

# Apache News Online: Apache Hadoop

## Table of contents

13 February 2009 - Apache ZooKeeper version 3.1.0 released.....	2
24 January 2008 - Apache Hadoop Now TLP.....	2

## 13 February 2009 - Apache ZooKeeper version 3.1.0 released

[The Apache ZooKeeper team](#) is proud to announce Apache ZooKeeper version 3.1.0. ZooKeeper is a high-performance coordination service for distributed applications. It exposes common services - such as naming, configuration management, synchronization, and group services - in a simple interface so you don't have to write them from scratch. You can use it off-the-shelf to implement consensus, group management, leader election, and presence protocols. And you can build on it for your own, specific needs. Key features of the 3.1.0 release: \* Quota support \* BookKeeper - a system to reliably log streams of records \* JMX for server management \* many fixes, improvements, improved documentation, etc... A bit about BookKeeper: a system to reliably log streams of records. In BookKeeper, servers are "bookies", log streams are "ledgers", and each unit of a log (aka record) is a "ledger entry". BookKeeper is designed to be reliable; bookies, the servers that store ledgers can be byzantine, which means that some subset of the bookies can fail, corrupt data, discard data, but as long as there are enough correctly behaving servers the service as a whole behaves correctly; the meta data for BookKeeper is stored in ZooKeeper. For ZooKeeper release details and downloads, visit: <http://hadoop.apache.org/zookeeper/releases.html> ZooKeeper 3.1.0 Release Notes are at: <http://hadoop.apache.org/zookeeper/docs/r3.1.0/releasenotes.html> Regards, The ZooKeeper Team

### Note:

Apache Hadoop -- Posted by [Tetsuya Kitahata](#) at 2009-02-14T03:14:35

## 24 January 2008 - Apache Hadoop Now TLP

Apache Hadoop is now Top Level Project (TLP) in the Apache Software Foundation with its new website : <http://hadoop.apache.org/> Apache Hadoop is a software platform that lets one easily write and run applications that process vast amounts of data - it had been developed under the [Apache Lucene Project](#).

### Note:

Apache Hadoop -- Posted by [Tetsuya Kitahata](#) at 2008-01-24T02:41:05