

KIP-800: Add reason to JoinGroupRequest and LeaveGroupRequest

- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
 - [Protocol](#)
 - [Admin API](#)
 - [Consumer API](#)
- [Proposed Changes](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Rejected Alternatives](#)

Status

Current state: *Accepted*

Discussion thread: [here](#)

JIRA: [here](#)

Please keep the discussion on the mailing list rather than commenting on the wiki (wiki discussions get unwieldy fast).

Motivation

When a consumer leave the consumer group, it logs the reason. However, there is not information on the broker side which makes troubleshooting rebalances triggered by `LeaveGroupRequest` hard. Propagating the reason to the broker via the `LeaveGroupRequest` would be really helpful in such situations.

Public Interfaces

Protocol

A new field called `Reason` is added to each member in the `JoinGroupRequest`.

```

{
  "apiKey": 11,
  "type": "request",
  "listeners": ["zkBroker", "broker"],
  "name": "JoinGroupRequest",
  // Version 1 adds RebalanceTimeoutMs.
  //
  // Version 2 and 3 are the same as version 1.
  //
  // Starting from version 4, the client needs to issue a second request to join group
  //
  // Starting from version 5, we add a new field called groupId to indicate member identity across
  restarts.
  // with assigned id.
  //
  // Version 6 is the first flexible version.
  //
  // Version 7 is the same as version 6.
  //
  // Version 8 adds the Reason field (KIP-800).
  "validVersions": "0-8",
  "flexibleVersions": "6+",
  "fields": [
    { "name": "GroupId", "type": "string", "versions": "0+", "entityType": "groupId",
      "about": "The group identifier." },
    { "name": "SessionTimeoutMs", "type": "int32", "versions": "0+",
      "about": "The coordinator considers the consumer dead if it receives no heartbeat after this timeout in
  milliseconds." },
    // Note: if RebalanceTimeoutMs is not present, SessionTimeoutMs should be
    // used instead. The default of -1 here is just intended as a placeholder.
    { "name": "RebalanceTimeoutMs", "type": "int32", "versions": "1+", "default": "-1", "ignorable": true,
      "about": "The maximum time in milliseconds that the coordinator will wait for each member to rejoin when
  rebalancing the group." },
    { "name": "MemberId", "type": "string", "versions": "0+",
      "about": "The member id assigned by the group coordinator." },
    { "name": "GroupId", "type": "string", "versions": "5+",
      "nullableVersions": "5+", "default": "null",
      "about": "The unique identifier of the consumer instance provided by end user." },
    { "name": "ProtocolType", "type": "string", "versions": "0+",
      "about": "The unique name the for class of protocols implemented by the group we want to join." },
    { "name": "Protocols", "type": "[JoinGroupRequestProtocol]", "versions": "0+",
      "about": "The list of protocols that the member supports.", "fields": [
        { "name": "Name", "type": "string", "versions": "0+", "mapKey": true,
          "about": "The protocol name." },
        { "name": "Metadata", "type": "bytes", "versions": "0+",
          "about": "The protocol metadata." }
      ]
    },
    // New Field //
    { "name": "Reason", "type": "string",
      "versions": "8+", "nullableVersions": "8+", "default": "null",
      "about": "The reason why the member (re-)joins the group." }
  ]
}

```

A new field called Reason is added to each member in the LeaveGroupRequest.

```
{
  "apiKey": 13,
  "type": "request",
  "listeners": ["zkBroker", "broker"],
  "name": "LeaveGroupRequest",
  // Version 1 and 2 are the same as version 0.
  //
  // Version 3 defines batch processing scheme with group.instance.id + member.id for identity
  //
  // Version 4 is the first flexible version.
  //
  // Version 5 adds the Reason field (KIP-800).
  "validVersions": "0-5",
  "flexibleVersions": "4+",
  "fields": [
    { "name": "GroupId", "type": "string", "versions": "0+", "entityType": "groupId",
      "about": "The ID of the group to leave." },
    { "name": "MemberId", "type": "string", "versions": "0-2",
      "about": "The member ID to remove from the group." },
    { "name": "Members", "type": "[]MemberIdentity", "versions": "3+",
      "about": "List of leaving member identities.", "fields": [
        { "name": "MemberId", "type": "string", "versions": "3+",
          "about": "The member ID to remove from the group." },
        { "name": "GroupInstanceId", "type": "string",
          "versions": "3+", "nullableVersions": "3+", "default": "null",
          "about": "The group instance ID to remove from the group." },
        // New Field //
        { "name": "Reason", "type": "string",
          "versions": "5+", "nullableVersions": "5+", "default": "null",
          "about": "The reason why the member left the group." }
      ]
    }
  ]
}
```

Admin API

The reason is added as an optional attribute to `RemoveMembersFromConsumerGroupOptions` class which is used by the `Admin#removeMembersFromConsumerGroup` method.

```
@InterfaceStability.Evolving
public class RemoveMembersFromConsumerGroupOptions extends
AbstractOptions<RemoveMembersFromConsumerGroupOptions> {
  /**
   * Sets an optional reason.
   */
  void reason(final String reason);

  /**
   * Returns the optional reason or an empty String if not defined.
   */
  String reason();
}
```

Consumer API

The reason is added as an argument to `enforceRebalance` method.

```
/**
 * @see KafkaConsumer#enforceRebalance()
 */
void enforceRebalance(final String reason);
```

Proposed Changes

The KIP adds a per-member field called `Reason` to the `LeaveGroupRequest`. The field will be populated with the leaving reason by the Consumer when it wants to leave the group. The consumer already logs the reason so we will reuse the same message. The field will be populated by the AdminClient as well when `removeMembersFromConsumerGroup` is called. In this case, the reason will be "member was removed by an admin" by default or the provided reason.

The KIP also adds a field called `Reason` to the `JoinGroupRequest`. The field will be populated with the reason why the joining reason by the Consumer. This will be really helpful to understand why a rebalance is triggered. The field will be populated by the Consumer as well when `enforceRebalance` is called. In this case, the reason will be "rebalance enforced by user" by default or the provided reason.

Compatibility, Deprecation, and Migration Plan

The new field defaults to `null` and is optional so compatibility is not a concern.

Rejected Alternatives

The alternative would be to do nothing and to keep relying on client logs in this case.