

BookKeeper Administrator's Guide

Setup Guide

by

Table of contents

1 Deployment.....	2
1.1 System requirements.....	2
1.2 Running bookies.....	2
1.3 ZooKeeper Metadata.....	3

1 Deployment

This section contains information about deploying BookKeeper and covers these topics:

- [System requirements](#)
- [Running bookies](#)
- [ZooKeeper Metadata](#)

The first section tells you how many machines you need. The second explains how to bootstrap bookies (BookKeeper storage servers). The third section explains how we use ZooKeeper and our requirements with respect to ZooKeeper.

1.1 System requirements

A typical BookKeeper installation comprises a set of bookies and a set of ZooKeeper replicas. The exact number of bookies depends on the quorum mode, desired throughput, and number of clients using this installation simultaneously. The minimum number of bookies is three for self-verifying (stores a message authentication code along with each entry) and four for generic (does not store a message authentication code with each entry), and there is no upper limit on the number of bookies. Increasing the number of bookies, in fact, enables higher throughput.

For performance, we require each server to have at least two disks. It is possible to run a bookie with a single disk, but performance will be significantly lower in this case. Of course, it works with one disk, but performance is significantly lower.

For ZooKeeper, there is no constraint with respect to the number of replicas. Having a single machine running ZooKeeper in standalone mode is sufficient for BookKeeper. For resilience purposes, it might be a good idea to run ZooKeeper in quorum mode with multiple servers. Please refer to the ZooKeeper documentation for detail on how to configure ZooKeeper with multiple replicas

1.2 Running bookies

To run a bookie, we execute the following command:

```
java -cp ../zookeeper-<version>-bookkeeper.jar:./
zookeeper-<version>.jar\ :../log4j/apache-log4j-1.2.15/
log4j-1.2.15.jar -Dlog4j.configuration=log4j.properties\
org.apache.bookkeeper.proto.BookieServer 3181 127.0.0.1:2181 /
path_to_log_device/\ /path_to_ledger_device/
```

The parameters are:

- Port number that the bookie listens on;
- Comma separated list of ZooKeeper servers with a hostname:port format;

- Path for Log Device (stores bookie write-ahead log);
- Path for Ledger Device (stores ledger entries);

Ideally, `/path_to_log_device/` and `/path_to_ledger_device/` are each in a different device.

1.3 ZooKeeper Metadata

For BookKeeper, we require a ZooKeeper installation to store metadata, and to pass the list of ZooKeeper servers as parameter to the constructor of the BookKeeper class (`org.apache.bookkeeper.client.BookKeeper`). To setup ZooKeeper, please check the [ZooKeeper documentation](#).