HelixProposal

Abstract

Helix is a cluster management system for managing partitioned and replicated resources in distributed data systems.

Proposal

Helix provides an abstraction that separates coordination and management tasks from functional tasks of a distributed system. The developer defines the system behavior via a state machine, the transitions between those states, and constraints on states and transitions that govern the system's valid settings. Helix ensures the distributed system satisfies the state machine, controlling state changes as appropriate during common operational activities such as upgrades, component failures, bootstrapping, running maintenance tasks, and adding capacity.

Background

Helix was developed at LinkedIn to manage large clusters for several diverse applications, including a distributed, partitioned, replicated, highly available document store with a master-slave model, a search service with multiple replicas that are updated atomically and in near real-time, and a change data capture service for reliably transporting database changes to caches, other dependent databases and indexes.

These services use Helix to reliably manage dozens of clusters in multiple data centers. These services meet stringent SLAs at large scale for missioncritical production applications such as search, social gestures, and profiles. Helix has proven to be flexible for a wide variety of system configurations and operational patterns, is easy to integrate, with pluggable interfaces enabling custom behavior. It depends on Apache Zookeeper for coordination and tracking of system state across the cluster, as well as providing fault tolerance. Helix is written in Java. It was developed internally at LinkedIn to meet our particular use cases, but will be useful to many organizations facing a similar need to manage large clusters. Therefore, we would like to share it the ASF and begin developing a community of developers and users within Apache.

Rationale

Many organizations can benefit from a generalized cluster management system such as Helix. While our distributed data systems use-cases for a very large website like LinkedIn has driven the design of Helix, its uses are varied and we expect many new use cases to emerge.

Current Status

Meritocracy

Our intent with this incubator proposal is to start building a diverse developer community around Helix following the Apache meritocracy model. Since Helix was initially developed in late 2011, we have had fast adoption and contributions by multiple teams at LinkedIn. We plan to continue support for new contributors and work with those who contribute significantly to the project to make them committers.

Community

Helix is currently being used internally at LinkedIn and is in production in that company for customer-facing features. Recent public presentations of Helix and its goals garnered much interest from potential contributors. We hope to extend our contributor base significantly and invite all those who are interested in building large-scale distributed systems to participate. To further this goal, we use GitHub issue tracking and branching facilities.

Core Developers

Helix is currently being developed by three engineers at LinkedIn: Kishore Gopalakrishna, Shi Lu and Jason Zheng, and Adam Silberstein, an engineer at Trifacta. Kishore, the lead developer and architect, has experience within Apache as an S4 committer. Shi developed the partition to node mapping and rebalancing algorithm, cluster admin APIs, and the health check framework. Jason developed the cluster controller and most of the test framework. Adam developed the rich alerting framework that enables cluster-wide, "intelligent" alerts.

Alignment

The ASF is the natural choice to host the Helix project as its goal of encouraging community-driven open-source projects fits with our vision for Helix. Many projects that can benefit from Helix will rely on Apache ZooKeeper for cluster state management, and can far more easily achieve their operational goals by using Helix.

Known Risks

Orphaned Products

The core developers plan to work full time on the project. There is very little risk of Helix being abandoned as it is a critical part of LinkedIn's internal infrastructure and is in production use.

Inexperience with Open Source

Only one of the core developers has experience with open source development. Kishore has been actively involved with the ASF as a committer and lead developer of S4.

Homogeneous Developers

The current core developers are all from LinkedIn. However, we hope to establish a developer community that includes contributors from several corporations and we are actively encouraging new contributors via the mailing lists and public presentations of Helix.

Reliance on Salaried Developers

Currently, the developers are paid to do work on Helix. However, once the project has a community built around it, we expect to get committers, developers and community from outside the current core developers. However, because LinkedIn relies on Helix internally, the reliance on salaried developers is unlikely to change.

Relationships with Other Apache Products

Helix uses Apache ZooKeeper to coordinate its state amongst the managed cluster components and for leader election to provide fault tolerance, and uses Apache Maven for build management.

An Excessive Fascination with the Apache Brand

While we respect the reputation of the Apache brand and have no doubts that it will attract contributors and users, our interest is primarily to give Helix a solid home as an open source project following an established development model. We have also given reasons in the Rationale and Alignment sections.

Documentation

Information about Helix can be found at [https://github.com/linkedin/helix/wiki]. The following links provide more information about the project:

- Home Page: http://linkedin.github.com/helix/
- Github source: https://github.com/linkedin/helix
- Documentation: https://github.com/linkedin/helix/wiki
- Javadocs: http://linkedin.github.com/helix/apidocs/

Initial Source

Helix has been under development at [LinkedIn] since April 2011. It is currently hosted on github under the Apache license 2 at [https://github.com/linkedin /helix] Helix is written in Java. Its source tree is entirely self-contained and relies on Maven as its build system and dependency resolution mechanism.

External Dependencies

The dependencies all have Apache compatible licenses.

- log4j
- zookeeper
- xstream
- jackson-core-asl
- jackson-mapper-asl
- commons-io
- commons-cli
- commons-math
- zkclient
- camel-josql
- camel-core
- gentlyweb-utils
- josql
- commons-management
- commons-logging-api
- org.restlet
- com.noelios.restlet
 net.sf.jsqlparser

Non-Apache build tools that are used by Helix are as follows:

Cobertura: GNU GPLv2

Note that Cobertura is optional and is only used for calculating unit test coverage.

Cryptography

Not applicable.

Required Resources

Mailing Lists

- helix-private for private PMC discussions (with moderated subscriptions)
- helix-dev
- helix-commits
- helix-user

Git Directory

Since Git is now available to be used as primary repo type, Helix would be available in the git repository instead of svn. [https://git.apache.org/helix.git]

Issue Tracking

JIRA Helix (HELIX)

Other Resources

The existing code already has unit and integration tests, so we would like a Jenkins instance to run them whenever a new patch is submitted. This can be added after project creation.

Initial Committers

- Kishore Gopalakrishna
- Shi Lu
- Zhen Zheng
- Adam Silberstein
- Kapil Surlaker
- Bob Schulman
- Swaroop Jagadish
- Rahul Aggarwal
- Terence Yim
- Santiago Perez

Affiliations

- Kishore Gopalakrishna (LinkedIn)
- Shi Lu (LinkedIn)
- Jason Zheng (LinkedIn)
- Adam Silberstein (Trifacta)
- Kapil Surlaker (LinkedIn)
- Bob Schulman (LinkedIn)
- Swaroop Jagadish (LinkedIn)
- Rahul Aggarwal (LinkedIn)
- Terence Yim (LinkedIn)
- Santiago Perez (LinkedIn)

Sponsors

Champion

• Patrick Hunt (Apache Member)

Nominated Mentors

- Patrick Hunt (Apache Member)
- Mahadev Konar (Apache Member)
- Owen O'Malley (Apache Member)
- Olivier Lamy (Apache Member)

Sponsoring Entity

We are requesting the Incubator to sponsor this project.