

GossipProposal

⚠ FINAL ⚠

This proposal is now complete and has been submitted for a VOTE.

Abstract

Apache Gossip will be an implementation of the Gossip Protocol based on code available here: <https://github.com/edwardcapriolo/gossip/> which is already licenced using the glorious Apache V2 License.

Proposal

Apache Gossip aims to provide a gossip based consensus protocol written in Java for peer-to-peer communication to the Apache Incubator (<http://incubator.apache.org/>). This implementation will effectively scale from one to one-thousand node clusters. In addition to the code implementation, the project should produce specifications of the wire protocol, features, and expected behavior of the system such that compatible implementations can communicate.

Background

The gossip protocol has been implemented to varying levels of rigor by a number of entities. In particular, Apache Cassandra uses an implementation of gossip to locate peers and transmit up/down state. Apache Spark leverages tooling in Akka which provides peer-to-peer node discovery capabilities.

- <http://highscalability.com/blog/2011/11/14/using-gossip-protocols-for-failure-detection-monitoring-mess.html>
- https://en.wikipedia.org/wiki/Gossip_protocol

Rationale

With distributed computing becoming extremely widespread, and the growth of the buzz-factor of 'the-internet-of-things' it is increasingly important that networks of IP addressable devices can form a peer-to-peer network. Applications of peer-to-peer networks include generating crypto currency, managing hardware such as solar power micro-grids, and more traditional roles like grid/High Performance Computing and distributed storage systems. Different implementations of gossip based consensus protocols have been implemented in numerous languages or as part of more complex software stacks. The Apache Software Foundation should lead the effort of producing a purpose built tool that can be used by downstream projects to form peer-to-peer networks.

Initial Goals

- Migration of current code <https://github.com/edwardcapriolo/gossip> and existing community to the Apache Software Foundation infrastructure
- Secure communications
 - Transport security using a pre-shared key
 - Public Key Infrastructure
- Introduce a cluster name to wire protocol to avoid misconfigurations
- Effectively operate when systems have multiple network interfaces by controlling IP binding settings
- Effectively operate when systems have Network Address Translations devices between them using a broadcast IP settings
- Develop advanced integration testing from cluster sizes of 1-1000 nodes
 - Test convergence times
 - Demonstrate the tradeoffs of different settings in regard to bandwidth/cpu/convergence time/accuracy
- Gossip data other than cluster state such as application/user data
- Provide detailed specifications such that others can implement the protocol in other programming languages
- Explore HTTP transport as an alternative to UDP

Current Status

The current code has been around for some time. Previously it was a Google code project. Since the fork in January 2015 there have been 55 commits and 4 releases.

Meritocracy

We believe in meritocracy. All suggestions are taken seriously. We enjoy helping new people become part of process. For other projects available on our Github, once a user shows enough activity we grant them collaborator status.

Community

In a relatively short amount of time, with a small amount of promotion on twitter and through blogging, we have 50+ followers on Github and several forks of the project. With the Apache brand we should be able to attract more. Once we have entered the incubator we believe it will be easier to attempt to unify with other similar implementations.

Core Developers

The code was forked on Jan 9th 2015, since then there have been 4 releases and 55 commits. Since that period, the majority of the work was undertaken by Edward Capriolo. Several people are interested in the features of this proposal and have indicated they will volunteer their time.

Alignment

Apache is the perfect organization to take on the Gossip project. Besides benefiting a number of projects directly, the active development and outreach will increase adoption of Gossip with the aim of it becoming a leader in the space.

Known Risks

Several existing implementations of similar cluster membership systems (gossip based and otherwise) exist. A key challenge is moving from a relatively niche technical audience to a more general tool for solving a common problem. Differentiating when Apache Gossip may be the optimal solution versus a clouded landscape of coordination services such as etcd or zookeeper, and distributed data stores via feature set will also be key. We believe that users be attracted to the peer-to-peer distributed service toolkit that gossip will provide.

Orphaned products

We plan on building on the current code by developing discrete features with a focus on testing. Up until this point the project has been maintained by a single person. However the project currently releases artifacts to maven central, is tested using travis CI, and follows controlled development practice. This level of dedication will see the process through the initial stages.

Inexperience with Open Source

We are very familiar with Open Source development and the Apache Foundation. The current code base already carries an Apache V2 Licence.

Homogenous Developers

Multiple people have made contributions to the current code base. This proposal has generated more interest and several more are offering to volunteer time. These volunteers are from diverse corporate entities and many of them are also affiliated with existing Apache projects (both top level and incubating) so there is a already experience and degree of incubating knowledge within the proposed incubating community.

Reliance on Salaried Developers

We wish to create Apache Gossip for the challenge of producing great software. Initially all members of the project will volunteer their time and no one will be expressly salaried to work only on this project.

Relationships with Other Apache Products

If the Apache Gossip project is successful, other products in Apache such as Storm or Cassandra could adopt it. However, adoption by those specific projects is not our criteria for success. There are a large number of applications for this system. One example is Apache JMeter could be built with Apache Gossip as a backend for distributed testing. Another example is a polyglot registry of thrift services with gossip based discovery.

A Excessive Fascination with the Apache Brand

We care about the Apache foundation. Having the recognition of the Apache incubator will undoubtedly help the project. We do not seek the Apache brand to be used as legal shield or personal glory. We believe in the Apache foundation and will manage the project with esprit de corps, welcoming all through meritocracy while using bylaws as guiding values.

Initial Source

The initial source for <https://github.com/edwardcapriolo/gossip> was forked from <https://code.google.com/p/java-gossip/>. The code base always Apache V2 licensed.

Source and Intellectual Property Submission Plan

During the course of proposal development, the two original authors of the software were contacted to see if they would be interested in joining the project as initial committers, and if they are willing to submit an SGA to the ASF. The first author has responded positively and has been added as an initial committer. We would prefer to have an SGA from both authors, but given second author has not contributed to the codebase in 6 years, the ability to obtain an SGA is not certain. If we do not hear back from the second author within a reasonable time frame, we intend to proceed with the Incubator IP Clearance process.

1. Make final commit on the gossip [GitHub](#) project explaining the move to ASF.
2. Complete the Incubator IP Clearance process.
3. Move code into ASF repo.
4. Rename references to old name.
5. Apply Apache V2 licence to all source files.

External Dependencies

Currently the project encodes messages into JSON for network transmission. This is done using JSONorg, but will switch to jackson (potentially before the move to the ASF).

Cryptography

The current code is not using cryptography. It is on the road map to add security through transport encryption (SSL) and transport encryption.

Required Resources

Mailing lists

- dev@gossip.incubator.apache.org
- commits@gossip.incubator.apache.org
- private@gossip.incubator.apache.org

The user list will be added when we have broader adoption.

Git Repository

<https://git-wip-us.apache.org/repos/asf/incubator-gossip.git>

Issue Tracking

JIRA tracker: GOSSIP

Initial Committers

- Edward Capriolo (ecapriolo at apache dot org)
- Josh Clemm (clemm22 at gmail dot com)
- P. Taylor Goetz (ptgoetz at apache dot org)
- Gary Dusbabek (gdusbabek at apache dot org)
- Dorian Ellerbe (Doellerbe06 at gmail dot com)(requires CLA)
- Sathish Dhinakaran (requires CLA)
- Joe Price (pricejosephd at gmail dot com)(requires CLA)

Affiliations

- Edward Capriolo - The Huffington Post
- P. Taylor Goetz - Hortonworks
- Gary Dusbabek - Silicon Valley Data Science
- Dorian Ellerbe - Dstillery
- Sathish Dhinakaran - Dstillery
- Sean Busbey - Cloudera
- Josh Elser - Hortonworks

Additional Interested Contributors

Those interested in getting involved with the project as it starts are encourage to list themselves here.

- Suneel Marthi (smarthi at apache dot org) - Red hat Inc.
- Debo Dutta (ddutta at apache dot org) - Cisco

Sponsors

Champion

P. Taylor Goetz (ASF Member, IPMC)

Proposed Mentors

- Sean Busbey (ASF Member, IPMC)
- Josh Elser (ASF Member, IPMC)
- P. Taylor Goetz (ASF Member, IPMC)

Sponsoring Entity

The Apache Incubator