

KIP-700: Add Describe Cluster API

- [Status](#)
- [Motivation](#)
- [Public Interfaces](#)
 - [Describe Cluster API](#)
 - [DescribeClusterRequest](#)
 - [DescribeClusterResponse](#)
 - [Metadata API](#)
- [Proposed Changes](#)
- [Compatibility, Deprecation, and Migration Plan](#)
- [Rejected Alternatives](#)
- [Future Work](#)

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

The AdminClient introduced by KIP-117 has a method which describe the cluster. It provides:

- the cluster id
- the controller id
- detailed information about every broker in the cluster
- the cluster authorized operations (KIP-430)

Historically, the AdminClient has been using the Metadata API to get this metadata about the cluster. We did so because it was convenient as all the required metadata where there. However, the Metadata API was designed around the needs of producers and consumers, which need to periodically refresh detailed information about the cluster and each topics they are producing to or consuming from. Therefore extending the Metadata API to service information only used by the AdminClient does not align well with its primary purpose. We already did this in the past with KIP-430 which added the authorized operations to the response. The authorized operations are not used at all by producers nor consumers.

This KIP proposes to introduce a new Describe Cluster API which will service the AdminClient. This will allow us to service more information in the future without having to worry about hijacking the Metadata API.

Public Interfaces

Describe Cluster API

The initial version of the Describe Cluster API is a subset of the Metadata API. It contains all the fields currently used by the AdminClient.

DescribeClusterRequest

```
{
  "apiKey": "TBD",
  "type": "request",
  "name": "DescribeClusterRequest",
  "validVersions": "0",
  "flexibleVersions": "0+",
  "fields": [
    { "name": "IncludeClusterAuthorizedOperations", "type": "bool", "versions": "0+",
      "about": "Whether to include cluster authorized operations." }
  ]
}
```

DescribeClusterResponse

```

{
  "apiKey": "TBD",
  "type": "response",
  "name": "DescribeClusterResponse",
  "validVersions": "0",
  "flexibleVersions": "0+",
  "fields": [
    { "name": "ThrottleTimeMs", "type": "int32", "versions": "0+",
      "about": "The duration in milliseconds for which the request was throttled due to a quota violation, or zero if the request did not violate any quota." },
    { "name": "ClusterId", "type": "string", "versions": "0+",
      "about": "The cluster ID that responding broker belongs to." },
    { "name": "ControllerId", "type": "int32", "versions": "0+", "default": "-1", "entityType": "brokerId",
      "about": "The ID of the controller broker." },
    { "name": "Brokers", "type": "[]DescribeClusterBroker", "versions": "0+",
      "about": "Each broker in the response.", "fields": [
        { "name": "NodeId", "type": "int32", "versions": "0+", "mapKey": true, "entityType": "brokerId",
          "about": "The broker ID." },
        { "name": "Host", "type": "string", "versions": "0+",
          "about": "The broker hostname." },
        { "name": "Port", "type": "int32", "versions": "0+",
          "about": "The broker port." },
        { "name": "Rack", "type": "string", "versions": "0+", "nullableVersions": "0+", "default": "null",
          "about": "The rack of the broker, or null if it has not been assigned to a rack." }
      ]
    },
    { "name": "ClusterAuthorizedOperations", "type": "int32", "versions": "0+", "default": "-2147483648",
      "about": "32-bit bitfield to represent authorized operations for this cluster." }
  ]
}

```

Metadata API

The version of the MetadataRequest and the MetadataResponse will be bumped to respectively deprecate the IncludeClusterAuthorizedOperations field and the ClusterAuthorizedOperations field.

Proposed Changes

The Describe Cluster API will be implemented by the brokers. On the client side, Admin#describeCluster should try to use it if available. Admin#describeCluster remains unchanged.

Compatibility, Deprecation, and Migration Plan

When communicating with older brokers that do not implement the DescribeCluster API, the AdminClient will fail back to using the Metadata API.

Rejected Alternatives

The alternative is the keep using the Metadata API to describe the cluster.

Future Work

- With the removal of Zookeeper, administrators will lose their ability to retrieve some metadata about the clusters from Zookeeper. For instance, the JMX port or the entire list of end-points for each broker. The new DescribeCluster API will enable us to provide more metadata about the cluster and about the brokers in the future.
- More generally, the new Describe Cluster API will enable us to provide more metadata about the cluster such as the status of each brokers, their respective version, their respective epochs, or various other things.