

PoweredBy

Applications and organizations using ZooKeeper include (alphabetically):

- company w/ link if available use

Free Software Projects

- [AdroitLogic UltraESB](#) - Uses ZooKeeper to implement node coordination, in clustering support. This allows the management of the complete cluster, or any specific node - from any other node connected via JMX. A Cluster wide command framework developed on top of the ZooKeeper coordination allows commands that fail on some nodes to be retried etc. We also support the automated graceful round-robin-restart of a complete cluster of nodes using the same framework. More information on how we use ZooKeeper can be found at the [clustering concepts](#) and [clustering guide](#) articles.
- [Akka](#) Akka is the platform for the next generation event-driven, scalable and fault-tolerant architectures on the JVM. Or: Akka is a toolkit and runtime for building highly concurrent, distributed, and fault tolerant event-driven applications on the JVM.
- [Eclipse Communication Framework](#) - The Eclipse ECF project provides an implementation of its Abstract Discovery services using Zookeeper. ECF itself is used in many projects providing base functionality for communication, all based on OSGi.
- [Eclipse Gyrex](#) - The Eclipse Gyrex project provides a platform for building your own Java OSGi based clouds. ZooKeeper is used as the core cloud component for node membership and management, coordination of jobs executing among workers, a lock service and a simple queue service and a lot more.
- [GoldenOrb](#) - massive-scale Graph analysis
- [Juju](#) Service deployment and orchestration framework, formerly called Ensemble. [How is ZooKeeper used](#) by Juju.
- [Katta](#) - Katta serves distributed Lucene indexes in a grid environment. Zookeeper is used for node, master and index management in the grid.
- [KeptCollections](#) - KeptCollections is a library of drop-in replacements for the data structures in the Java Collections framework. KeptCollections uses Apache ZooKeeper as a backing store, thus making its data structures distributed and scalable.
- [Mesos](#) Mesos is a cluster management platform that provides resource sharing and isolation across distributed applications, or frameworks. For example, Hadoop MapReduce, MPI, Hypertable, and Spark
- [Neo4j](#) - Neo4j is a Graph Database. It's a disk based, ACID compliant transactional storage engine for big graphs and fast graph traversals, using external indices like Lucene/Solr for global searches. We use ZooKeeper in the Neo4j High Availability components for write-master election, read slave coordination and other cool stuff. ZooKeeper is a great and focused project - we like!
- [Norbert](#) Partitioned routing and cluster management.
- [Talend ESB](#) - Talend ESB is a versatile and flexible, enterprise service bus. It uses ZooKeeper as endpoint repository of both REST and SOAP Web services. By using ZooKeeper Talend ESB is able to provide failover and load balancing capabilities in a very light-weight manner. .
- [redis_failover](#) Redis Failover is a ZooKeeper-based automatic master/slave failover solution for Ruby.

Apache Projects

- [Apache Accumulo](#) - Accumulo is a distributed key/value store that provides expressive, cell-level access labels.
- [Apache BookKeeper](#) (ZooKeeper subproject) distributed logging service
- [Apache CXF DOSGi](#) The Distributed OSGi implementation at Apache CXF uses ZooKeeper for its Discovery functionality. See here for more information: <http://cxf.apache.org/dosgi-discovery-demo-page.html>
- planning: [Apache Flume](#) Service for efficiently collecting, aggregating, and moving large amounts of log data. The next major version is planned to use ZooKeeper: [Jira issue](#)
- [Apache Hadoop MapReduce](#) The next generation of Hadoop MapReduce (colled "Yarn") uses ZooKeeper. [Jira issue](#), [Design paper](#)
- [Apache HBase](#) - HBase is the Hadoop database. Its an open-source, distributed, column-oriented store modeled after the Google paper, Bigtable: A Distributed Storage System for Structured Data by Chang et al. We use ZooKeeper for master election, server lease management, bootstrapping, and coordination between servers.
- [Apache Hedwig](#) Hedwig is a publish-subscribe system designed to carry large amounts of data across the internet in a guaranteed-delivery fashion from those who produce it (publishers) to those who are interested in it (subscribers).
- [Apache Kafka \(incubating\)](#) Kafka is a distributed publish/subscribe messaging system
- [Apache S4 \(incubating\)](#) S4 is a general-purpose, distributed, scalable, partially fault-tolerant, pluggable platform that allows programmers to easily develop applications for processing continuous unbounded streams of data.
- [Apache Solr](#) - In the "Cloud" edition (v4.x and up) of enterprise search engine Apache Solr, ZooKeeper is used for configuration, leader election and more.

Companies

- [AGETO](#) - The AGETO RnD team uses ZooKeeper in a variety of internal as well as external consulting projects.
- [Benipal Technologies](#) - ZooKeeper is used for internal application development with Solr and Hadoop with Hbase.
- [Box](#) - Box uses ZooKeeper for service discovery, service coordination, Solr and Hadoop support, etc.
- [Deepdyve](#) - We do search for research and provide access to high quality content using advanced search technologies Zookeeper is used to manage server state, control index deployment and a myriad other tasks.
- [Idium Portal](#) - Iidium Portal is a hosted web-publishing system delivered by Norwegian company, Iidium AS. ZooKeeper is used for cluster messaging, service bootstrapping, and service coordination.
- [Makara](#)
 - Using ZooKeeper on 2-node cluster on VMware workstation, Amazon EC2, Zen
 - Using zkpython
 - Looking into expanding into 100 node cluster
- [Midokura](#) - We do virtualized networking for the cloud computing era. We use ZooKeeper for various aspects of our distributed control plane.
- [Rackspace](#) - The Email & Apps team uses ZooKeeper to coordinate sharding and responsibility changes in a distributed e-mail client that pulls and indexes data for search. ZooKeeper also provides distributed locking for connections to prevent a cluster from overwhelming servers.
- [Sematext](#) - Uses ZooKeeper in [SPM](#) (which includes ZooKeeper monitoring component, too!), [Search Analytics](#), and [Logsene](#).
- [Tubemogul](#) - Uses ZooKeeper for leader election, configuration management, locking, group membership

- [Vast.com](#) - Used internally as a part of sharding services, distributed synchronization of data/index updates, configuration management and failover support
- [Wealthfront](#)
 - Wealthfront uses ZooKeeper for service discovery, leader election and distributed locking among its many backend services. ZK is an essential part of Wealthfront's [continuous deployment](#) infrastructure.
- [Yahoo!](#) - ZooKeeper is used for a myriad of services inside Yahoo! for doing leader election, configuration management, sharding, locking, group membership etc.
- [Zynga](#) - ZooKeeper at Zynga is used for a variety of services including configuration management, leader election, sharding and more.... A brief introduction to our usage at Zynga can be seen at the following blog post. <http://code.zynga.com/2011/08/updating-thousands-of-configuration-files-in-under-a-second/>

Transitive Closure

- powered by Kafka
- powered by HBase