throws NullPointerException If null is used for any key.
 */
default KeyValueTypeIterator<Windowed<K>, AGG> backwardFindSessions(final K keyFrom,
                                                                    final K keyTo,
                                                                    final long earliestSessionEndTime,
                                                                    final long latestSessionStartTime)
    {
        throw new UnsupportedOperationException(
            "This API is not supported by this implementation of ReadOnlySessionStore.");
    }

/**
 * Fetch any sessions in the given range of keys and the sessions end is &ge;
 * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
 * latest to earliest.
 * <p>
 * This iterator must be closed after use.
 *
 * @param keyFrom          The first key that could be in the range
 * @param keyTo           The last key that could be in the range
 * @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
 *                          iteration ends.
 * @param latestSessionStartTime the end timestamp of the latest session to search for, where
 *                          iteration starts.
 * @return backward iterator of sessions with the matching keys and aggregated values, from
 * latest to earliest session time.
 * @throws NullPointerException If null is used for any key.
 */
default KeyValueTypeIterator<Windowed<K>, AGG> backwardFindSessions(final K keyFrom,
                                                                    final K keyTo,
                                                                    final Instant
                                                                    earliestSessionEndTime,
                                                                    final Instant
                                                                    latestSessionStartTime) {
    throw new UnsupportedOperationException(
        "This API is not supported by this implementation of ReadOnlySessionStore.");
}

```

```

/**
 * Retrieve all aggregated sessions for the given range of keys. This iterator must be closed
 * after use.
 * <p>
 * For each key, the iterator guarantees ordering of sessions, starting from the oldest/earliest
 * available session to the newest/latest session.
 *
 * @param keyFrom first key in the range to find aggregated session values for
 * @param keyTo last key in the range to find aggregated session values for
 * @return KeyValueTypeIterator containing all sessions for the provided key, from oldest to newest
 * session.
 * @throws NullPointerException If null is used for any of the keys.
 */
KeyValueTypeIterator<Windowed<K>, AGG> fetch(final K keyFrom, final K keyTo);

/**
 * Retrieve all aggregated sessions for the given range of keys. This iterator must be closed
 * after use.
 * <p>
 * For each key, the iterator guarantees ordering of sessions, starting from the newest/latest
 * available session to the oldest/earliest session.
 *
 * @param keyFrom first key in the range to find aggregated session values for
 * @param keyTo last key in the range to find aggregated session values for
 * @return backward KeyValueTypeIterator containing all sessions for the provided key, from newest
 * to oldest session.
 * @throws NullPointerException If null is used for any of the keys.
 */
default KeyValueTypeIterator<Windowed<K>, AGG> backwardFetch(final K keyFrom, final K keyTo) {
    throw new UnsupportedOperationException(
        "This API is not supported by this implementation of ReadOnlySessionStore.");
}

```

The requested `fetch/findSessions` is bounded because both the `keyFrom` and the `keyTo` parameter needs to be a valid object of key type `K` and must not be null.

Unlike bounded queries, unbounded queries are currently not supported in an efficient manner. For instance, the only way to retrieve all the records with a key smaller or greater a certain constant is to query the full range range (e.g., using the `all()` method) and then filter out elements greater or smaller the given value. In this KIP we propose to a solution to directly query unbounded ranges of records from the store store, without the need for extra filtering.

Public Interfaces

We propose to allow the use of `null` values as a way to represent unbounded ranges. With this change, the range interface signature will remain the same, only the set of accepted values for the parameters `keyFrom` and `keyTo` would change, thus, requiring a modifications to the JavaDoc method description. The proposed change is shown below:

In the `ReadOnlyWindowStore` looks as follows:

```

1
2
3 /**
4  * Get all the key-value pairs in the given key range and time range from all the existing windows.
5  * <p>
6  * This iterator must be closed after use.
7  *
8  * @param keyFrom      the first key in the range
9  * A null value indicates a starting position from the first element in the store <-- new added
10 * @param keyTo the last key in the range
11 * A null value indicates that the range ends with the last element in the store. <-- new added
12 * @param timeFrom time range start (inclusive), where iteration starts.
13 * @param timeTo time range end (inclusive), where iteration ends.
14 * @return an iterator over windowed key-value pairs {@code <Windowed<K>, value>}, from beginning
15 to end of time.
16 * @throws InvalidStateStoreException if the store is not initialized
17 * @throws IllegalArgumentException if duration is negative or can't be represented as {@code long
18 milliseconds}
19 */
20 KeyValueTypeIterator<Windowed<K>, V> fetch(K keyFrom, K keyTo, Instant timeFrom, Instant timeTo)
21 throws IllegalArgumentException;
22 /**
23  * Get all the key-value pairs in the given key range and time range from all the existing windows
24  * in backward order with respect to time (from end to beginning of time).
25  * <p>
26  * This iterator must be closed after use.
27  *
28  * @param keyFrom the first key in the range
29  * A null value indicates a starting position from the first element in the store <-- new added
30 * @param keyTo the last key in the range
31 * A null value indicates that the range ends with the last element in the store. <-- new added
32 * @param timeFrom time range start (inclusive), where iteration ends.
33 * @param timeTo time range end (inclusive), where iteration starts.
34 * @return an iterator over windowed key-value pairs {@code <Windowed<K>, value>}, from end to
35 beginning of time.
36 * @throws InvalidStateStoreException if the store is not initialized
37 * @throws IllegalArgumentException if duration is negative or can't be represented as {@code long
38 milliseconds}
39 */
40 default KeyValueTypeIterator<Windowed<K>, V> backwardFetch(K keyFrom, K keyTo, Instant timeFrom,
41 Instant timeTo)
42 throws IllegalArgumentException {
43     throw new UnsupportedOperationException();
44 }

```

In the `ReadOnlySessionStore` looks as follows:

```

1 /**
2  * Fetch any sessions in the given range of keys and the sessions end is &ge;
3  * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
4  * earliest to latest.
5  * <p>
6  * This iterator must be closed after use.
7  *
8  * @param keyFrom      The first key that could be in the range
9  * A null value indicates a starting position from the first element in the store <-- new added
10 * @param keyTo      The last key that could be in the range
11 * A null value indicates that the range ends with the last element in the store. <-- new added
12 * @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
13 * iteration starts.
14 * @param latestSessionStartTime the end timestamp of the latest session to search for, where
15 * iteration ends.
16 * @return iterator of sessions with the matching keys and aggregated values, from earliest to
17 * latest session time.
18 */
19 default KeyValueTypeIterator<Windowed<K>, AGG> findSessions(final K keyFrom,
20 final K keyTo,
21 final long earliestSessionEndTime,
22 final long latestSessionStartTime) {
23     throw new UnsupportedOperationException(
24         "This API is not supported by this implementation of ReadOnlySessionStore.");
25 }
26 /**
27  * Fetch any sessions in the given range of keys and the sessions end is &ge;
28  * earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
29  * earliest to latest.

```

```

* <p>
* This iterator must be closed after use.
*
* @param keyFrom The first key that could be in the range
* A null value indicates a starting position from the first element in the store <-- new added
* @param keyTo The last key that could be in the range
* A null value indicates that the range ends with the last element in the store. <-- new added
* @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
* iteration starts.
* @param latestSessionStartTime the end timestamp of the latest session to search for, where
* iteration ends.
* @return iterator of sessions with the matching keys and aggregated values, from earliest to
* latest session time.
*/
default KeyValueIterator<Windowed<K>, AGG> findSessions(final K keyFrom,
final K keyTo,
final Instant earliestSessionEndTime,
final Instant latestSessionStartTime) {
throw new UnsupportedOperationException(
"This API is not supported by this implementation of ReadOnlySessionStore.");
}

/**
* Fetch any sessions in the given range of keys and the sessions end is &ge;
* earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
* latest to earliest.
* <p>
* This iterator must be closed after use.
*
* @param keyFrom The first key that could be in the range
* A null value indicates a starting position from the first element in the store <-- new added
* @param keyTo The last key that could be in the range
* A null value indicates that the range ends with the last element in the store. <-- new added
* @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
* iteration ends.
* @param latestSessionStartTime the end timestamp of the latest session to search for, where
* iteration starts.
* @return backward iterator of sessions with the matching keys and aggregated values, from
* latest to earliest session time.
*/
default KeyValueIterator<Windowed<K>, AGG> backwardFindSessions(final K keyFrom,
final K keyTo,
final long earliestSessionEndTime,
final long latestSessionStartTime) {
throw new UnsupportedOperationException(
"This API is not supported by this implementation of ReadOnlySessionStore.");
}

/**
* Fetch any sessions in the given range of keys and the sessions end is &ge;
* earliestSessionEndTime and the sessions start is &le; latestSessionStartTime iterating from
* latest to earliest.
* <p>
* This iterator must be closed after use.
*
* @param keyFrom The first key that could be in the range
* A null value indicates a starting position from the first element in the store <-- new added
* @param keyTo The last key that could be in the range
* A null value indicates that the range ends with the last element in the store. <-- new added
* @param earliestSessionEndTime the end timestamp of the earliest session to search for, where
* iteration ends.
* @param latestSessionStartTime the end timestamp of the latest session to search for, where
* iteration starts.
* @return backward iterator of sessions with the matching keys and aggregated values, from
* latest to earliest session time.
*/
default KeyValueIterator<Windowed<K>, AGG> backwardFindSessions(final K keyFrom,
final K keyTo,
final Instant earliestSessionEndTime,
final Instant latestSessionStartTime) {
throw new UnsupportedOperationException(
"This API is not supported by this implementation of ReadOnlySessionStore.");
}

```

```

/**
 * Retrieve all aggregated sessions for the given range of keys. This iterator must be closed
 * after use.
 * <p>
 * For each key, the iterator guarantees ordering of sessions, starting from the oldest/earliest
 * available session to the newest/latest session.
 *
 * @param keyFrom first key in the range to find aggregated session values for
 * A null value indicates a starting position from the first element in the store <-- new added
 * @param keyTo last key in the range to find aggregated session values for
 * A null value indicates that the range ends with the last element in the store. <-- new added
 * @return KeyValueIterator containing all sessions for the provided key, from oldest to newest
 * session.
 */
KeyValueIterator<Windowed<K>, AGG> fetch(final K keyFrom, final K keyTo);

/**
 * Retrieve all aggregated sessions for the given range of keys. This iterator must be closed
 * after use.
 * <p>
 * For each key, the iterator guarantees ordering of sessions, starting from the newest/latest
 * available session to the oldest/earliest session.
 *
 * @param keyFrom first key in the range to find aggregated session values for
 * A null value indicates a starting position from the first element in the store <-- new added
 * @param keyTo last key in the range to find aggregated session values for
 * A null value indicates that the range ends with the last element in the store. <-- new added
 * @return backward KeyValueIterator containing all sessions for the provided key, from newest
 * to oldest session.
 */
default KeyValueIterator<Windowed<K>, AGG> backwardFetch(final K keyFrom, final K keyTo) {
    throw new UnsupportedOperationException(
        "This API is not supported by this implementation of ReadOnlySessionStore.");
}

```

Proposed Changes

Allow the use of *null* values as a way to represent unbounded ranges. With this change, the range interface signature will remain the same, only the set of accepted values for the parameters *keyFrom* and *keyTo* would change.

Note: we didn't change the time parameters because currently user can always *fetch/findSessions* from/to the earliest/latest time by setting the time as *OL/Long.MAX_VALUE* or *Instant.ofEpochMilli(0)/Instant.ofEpochMilli(Long.MAX_VALUE)* to achieve it. So, we focus on the key parameters here.

For example, in **ReadOnlyWindowStore**, the following code snippet illustrates how unbounded range queries can be issued after those changes:

```

1 //return an iterator for all the records with a key greater (or equal) than "key-x"
2 //within time range from all the existing windows. starting with the smallest record
3 KeyValueIterator<String, String> iter = windowStore.fetch("key-x", null, Instant.ofEpochMilli(start_time), In
4 stant.ofEpochMilli(end_time));
5
6 //return an iterator for all the records with a key less (or equal) than "key-x"
7 //within time range from all the existing windows. starting with the smallest record
8 KeyValueIterator<String, String> iter = windowStore.fetch(null, "key-x", Instant.ofEpochMilli(start_time), In
9 stant.ofEpochMilli(end_time));
10
11 //return an iterator for all the records with a key greater (or equal) than "key-x"
12 //within time range from all the existing windows. starting with the largest record
13 KeyValueIterator<String, String> iter = windowStore.backwardFetch(null, "key-x", Instant.ofEpochMilli
14 (start_time), Instant.ofEpochMilli(end_time));
15
16 //return all the records in the store, equivalent to windowStore.all()
17 KeyValueIterator<String, String> iter = windowStore.fetch(null, null);

```

in **ReadOnlySessionStore**, the following code snippet illustrates how unbounded range queries can be issued after those changes:

```

1 //return an iterator for all the sessions with a key greater (or equal) than "key-x"
2 //within time range from all the existing windows. starting with the smallest record
3 KeyValueIterator<String, String> iter = sessionStore.findSessions("key-x", null, Instant.ofEpochMilli
4 (start_time), Instant.ofEpochMilli(end_time));
5
6 KeyValueIterator<String, String> iter = sessionStore.findSessions("key-x", null, start_time_ms, end_time_ms);
7
8 //return an iterator for all the sessions with a key less (or equal) than "key-x"
9 //within time range from all the existing windows. starting with the smallest record
10 KeyValueIterator<String, String> iter = sessionStore.findSessions(null, "key-x" Instant.ofEpochMilli
11 (start_time), Instant.ofEpochMilli(end_time));
12
13 KeyValueIterator<String, String> iter = sessionStore.findSessions(null, "key-x", start_time_ms, end_time_ms);
14
15 //return an iterator for all the sessions with a key greater (or equal) than "key-x"
16 //within time range from all the existing windows. starting with the largest record
17 KeyValueIterator<String, String> iter = sessionStore.backwardFindSessions(null, "key-x" Instant.ofEpochMilli
18 (start_time), Instant.ofEpochMilli(end_time));

```

Compatibility, Deprecation, and Migration Plan

This change is completely backward compatible.

Rejected Alternatives

Extended from [KIP-763: Range queries with open endpoints](#)

Adding new method methods for unbounded range queries

One alternative to the proposed solution is to add new methods for querying unbounded ranges:

```

1 /**
2  * Get an iterator over all keys less than or equal to the key passed in.
3  * @param to The last key that could be in the range
4  * @return The iterator for this range.
5  * @throws NullPointerException If null is used for to.
6  */
7 default KeyValueIterator<K, V> rangeUntil(K to) { throw new UnsupportedOperationException(); }
8
9 /**
10 * Get an iterator over all keys greater than or equal to the key passed in.
11 * @param from The first key that could be in the range
12 * @return The iterator for this range.
13 * @throws NullPointerException If null is used for from.
14 */
15 default KeyValueIterator<K, V> rangeFrom(K from) { throw new UnsupportedOperationException(); }

```

We rejected this solution because of two reasons. First, because of the nested architecture of the Kafka streams state stores, adding new methods to the state store interface requires modifications to more than a dozen of interfaces which increases the potential for bugs and makes code maintenance and testing more difficult. These issues are further amplified if in the future we want to add variations to those range calls, such as the ability to configure one of the bounds to be inclusive or exclusive.

Modify existing range interface

Yet another solution would be replace the existing range interface with something more powerful that can support different types of range queries with a single method. An example of such an interface is given blow.

```

1 void range(Range<K> range)
2
3 class Range<K> {
4     K from;
5     K to;
6     ...
7     static <K> from(K key) ...
8     static <K> until(K key) ..
9     // etc
10 }

```

While a solution like this would be the cleanest, it would require that we deprecate the old range interface and migrate existing code to the new interface. Compared to the solution proposed in this KIP, which does not require modifications to the interface and no migration is needed, we felt that making bigger changes to the range interface is not necessary at this point.