

KIP-142: Add ListTopicsRequest to efficiently list all the topics in a cluster

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

Status

Current state: Under Discussion

Discussion thread: <https://www.mail-archive.com/dev@kafka.apache.org/msg70723.html>

JIRA: *TDB*

Motivation

The **AdminClient** introduced by KIP-117 has a method which lists all the topics in the cluster. However, we currently don't have an efficient way to implement this. The **AdminClient** can make a **MetadataRequest** which will return information about all the topics in the cluster. However, the request will also return:

- detailed information about every broker in the cluster
- the cluster id
- the controller id
- detailed information about each topic.

In a cluster with thousands of topics and dozens of brokers, the **MetadataResponse** will be quite large. Most of this information is not needed for **listTopics**.

MetadataRequest was designed around the needs of consumers and producers, which need to periodically refresh detailed information about each topic they are producing to or consuming from. In contrast, the **AdminClient** often just needs to know about the presence or absence of a topic. It would be better to have a method which just listed the names of the topics which exist without adding all the other information.

Public Interfaces

We will add a new request for listing the topics in a cluster.

```
ListTopicsRequest (Version 0) => list_internal
    list_internal => BOOLEAN

ListTopicsResponse (Version 0) => error_code [topic_listing]
    topic_listing => name internal
    name => STRING
    internal => BOOLEAN
```

When creating a **ListTopicsRequest**, we can decide whether or not we want to see internal topics in the results. If `list_internal` is false, internal topics will be excluded from the results.

The result of **ListTopicsResponse** is only the names of the topics and whether each one is internal. We do not return information about partitions, leaders, replicas, ISRs, or the list of brokers.

ListTopicsResponse will not return information about topics that the current user does not have permission to describe.

Proposed Changes

ListTopicsRequest will be implemented by the brokers. On the client side, **AdminClient#listTopics** should try to use the new request if available.

Compatibility, Deprecation, and Migration Plan

When communicating with older brokers that do not implement **ListTopicsRequest**, the **AdminClient** can fall back to making a **MetadataRequest**.

Rejected Alternatives

Rather than creating a new request type, we could add flags to **MetadataRequest** directing it to leave out certain information in its response. This would result in a smaller **MetadataResponse**, but not quite as small as the **ListTopicsResponse**. **MetadataResponse** has fields that take up space even if they're empty, such as arrays of partitions, leaders, ISRs, and so forth. It would also add complexity to a request which is already somewhat complex. It seems better to simply create a new request type.