

ReefProposal

REEFPproposal - Incubator

Abstract

REEF (Retainable Evaluator Execution Framework) is a scale-out computing fabric that eases the development of Big Data applications on top of resource managers such as Apache YARN and Mesos.

Proposal

REEF is a Big Data system that makes it easy to implement scalable, fault-tolerant runtime environments for a range of data processing models (e.g., graph processing and machine learning) on top of resource managers such as Apache YARN and Mesos. REEF provides capabilities to run multiple heterogeneous frameworks and workflows of those efficiently.

Additionally, REEF contains two libraries that are of independent value: Wake is an event-based-programming framework inspired by Rx and SEDA. Tang is a dependency injection framework inspired by Google Guice, but designed specifically for configuring distributed systems.

Background

The resource management layer such as Apache YARN and Mesos has emerged as a critical layer in the new scale-out data processing stack; resource managers assume the responsibility of multiplexing a cluster of shared-nothing machines across heterogeneous applications. They operate behind an interface for leasing containers

- a slice of a machine's resources - to computations in an elastic fashion. However, building data processing frameworks directly on this layer comes at a high cost: each framework must tackle the same challenges (e.g., fault-tolerance, task scheduling and coordination) and reimplement common mechanisms (e.g., caching, bulk transfers). REEF provides a reusable control-plane for scheduling and coordinating task-level work on cluster resource managers. The REEF design enables sophisticated optimizations, such as container re-use and data caching, and facilitates workflows that span multiple frameworks. Examples include pipelining data between different operators in a relational system, retaining state across iterations in iterative or recursive data flow, and passing the result of a [MapReduce](#) job to a Machine Learning computation.

Rationale

Since REEF is a library that makes it easy to write distributed applications on top of Apache YARN or Mesos, the Apache Software Foundation is the perfect home for hosting REEF.

Current Status

REEF has been developed mostly by Microsoft, UCLA and the Seoul National University. The REEF codebase is open-sourced under Apache License 2.0 and is currently hosted in a public repository at [github.com](#).

Meritocracy

We plan to build a strong open community by following the Apache meritocracy principles. We will work with those who contribute significantly to the project and invite them to be its committers.

Community

REEF is currently being used internally at Microsoft. Also, SK Telecom builds their data analytics infrastructure on top of REEF in collaboration with Seoul National University. We hope to extend our contributor base by becoming an Apache incubator project. REEF will attract developers who are interested in creating common building blocks for simplifying the development of large-scale big data applications.

Core Developers

Core developers are engineers from Microsoft, Purestorage, UCB, UCLA, UW and Seoul National University.

Alignment

REEF depends on many Apache projects and dependencies. REEF is built on resource managers such as Apache YARN and Apache Mesos. REEF also uses HDFS as a distributed storage layer.

Known Risks

Orphaned Products

The risk of REEF being orphaned is small because Microsoft products are built on REEF. The core REEF developers continue to work on REEF at Microsoft, UCLA, and Seoul National University. The REEF project is gaining interest from other institutions to be used as their infrastructure.

Inexperience with Open Source

Several core developers have experience with open source development. REEF committers will be guided by the mentors with strong Apache open source project backgrounds.

Homogeneous Developers

The initial committers include developers from several institutions including Microsoft, Purestorage, UCB, UCLA, and Seoul National University.

Reliance on Salaried Developers

Developers from Microsoft are paid to work on REEF. Since the work is used internally at Microsoft, Microsoft will keep supporting the developers to work on REEF. There are also engineers and graduate students that contribute to REEF from UCLA, UCB, UW and Seoul National University. We plan to attract active developers from other institutions.

Relationships with Other Apache Products

Given REEF's position in the big data stack, there are three relationships to consider: Projects that fit below, on top of, or alongside REEF in the stack.

Below REEF: Mesos and YARN

REEF is designed to facilitate application development on top of resource managers. Hence, its relationship with the aforementioned resource managers is symbiotic by design.

On Top of REEF

Apache Spark, Giraph, [MapReduce](#) and Flink are only some of the projects that logically belong at a higher layer of the big data stack than REEF. Of course, none of these today actually are leveraging REEF and had to each individually solve some of the issues REEF addresses. It is our goal that REEF will help developers create an even richer set of future big data frameworks.

Alongside REEF

Apache hosts several projects building intermediate, library layers on top of a resource management platform. Twill, Slider, and Tez are notable examples in the incubator. These projects share many objectives with REEF (and each other). We expect these parallel explorations to converge and differentiate within Apache, as the space for distributed applications and deployment is too vast for a single answer.

Apache Twill and REEF both aim to simplify application development on top of resource managers. However, REEF and Twill go about this in different ways: Twill simplifies programming by exposing a programming model, Java Threads. REEF on the other hand provides a set of common building blocks (e.g., job coordination, state passing, cluster membership) for building big data processing applications and virtualizes underlying resources managers. None of this prescribes a specific programming model. As such, REEF occupies a slot ever so slightly below Twill in an architecture stack. Apache Slider is a framework to make it easy to deploy and manage long-running static applications in a YARN cluster. The focus is to adapt existing applications such as HBase and Accumulo to run on YARN with little modification. Therefore, the goals of Slider and REEF are different.

Apache Tez is a project to develop a generic Directed Acyclic Graph (DAG) processing framework with a reusable set of data processing primitives. The initial focus is to provide improved data processing capabilities for projects like Apache Hive, Apache Pig, and Cascading. Tez is still a single framework for DAG processing. In contrast, REEF provides a generic layer on which diverse computation models (DAG, ML, Graph processing, and Interactive query processing) can be built. More importantly, REEF provides a layer that facilitates inter-framework resource and in-memory state use and virtualizes resource managers. Regarding re-usable data processing primitives, Tez and REEF share the same goal. We hope to collaborate on features which can be shared between Tez and REEF.

Apache Helix automates application-wide management operations which require global knowledge and coordination, such as repartitioning of resources and scheduling of maintenance tasks. Helix separates global coordination concerns from the functional tasks of the application with a state machine abstraction. REEF's generic layer makes it easy to program the functional and management tasks, which may span small or large groups within the application. Helix can work hand-in-hand with REEF, by providing the global management component for REEF applications.

An Excessive Fascination with the Apache Brand

The Apache Software Foundation has a reputation of being the best place to host open source projects. We believe that we will attract many developers who want to contribute to innovating in the Big Data platform space by joining the Apache Software Foundation.

Documentation

The current documentation for REEF is at <https://github.com/Microsoft-CISL/REEF> as well as on <http://www.reef-project.org>

Initial Source

The REEF codebase is currently hosted at <https://github.com/Microsoft-CISL/REEF>.

External Dependencies

REEF makes extensive use of the vast array of Java libraries from the Apache Software Foundation, namely:

- avro (Apache 2.0)
- hadoop (Apache 2.0)
- hdfs (Apache 2.0)
- yarn (Apache 2.0)
- commons-cli (Apache 2.0)
- commons-configuration (Apache 2.0)
- commons-lang (Apache 2.0)
- commons-logging (Apache 2.0)

To the best of our knowledge, the external dependencies of REEF are distributed under Apache compatible licenses:

- guava-libraries (Apache 2.0)
- protobuf (BSD)
- asm (BSD)
- netty (Apache 2.0)
- mockito (MIT)
- junit (EPL 1.0)
- slf4j (MIT)

Cryptography

REEF will depend on secure Hadoop, which can optionally use Kerberos.

Required Resources

Mailing Lists

- private@reef for private PMC discussions
- dev@reef for technical discussions among contributors and notification about commits

Subversion Directory

The REEF team uses Git for source version control:
[git://git.apache.org/reef](https://git.apache.org/reef)

Issue Tracking

JIRA REEF (REEF)

Other Resources

Jenkins continuous integration testing

Initial Committers

- Markus Weimer
- Sergiy Matushevych
- Julia Wang
- Shravan M Narayanamurthy
- Yingda Chen
- Tony Majestro
- Beysim Sezgin
- Boris Shulman
- Russell Sears
- Jung Ryong Lee
- You Sun Jung
- Dong Joon Hyun
- Josh Rosen
- Tyson Condie
- Brandon Myers
- Yunseong Lee
- Taegeon Um
- Youngseok Yang
- Brian Cho
- Byung-Gon Chun

Affiliations

- Microsoft:
 - Markus Weimer
 - Sergiy Matusevych
 - Julia Wang
 - Shravan M Narayanamurthy
 - Yingda Chen
 - Tony Majestro
 - Beysim Sezgin
 - Boris Shulman
- Purestorage:
 - Russell Sears
- SK Telecom:
 - Jung Ryong Lee
 - You Sun Jung
 - Dong Joon Hyun
- University of California:
 - Josh Rosen (Berkeley)
 - Tyson Condie (LA)
- University of Washington:
 - Brandon Myers
- Seoul National University:
 - Yunseong Lee
 - Taegeon Um
 - Youngseok Yang
 - Brian Cho
 - Byung-Gon Chun

Sponsors

Champions

Chris Douglas <cdouglas@apache.org>

Nominated Mentors

- Chris Mattmann <mattmann@apache.org>
- Ross Gardler <rgardler@apache.org>
- Owen O'Malley <omalley@apache.org>

Sponsoring Entity

The Apache Incubator