

# MEP 31 -- Introducing Leading/Alternative Mode for Milvus Cluster

## Motivation

There are three ways to deploy Milvus for production use: standalone version, clustered version and cloud version.

In production, high availability (HA) is one of the most important things that people care about. In the cloud version, everything is fully managed and you do not need to worry about it. But there are different stories when it comes to local standalone deployment and clustered deployment.

We here introduce our current HA plan for Milvus standalone and clustered deployment. We will then focus on the HA problem for Milvus clustered deployment and introduce a master-slave architecture to achieve high availability.

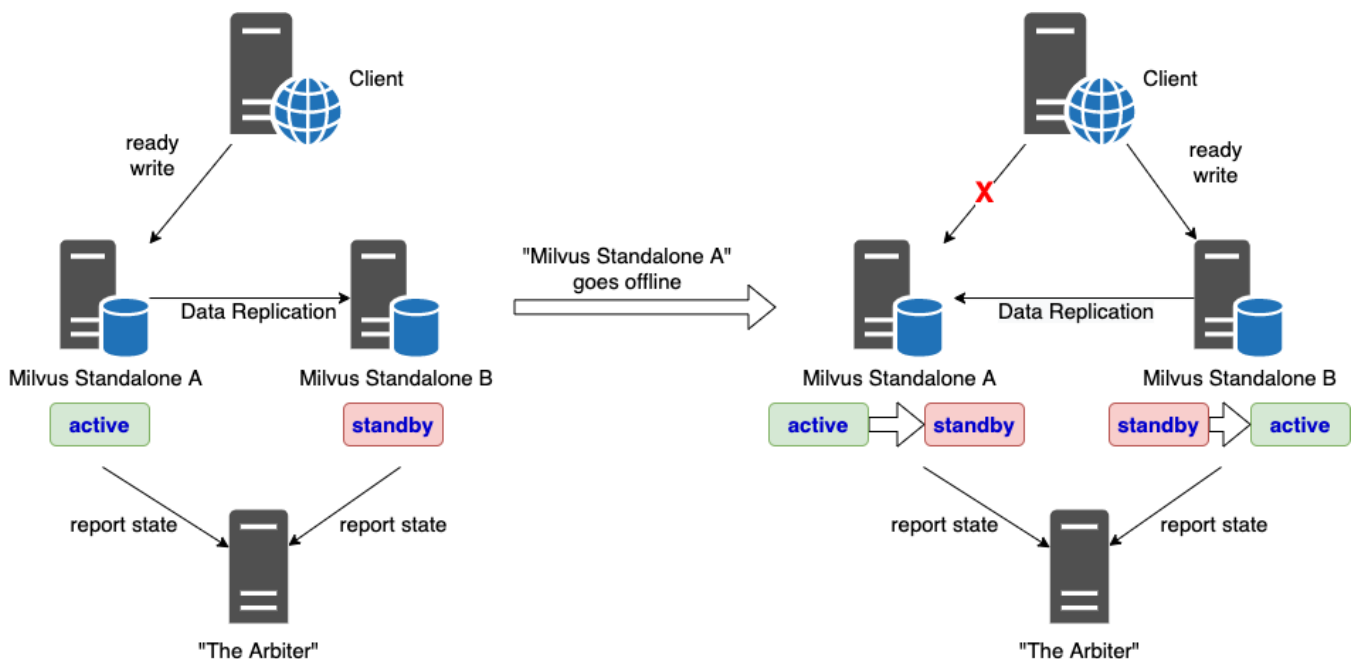
## For Standalone HA: Active/Standby

In standalone mode, we will use the rather straightforward HA solution, i.e. active-standby mode.

In addition to a normal Milvus standalone instance, or the "active" instance, we also start another identical instance, or the "standby" instance.

Here's how it works:

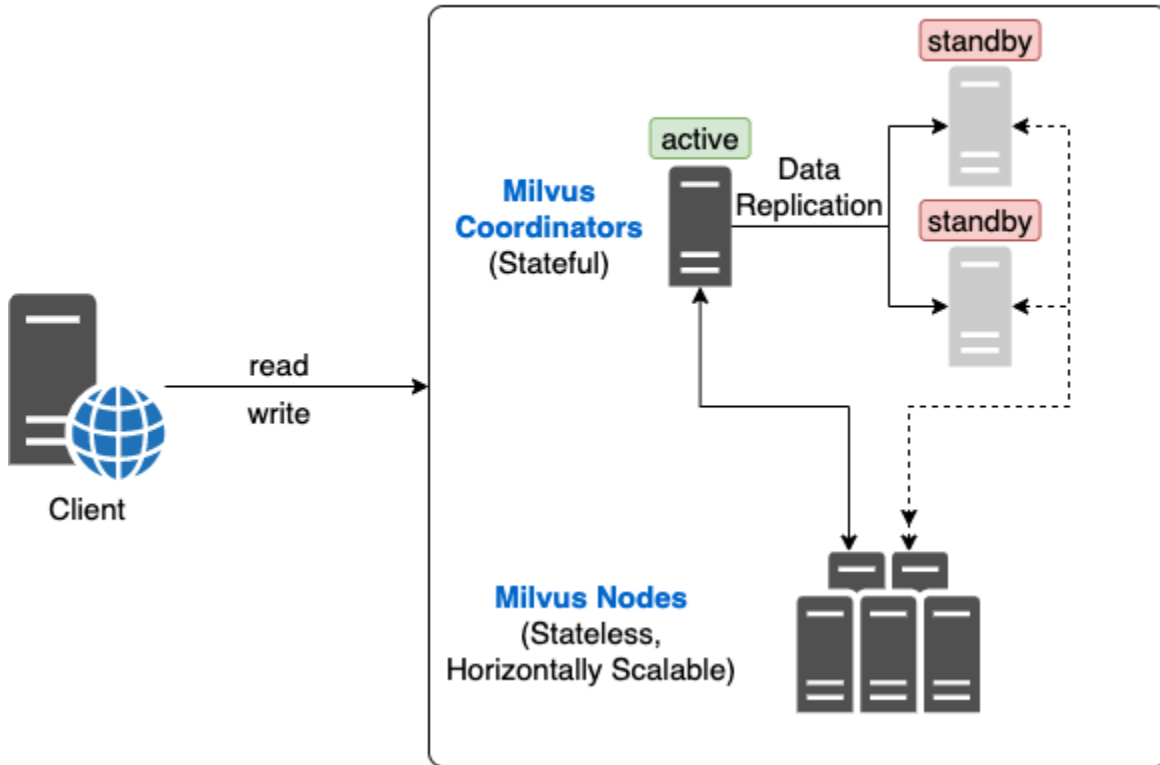
- There could be multiple standby instances but only one active instance at a time.
- The standby instance merely serves as a backup. It does not accept any traffic from the client in normal cases. The active instance will continuously replicate all data to the standby instance.
- Both active and standby instances will keep reporting their status to a third-party service (or the "arbiter"). It is up to this arbiter that decides the liveness of all instances. (Note that this arbiter itself needs to have HA in the first place)
- When the active instance, for example, instance A in the figure below, goes offline, the arbiter will immediately sense this and notifies the standby instance to promote itself as the new active instance.
- When the previous active instance resurrects from whatever situation, it demotes itself to a standby instance and starts reporting to the arbiter with its new identity.



## Clustered Milvus(Single Instance) HA: Coordinator Active/Standby

Milvus can be deployed as a single Milvus cluster. To reach high availability in this case, we'd introduce the Active/Standby mode to all Coordinators (i.e. RootCoord, DataCoord, QueryCoord, IndexCoord). Please refer to [MEP 30 -- Support Coordinators Primary-backup Mechanism](#) to see how it works.

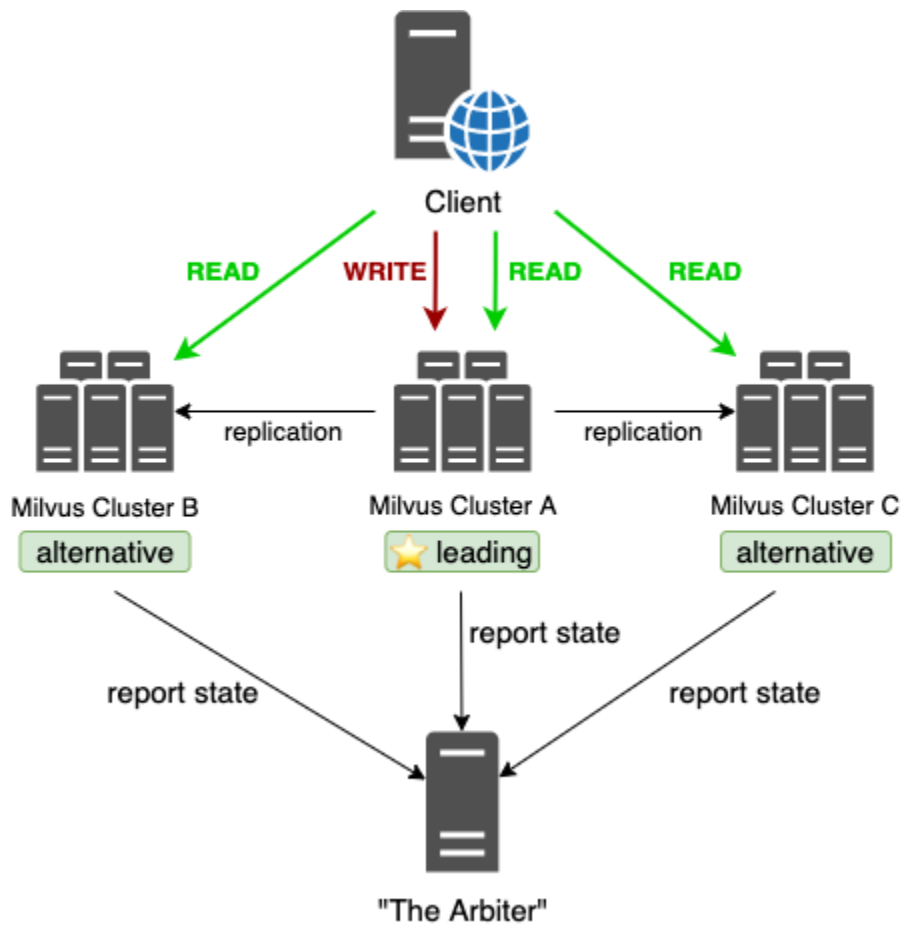
## Single Milvus Cluster



## Clustered Milvus(Multi Instances) HA: Leading/Alternative

The leading/alternative mode is similar to the active/standby mode, with the following differences:

- The client is now able to sense both leading instances and alternative instances.
- The alternative instances can serve read requests from clients, but cannot serve write requests. This can fully utilize the hardware resources of machines for alternative instances.
- Since the alternative instances also serve as read services, all data in alternative instances should be ready at all time and should always be consistent with the leading instance, which is far more strict than active/standby mode.



Data exporting is highly required to fully accomplish the leading/alternative mode for clustered Milvus, where we need to support exporting in-real time all historical data, live data and DDL data. Once we have this, we will have alternative instance subscribe to the leading instance for all data. Expecting more details to be enclosed.