

ConcurProposal

Concur Proposal

Abstract

Concur is an open source java implementation for RAFT consensus protocol[1]. RAFT is being used successfully as an alternative to Paxos to implement a consistently replicated log. RAFT is proven to be safe and is designed to be simpler to understand.

Proposal

Concur is implemented as an independent RAFT library that can be used by any application to manage their replicated logs or replicated state machines. The implementation closely follows the original design proposed in the RAFT paper.

Background

Replicated log is a frequently used technique in distributed systems for reliability and parallelism. The consistency of replicas is an important requirement to ensure the correctness of the applications accessing the data in these replicas. Many algorithms have been proposed and used to maintain consistency and correctness. RAFT paper was proposed in 2012 and has gained significant popularity since then to manage the replicated logs and replicated state in many popular projects. Data replication is a requirement for many projects in ASF ecosystem for example Apache Hadoop, Apache HBase etc. We believe the Concur project can provide an implementation of RAFT that can fulfill the needs similar to these systems and strengthen the ASF ecosystem.

Rationale

There are a few RAFT implementations[2], but none are a part of ASF. Concur is significantly different because one of the critical goals of the project is to use it as a library with pluggability for different RPC, Raft log and state machine implementations. Another important goal of Apache Concur is to provide high throughput for large ingest rates of data, with data pipeline support. This will simplify the problem of data replication across the board. Apache Kudu uses RAFT protocol, but it has its own C++ implementation. Apache [DistributedLog] project[3] (in incubation) provides a replicated log service. However, Apache Concur is different as it provides a java library that other projects can use to implement their own replicated state machine, without deploying another service.

Current Status

The source code is available with Apache v2 license at <https://github.com/hortonworks/concur>. Significant amount of code has been added and basic functionality of RAFT is available for testing. The code is still in pre-alpha stage and is currently being tried out in POC mode.

Meritocracy

We plan to invest in supporting meritocracy. We intend to invite additional developers to participate. We will encourage and monitor community participation so that privileges can be extended to those that contribute.

Community

The developers on the initial committers list are experienced in the ASF ecosystem:

- Jing Zhao (jing9 at apache dot org)
- Tsz Wo Nicholas Sze (szetetszwo at apache dot org)
- Li Lu (llu at hortonworks dot com)
- Enis Soztutar (enis at apache dot org)
- Anu Engineer (anu at apache dot org)
- Arpit Agarwal (arpitagarwal at apache dot org)
- Chris Nauroth (cnauroth at apache dot org)
- Jakob Homan (jghoman at apache dot org)
- Mayank Bansal (mayank_bansal at apache dot org)
- Xiaoyu Yao (xyao at apache dot org)
- Chen Liang (cliang at hortonworks dot com)
- Hanisha Koneru (hkoneru at hortonworks dot com)
- Xiaobing Zhou (xiaobingo at gmail dot com)
- Mingliang Liu (liuml07 at apache dot org)
- Uma Maheswara Rao G (umamahesh at apache dot org)
- Jitendra Pandey (jitendra at apache dot org)

Affiliations

- Jing Zhao, Tsz Wo Nicholas Sze, Li Lu, Enis Soztutar, Anu Engineer, Arpit Agarwal, Xiaoyu Yao, Chen Liang, Hanisha Koneru, Mingliang Liu, Jitendra Pandey, Devaraj Das
 - Hortonworks

- Uma Gangumalla, Intel
- Mayank Bansal, Uber
- Jakob Homan, [OfferUp](#)
- Chris Nauroth, Disney

Alignment

We believe that Concur will address the requirements of several projects and communities in the Apache ecosystem and will gain great adoption.

Known Risks

Orphaned Products

The contributors are leading users and vendors in the Apache Hadoop ecosystem, with significant open source experience, so the risk of being orphaned is relatively low. The project could be at risk if vendors decided to change their strategies in the market. In such an event, the current committers plan to continue working on the project on their own time, though the progress will likely be slower. We plan to mitigate this risk by encouraging and recruiting additional committers. Since replicated log is a very common technique useful for large number of applications, we believe many developers would like to join and contribute to the project.

Inexperience with Open Source

The initial committers include veteran Apache Members (Committers, PMC Members and Apache Members) and other developers who have varying degrees of experience with open source projects. All have been involved with source code that has been released under an open source license, and also have experience developing code with an open source development process.

Homogenous Developers

The initial list of committers include senior Apache developers from multiple different organizations including Hortonworks, [OfferUp](#), Intel, Uber etc. We believe this project will benefit many different projects and therefore more and more diverse set of developers will join over time.

Reliance on Salaried Developers

It is expected that Concur development will occur on both salaried time and on volunteer time, after hours. The majority of initial committers are paid by their employer to contribute to this project. However, they are all passionate about the project, and we are confident that the project will continue even if no salaried developers contribute to the project. We are committed to recruiting additional committers including non-salaried developers.

Relationships with Other Apache Products

Most of the initial developers are active participants in Apache Hadoop community and we expect adoption of Concur by Hadoop Community. However, Concur is a generic RAFT library and we will strive to take it to more general adoption. This project depends on Apache Maven, Apache Commons {collections,configuration,io} and Apache Hadoop.

Initial Source

<https://github.com/hortonworks/concur>

External Dependencies

- Apache Maven
- Apache Commons {collections,configuration,io}
- Apache Hadoop
- Netty (Apache 2.0)
- Guava (Apache 2.0)
- Protobuf (Google 2014)
- gRPC (Google 2015)
- slf4j (MIT License)
- Mockito (MIT)
- junit (EPL 1.0)
- [FindBugs](#) (LGPL)

Required Resources

Mailing List

- concur-private
- concur-dev
- concur-user

- [concur-issues](#)

Git Repository

Git is the preferred source control system: [git://git.apache.org/concur](https://git.apache.org/concur)

Issue Tracking

JIRA Concur (Concur)

Sponsors

Champion

Jitendra Pandey ([jitendra](#))


Nominated Mentors

- Chris Nauroth ([cnauroth](#))
- Devaraj Das ([devaraj](#))
- Jakob Homan ([jghoman](#))
- Uma Maheswara Rao G ([umamahesh](#))

Sponsoring Entity

Incubator PMC

Reference

- [1] Diego Ongaro and John Ousterhout. 2014. In search of an understandable consensus algorithm. In Proceedings of the 2014 USENIX conference on USENIX Annual Technical Conference (USENIX ATC'14), Garth Gibson and Nickolai Zeldovich (Eds.). USENIX Association, Berkeley, CA, USA, 305-320.
- [2] <https://raft.github.io/> 
- [3] Apache [DistributedLog] (incubating) <http://distributedlog.incubator.apache.org/> 