

ZooKeeper 3.4.0 Release Notes

by

Table of contents

1 Changes Since ZooKeeper 3.3.0.....	2
--------------------------------------	---

These release notes include new developer and user facing incompatibilities, features, and major improvements.

- [Changes](#)

1 Changes Since ZooKeeper 3.3.0

Issue	Notes
Sub-Tasks	
ZOOKEEPER-1239	add logging/stats to identify fsync stalls.
ZOOKEEPER-1208	Ephemeral node not removed after the client session is long gone.
ZOOKEEPER-784	server-side functionality for read-only mode. This is not thoroughly tested. Avoid using it in production. This is also at risk of being removed from the feature set later.
ZOOKEEPER-798	Fixup loggraph for FLE changes
ZOOKEEPER-839	deleteRecursive does not belong to the other methods
ZOOKEEPER-908	Remove code duplication and inconsistent naming in ClientCnxn.Packet creation
ZOOKEEPER-909	Extract NIO specific code from ClientCnxn
ZOOKEEPER-966	Client side for multi
ZOOKEEPER-967	Server side decoding and function dispatch
ZOOKEEPER-968	Database multi-update
ZOOKEEPER-1042	Generate zookeeper test jar for maven installation
ZOOKEEPER-1081	modify leader/follower code to correctly deal with new leader
ZOOKEEPER-1082	modify leader election to correctly take into account current epoch
ZOOKEEPER-1150	fix for this patch to compile on windows...
ZOOKEEPER-1160	test timeouts are too small
ZOOKEEPER-1201	Clean SaslServerCallbackHandler.java
Bug Fixes	

Issue	Notes
ZOOKEEPER-1268	problems with read only mode, intermittent test failures and ERRORS in the log.
ZOOKEEPER-1271	testEarlyLeaderAbandonment failing on solaris - clients not retrying connection.
ZOOKEEPER-1192	Leader.waitForEpochAck() checks waitingForNewEpoch instead of checking electionFinished.
ZOOKEEPER-1246	Dead code in PrepRequestProcessor catch Exception block.
ZOOKEEPER-1264	FollowerResyncConcurrencyTest failing intermittently.
ZOOKEEPER-1270	testEarlyLeaderAbandonment failing intermittently, quorum formed, no serving.
ZOOKEEPER-1291	AcceptedEpoch not updated at leader before it proposes the epoch to followers.
ZOOKEEPER-1282	Learner.java not following Zab 1.0 protocol - setCurrentEpoch should be done upon receipt of NEWLEADER (before acking it) and not upon receipt of UPTODATE.
ZOOKEEPER-335	zookeeper servers should commit the new leader txn to their logs.
ZOOKEEPER-418	Need nifty zookeeper browser
ZOOKEEPER-603	zkpython should do a better job of freeing memory under error conditions
ZOOKEEPER-662	Too many CLOSE_WAIT socket state on a server
ZOOKEEPER-690	AsyncTestHammer test fails on hudson.
ZOOKEEPER-719	Add throttling to BookKeeper client
ZOOKEEPER-720	Use zookeeper-{version}-sources.jar instead of zookeeper-{version}-src.jar to publish sources in the Maven repository
ZOOKEEPER-722	zkServer.sh uses sh's builtin echo on BSD, behaves incorrectly.

Issue	Notes
ZOOKEEPER-731	Zookeeper#delete , #create - async versions miss a verb in the javadoc
ZOOKEEPER-734	QuorumPeerTestBase.java and ZooKeeperServerMainTest.java do not handle windows path correctly
ZOOKEEPER-735	cppunit test testipv6 assumes that the machine is ipv6 enabled.
ZOOKEEPER-737	some 4 letter words may fail with netcat (nc)
ZOOKEEPER-738	zookeeper.jute.h fails to compile with -pedantic
ZOOKEEPER-741	root level create on REST proxy fails
ZOOKEEPER-742	Deallocatng None on writes
ZOOKEEPER-746	learner outputs session id to log in dec (should be hex)
ZOOKEEPER-749	OSGi metadata not included in binary only jar
ZOOKEEPER-750	move maven artifacts into "dist-maven" subdir of the release (package target)
ZOOKEEPER-758	zkpython segfaults on invalid acl with missing key
ZOOKEEPER-763	Deadlock on close w/ zkpython / c client
ZOOKEEPER-764	Observer elected leader due to inconsistent voting view
ZOOKEEPER-766	forrest recipes docs don't mention the lock/queue recipe implementations available in the release
ZOOKEEPER-769	Leader can treat observers as quorum members
ZOOKEEPER-772	zkpython segfaults when watcher from async get children is invoked.
ZOOKEEPER-774	Recipes tests are slightly outdated: they do not compile against JUnit 4.8
ZOOKEEPER-777	setting acl on a non existant node should return no node error
ZOOKEEPER-782	Incorrect C API documentation for Watches

Issue	Notes
ZOOKEEPER-783	committedLog in ZKDatabase is not properly synchronized
ZOOKEEPER-787	groupId in deployed pom is wrong
ZOOKEEPER-790	Last processed zxid set prematurely while establishing leadership
ZOOKEEPER-792	zkpython memory leak
ZOOKEEPER-794	Callbacks are not invoked when the client is closed
ZOOKEEPER-795	eventThread isn't shutdown after a connection "session expired" event coming
ZOOKEEPER-796	zkServer.sh should support an external PIDFILE variable
ZOOKEEPER-800	zoo_add_auth returns ZOK if zookeeper handle is in ZOO_CLOSED_STATE
ZOOKEEPER-804	c unit tests failing due to "assertion cptr failed"
ZOOKEEPER-813	maven install is broken due to incorrect organisation
ZOOKEEPER-814	monitoring scripts are missing apache license headers
ZOOKEEPER-820	update c unit tests to ensure "zombie" java server processes don't cause failure
ZOOKEEPER-822	Leader election taking a long time to complete
ZOOKEEPER-831	BookKeeper: Throttling improved for reads
ZOOKEEPER-844	handle auth failure in java client
ZOOKEEPER-846	zookeeper client doesn't shut down cleanly on the close call
ZOOKEEPER-854	BookKeeper does not compile due to changes in the ZooKeeper code
ZOOKEEPER-855	clientPortBindAddress should be clientPortAddress
ZOOKEEPER-861	Missing the test SSL certificate used for running junit tests.
ZOOKEEPER-867	ClientTest is failing on hudson - fd cleanup
ZOOKEEPER-870	Zookeeper trunk build broken.

Issue	Notes
ZOOKEEPER-874	FileTxnSnapLog.restore does not call listener
ZOOKEEPER-880	QuorumCnxManager\$SendWorker grows without bounds
ZOOKEEPER-881	ZooKeeperServer.loadData loads database twice
ZOOKEEPER-882	Startup loads last transaction from snapshot
ZOOKEEPER-884	Remove LedgerSequence references from BookKeeper documentation and comments in tests
ZOOKEEPER-888	c-client / zkpython: Double free corruption on node watcher
ZOOKEEPER-893	ZooKeeper high cpu usage when invalid requests
ZOOKEEPER-897	C Client seg faults during close
ZOOKEEPER-898	C Client might not cleanup correctly during close
ZOOKEEPER-902	Fix findbug issue in trunk "Malicious code vulnerability"
ZOOKEEPER-904	super digest is not actually acting as a full superuser
ZOOKEEPER-913	Version parser fails to parse "3.3.2-dev" from build.xml.
ZOOKEEPER-917	Leader election selected incorrect leader
ZOOKEEPER-919	Ephemeral nodes remains in one of ensemble after deliberate SIGKILL
ZOOKEEPER-921	zkPython incorrectly checks for existence of required ACL elements
ZOOKEEPER-937	test -e not available on solaris /bin/sh
ZOOKEEPER-957	zkCleanup.sh doesn't do anything
ZOOKEEPER-958	Flag to turn off autoconsume in hedwig c++ client
ZOOKEEPER-961	Watch recovery after disconnection when connection string contains a prefix
ZOOKEEPER-962	leader/follower coherence issue when follower is receiving a DIFF
ZOOKEEPER-963	Make Forrest work with JDK6

Issue	Notes
ZOOKEEPER-965	Need a multi-update command to allow multiple znodes to be updated safely
ZOOKEEPER-975	new peer goes in LEADING state even if ensemble is online
ZOOKEEPER-976	ZooKeeper startup script doesn't use JAVA_HOME
ZOOKEEPER-981	Hang in zookeeper_close() in the multi-threaded C client
ZOOKEEPER-983	running zkServer.sh start remotely using ssh hangs
ZOOKEEPER-985	Test BookieRecoveryTest fails on trunk.
ZOOKEEPER-1006	QuorumPeer "Address already in use" -- regression in 3.3.3
ZOOKEEPER-1007	iarchive leak in C client
ZOOKEEPER-1013	zkServer.sh usage message should mention all startup options
ZOOKEEPER-1027	chroot not transparent in zoo_create()
ZOOKEEPER-1028	In python bindings, zookeeper.set2() should return a stat dict but instead returns None
ZOOKEEPER-1033	c client should install includes into INCDIR/zookeeper, not INCDIR/c-client-src
ZOOKEEPER-1034	perl bindings should automatically find the zookeeper c-client headers
ZOOKEEPER-1046	Creating a new sequential node results in a ZNODEEXISTS error
ZOOKEEPER-1049	Session expire/close flooding renders heartbeats to delay significantly
ZOOKEEPER-1051	SIGPIPE in Zookeeper 0.3.* when send'ing after cluster disconnection
ZOOKEEPER-1052	Findbugs warning in QuorumPeer.ResponderThread.run()
ZOOKEEPER-1055	check for duplicate ACLs in addACL() and create()
ZOOKEEPER-1058	fix typo in opToString for getData

Issue	Notes
ZOOKEEPER-1059	stat command issues on non-existing node causes NPE
ZOOKEEPER-1060	QuorumPeer takes a long time to shutdown
ZOOKEEPER-1061	Zookeeper stop fails if start called twice
ZOOKEEPER-1063	Dubious synchronization in Zookeeper and ClientCnxnSocketNIO classes
ZOOKEEPER-1068	Documentation and default config suggest incorrect location for Zookeeper state
ZOOKEEPER-1069	Calling shutdown() on a QuorumPeer too quickly can lead to a corrupt log
ZOOKEEPER-1073	address a documentation issue in ZOOKEEPER-1030
ZOOKEEPER-1074	zkServer.sh is missing nohup/sleep, which are necessary for remote invocation
ZOOKEEPER-1076	some quorum tests are unnecessarily extending QuorumBase
ZOOKEEPER-1083	Javadoc for WatchedEvent not being generated
ZOOKEEPER-1086	zookeeper test jar has non mavenised dependency.
ZOOKEEPER-1087	ForceSync VM arguement not working when set to "no"
ZOOKEEPER-1088	delQuota does not remove the quota node and subsequent setquota calls for that path will fail
ZOOKEEPER-1090	Race condition while taking snapshot can lead to not restoring data tree correctly
ZOOKEEPER-1091	when the chrootPath of ClientCnxn is not null and the Watches of zooKeeper is not null and the method primeConnection(SelectionKey k) of ClientCnxn Occurred again for some reason ,then the wrong watcher clientPath is sented to server
ZOOKEEPER-1097	Quota is not correctly rehydrated on snapshot reload
ZOOKEEPER-1101	Upload zookeeper-test maven artifacts to maven repository.
ZOOKEEPER-1108	Various bugs in zoo_add_auth in C

Issue	Notes
ZOOKEEPER-1109	Zookeeper service is down when SyncRequestProcessor meets any exception.
ZOOKEEPER-1111	JMXEnv uses System.err instead of logging
ZOOKEEPER-1119	zkServer stop command incorrectly reading comment lines in zoo.cfg
ZOOKEEPER-1124	Multiop submitted to non-leader always fails due to timeout
ZOOKEEPER-1136	NEW_LEADER should be queued not sent to match the Zab 1.0 protocol on the twiki
ZOOKEEPER-1138	release audit failing for a number of new files
ZOOKEEPER-1139	jenkins is reporting two warnings, fix these
ZOOKEEPER-1140	server shutdown is not stopping threads
ZOOKEEPER-1141	zkpython fails tests under python 2.4
ZOOKEEPER-1142	incorrect stat output
ZOOKEEPER-1144	ZooKeeperServer not starting on leader due to a race condition
ZOOKEEPER-1145	ObserverTest.testObserver fails at particular point after several runs of ant junt.run -Dtestcase
ZOOKEEPER-1146	significant regression in client (c/python) performance
ZOOKEEPER-1152	Exceptions thrown from handleAuthentication can cause buffer corruption issues in NIOServer
ZOOKEEPER-1154	Data inconsistency when the node(s) with the highest zxid is not present at the time of leader election
ZOOKEEPER-1156	Log truncation truncating log too much - can cause data loss
ZOOKEEPER-1165	better eclipse support in tests
ZOOKEEPER-1168	ZooKeeper fails to run with IKVM
ZOOKEEPER-1171	fix build for java 7
ZOOKEEPER-1174	FD leak when network unreachable

Issue	Notes
ZOOKEEPER-1181	Fix problems with Kerberos TGT renewal
ZOOKEEPER-1185	Send AuthFailed event to client if SASL authentication fails
ZOOKEEPER-1189	For an invalid snapshot file(less than 10bytes size) RandomAccessFile stream is leaking.
ZOOKEEPER-1190	ant package is not including many of the bin scripts in the package (zkServer.sh for example)
ZOOKEEPER-1195	SASL authorizedID being incorrectly set: should use getHostName() rather than getServiceName()
ZOOKEEPER-1203	Zookeeper systest is missing Junit Classes
ZOOKEEPER-1206	Sequential node creation does not use always use digits in node name given certain Locales.
ZOOKEEPER-1212	zkServer.sh stop action is not conformant with LSB para 20.2 Init Script Actions
ZOOKEEPER-1237	ERRORs being logged when queued responses are sent after socket has closed.
Improvements	
ZOOKEEPER-494	zookeeper should install include headers in /usr/local/include/zookeeper
ZOOKEEPER-500	Async methods shouldnt throw exceptions
ZOOKEEPER-631	zkpython's C code could do with a style clean-up
ZOOKEEPER-636	configure.ac has instructions which override the contents of CFLAGS and CXXFLAGS.
ZOOKEEPER-724	Improve junit test integration - log harness information
ZOOKEEPER-733	use netty to handle client connections
ZOOKEEPER-765	Add python example script
ZOOKEEPER-773	Log visualisation
ZOOKEEPER-788	Add server id to message logs
ZOOKEEPER-789	Improve FLE log messages

Issue	Notes
ZOOKEEPER-797	c client source with AI_ADDRCONFIG cannot be compiled with early glibc
ZOOKEEPER-809	Improved REST Interface
ZOOKEEPER-821	Add ZooKeeper version information to zkpython
ZOOKEEPER-853	Make zookeeper.is_unrecoverable return True or False and not an integer
ZOOKEEPER-862	Hedwig created ledgers with hardcoded Bookkeeper ensemble and quorum size. Make these a server config parameter instead.
ZOOKEEPER-864	Hedwig C++ client improvements
ZOOKEEPER-891	Allow non-numeric version strings
ZOOKEEPER-905	enhance zkServer.sh for easier zookeeper automation-izing
ZOOKEEPER-926	Fork Hadoop common's test-patch.sh and modify for Zookeeper
ZOOKEEPER-977	passing null for path_buffer in zoo_create
ZOOKEEPER-980	allow configuration parameters for log4j.properties
ZOOKEEPER-993	Code improvements
ZOOKEEPER-997	ZkClient ignores command if there are any space in front of it
ZOOKEEPER-1018	The connection permutation in get_addrs uses a weak and inefficient shuffle
ZOOKEEPER-1025	zkCli is overly sensitive to to spaces.
ZOOKEEPER-1030	Increase default for maxClientCnxns
ZOOKEEPER-1094	Small improvements to LeaderElection and Vote classes
ZOOKEEPER-1095	Simple leader election recipe
ZOOKEEPER-1103	In QuorumTest, use the same "for (.. try { break } catch { })" pattern in testFollowersStartAfterLeaders as in testSessionMove.

Issue	Notes
ZOOKEEPER-1104	CLONE - In QuorumTest, use the same "for (.. try { break } catch { })" pattern in testFollowersStartAfterLeaders as in testSessionMove.
ZOOKEEPER-1143	quorum send & rcv workers are missing thread names
ZOOKEEPER-1153	Deprecate AuthFLE and LE
ZOOKEEPER-1166	Please add a few svn:ignore properties
ZOOKEEPER-1169	Fix compiler (eclipse) warnings in (generated) jute code
ZOOKEEPER-1243	New 4lw for short simple monitoring ldck
Features	
ZOOKEEPER-464	Need procedure to garbage collect ledgers
ZOOKEEPER-465	Ledger size in bytes
ZOOKEEPER-546	add "diskless" ensemble support
ZOOKEEPER-712	Bookie recovery
ZOOKEEPER-729	Recursively delete a znode - zkCli.sh rmr /node
ZOOKEEPER-744	Add monitoring four-letter word
ZOOKEEPER-747	Add C# generation to Jute
ZOOKEEPER-775	A large scale pub/sub system
ZOOKEEPER-799	Add tools and recipes for monitoring as a contrib
ZOOKEEPER-808	Web-based Administrative Interface
ZOOKEEPER-859	Native Windows version of C client
ZOOKEEPER-938	Support Kerberos authentication of clients.
ZOOKEEPER-992	MT Native Version of Windows C Client
ZOOKEEPER-999	Create an package integration project
ZOOKEEPER-1012	support distinct JVMFLAGS for zookeeper server in zkServer.sh and zookeeper client in zkCli.sh

Issue	Notes
ZOOKEEPER-1020	Implement function in C client to determine which host you're currently connected to.
ZOOKEEPER-1107	automating log and snapshot cleaning
Tasks	
ZOOKEEPER-754	numerous misspellings "succesfully"
ZOOKEEPER-1149	users cannot migrate from 3.4->3.3->3.4 server code against a single datadir
Tests	
ZOOKEEPER-239	ZooKeeper System Tests

Table 1: Changes Since ZooKeeper 3.3.0