

SyncService

- [General](#)
 - [Architecture](#)
 - [Server](#)
 - [Client](#)
 - [Configuring](#)

General

Since Apache Hama 0.4.0 we provide the functionality to implement your own Synchronization Service.

Apache Hama already ships with a default implementation with Apache Zookeeper, however there may be other synchronization frameworks that can be used for it.

Architecture

There is a `SyncServer` and a `SyncClient`, both after the name are used for Apache Hama to do the barrier sync.

However, there are structural differences, for example in "normal" Hama the server is not launched by the framework. It has to be launched by the user. In Hama-YARN the server is part of the [ApplicationMaster](#) and will be started within.

The client is consistent across all the different implementations, it should do all the neat client stuff like registering a job/task and handling the sync by connecting to the synchronization server etc.

Server

The server consists of methods in the interface called `org.apache.hama.bsp.sync.SyncServer`. Consult the java docs for additional information.

The main implementation of Apache Hama is `org.apache.hama.bsp.sync.ZooKeeperSyncServerImpl` which abstracts the use of Zookeeper.

Client

The client is described in its interface `org.apache.hama.bsp.sync.SyncClient`. Note that the client takes care in registering the job as well as each task. It is also responsible to keep the hostnames of all tasks involved and of course the sync itself.

The main implementation of Apache Hama is `org.apache.hama.bsp.sync.ZooKeeperSyncClientImpl` which abstracts the use of Zookeeper.

Configuring

Both implementations can be set by the configuration values:

```
hama.sync.server.class
hama.sync.client.class
```

Make sure these two configurations are consistent on every host in every configuration. And make sure that all these classes have a public default constructor since they will be instantiated via reflection.

If you decide to implement your own synchronization server, please contact us at the dev-mailing list.