

Welcome to Curator

Pages

Please refer to the [Curator Website](#) for details on Curator.

- [Contributions to Curator](#)
- [Curator's History](#)
- [For Curator Committers](#)
 - [Example Emails](#)
- [New Committers](#)
- [Old Change Log](#)
- [Powered By](#)
- [Releases](#)
- [Submitting Pull Requests](#)
- [Tech Notes](#)
 - [TN1](#) — ZooKeeper watches are single threaded.
 - [TN2](#) — InterProcessMutex acquire() can be used to return immediately if lock can't be acquired.
 - [TN3](#) — Dealing with session failure.
 - [TN4](#) — ZooKeeper makes a very bad Queue source.
 - [TN5](#) — The APIs in Apache Curator are exactly the same as Netflix Curator.
 - [TN6](#) — Friends don't let friends write ZooKeeper recipes
 - [TN7](#) — Do not use paths passed to Curator recipes.
 - [TN8](#) — Curator logging can be customized.
 - [TN9](#) — NOTE: PathChildrenCache now uses getData() instead of checkExists().
 - [TN10](#) — JVM pauses can cause unexpected client state with improperly chosen session timeouts.
 - [TN11](#) — Curator internally wraps Watchers.
 - [TN12](#) — Curator connection semantics
 - [TN13](#) — Guava usage in Curator
 - [TN14](#) — Session expiration in Curator
 - [TN15](#) — API Changes in Curator 5.0.0 and how to workaround them if you need binary compatibility.
- [Third Party Articles](#)

Useful links

- [ZooKeeper](#)
- [Curator Mailing Lists](#)
- [Curator Issues](#)
- [Curator Team](#)
- [Getting started with Curator](#)
- [Javadoc](#)