

Accumulo Proposal

Abstract

Accumulo is a distributed key/value store that provides expressive, cell-level access labels.

Proposal

Accumulo is a sorted, distributed key/value store based on Google's BigTable design. It is built on top of Apache Hadoop, Zookeeper, and Thrift. It features a few novel improvements on the BigTable design in the form of cell-level access labels and a server-side programming mechanism that can modify key/value pairs at various points in the data management process.

Background

Google published the design of BigTable in 2006. Several other open source projects have implemented aspects of this design including HBase, CloudStore, and Cassandra. Accumulo began its development in 2008.

Rationale

There is a need for a flexible, high performance distributed key/value store that provides expressive, fine-grained access labels. The communities we expect to be most interested in such a project are government, health care, and other industries where privacy is a concern. We have made much progress in developing this project over the past 3 years and believe both the project and the interested communities would benefit from this work being openly available and having open development.

Current Status

Meritocracy

We intend to strongly encourage the community to help with and contribute to the code. We will actively seek potential committers and help them become familiar with the codebase.

Community

A strong government community has developed around Accumulo and training classes

have been ongoing for about a year. Hundreds of developers use Accumulo.

Core Developers

The developers are mainly employed by the National Security Agency, but we anticipate interest developing among other companies.

Alignment

Accumulo is built on top of Hadoop, Zookeeper, and Thrift. It builds with Maven. Due to the strong relationship with these Apache projects, the incubator is a good match for Accumulo.

Known Risks

Orphaned Products

There is only a small risk of being orphaned. The community is committed to improving the codebase of the project due to its fulfilling needs not addressed by any other software.

Inexperience with Open Source

The codebase has been treated internally as an open source project since its beginning, and the initial Apache committers have been involved with the code for multiple years. While our experience with public open source is limited, we do not anticipate difficulty in operating under Apache's development process.

Homogeneous Developers

The committers have multiple employers and it is expected that committers from different companies will be recruited.

Reliance on Salaried Developers

The initial committers are all paid by their employers to work on Accumulo and we expect such employment to continue. Some of the initial committers would continue as volunteers even if no longer employed to do so.

Relationships with Other Apache Products

Accumulo uses Hadoop, Zookeeper, Thrift, Maven, log4j, commons-lang, -net, -io, -jci, -collections, -configuration, -logging, and -codec.

Relationship to HBase

Accumulo and HBase are both based on the design of Google's BigTable, so there is a danger that potential users will have difficulty distinguishing the two. Some of the key areas in which Accumulo differs from HBase are discussed below. It may be possible to incorporate the desired features of Accumulo into HBase. However, the amount of work required would slow development of HBase and Accumulo considerably. We believe this warrants a podling for Accumulo at the current time. We expect active cross-pollination will occur between HBase and podling Accumulo and it is possible that the codebases and projects will ultimately converge.

Access Labels

Accumulo has an additional portion of its key that sorts after the column qualifier and before the timestamp. It is called column visibility and enables expressive cell-level access control. Authorizations are passed with each query to control what data is returned to the user. The column visibilities are boolean AND and OR combinations of arbitrary strings (such as "(A&B)|C") and authorizations are sets of strings (such as {C,D}).

Iterators

Accumulo has a novel server-side programming mechanism that can modify the data written to disk or returned to the user. This mechanism can be configured for any of the scopes where data is read from or written to disk. It can be used to perform joins on data within a single tablet.

Flexibility

HBase requires the user to specify the set of column families to be used up front. Accumulo places no restrictions on the column families. Also, each column family in HBase is stored separately on disk. Accumulo allows column families to be grouped together on disk, as does BigTable. This enables users to configure how their data is stored, potentially providing improvements in compression and lookup speeds. It gives Accumulo a row/column hybrid nature, while HBase is currently column-oriented.

Testing

Accumulo has testing frameworks that have resulted in its achieving a high level of correctness and performance. We have observed that under some configurations and conditions Accumulo will outperform HBase and provide greater data integrity.

Logging

HBase uses a write-ahead log on the Hadoop Distributed File System. Accumulo has its own logging service that does not depend on communication with the HDFS NameNode.

Storage

Accumulo has a relative key file format that improves compression.

Areas in which HBase features improvements over Accumulo

in memory tables, upserts, coprocessors, connections to other projects such as Cascading and Pig

Expectations

There is a risk that Accumulo will be criticized for not providing adequate security. The access labels in Accumulo do not in themselves provide a complete security solution, but are a mechanism for labeling each piece of data with the authorizations that are necessary to see it.

Apache Brand

Our interest in releasing this code as an Apache incubator project is due to its strong relationship with other Apache projects, i.e. Accumulo has dependencies on Hadoop, Zookeeper, and Thrift and has complementary goals to HBase.

Documentation

There is not currently documentation about Accumulo on the web, but a fair amount of documentation and training materials exists and will be provided on the Accumulo wiki at apache.org. Also, a paper discussing YCSB results for Accumulo will be presented at the 2011 Symposium on Cloud Computing.

Initial Source

Accumulo has been in development since spring 2008. There are hundreds of developers using it and tens of developers have contributed to it. The core codebase consists of 200,000 lines of code (mainly Java) and 100s of pages of documentation. There are also a few projects built on top of Accumulo that may be added to its contrib in the future. These include support for Hive, Matlab, YCSB, and graph processing.

Source and Intellectual Property Submission Plan

Accumulo core code, examples, documentation, and training materials will be submitted by the National Security Agency.

We will also be soliciting contributions of further plugins from MIT Lincoln Labs, Carnegie Mellon University, and others.

Accumulo has been developed by a mix of government employees and private companies under government contract. Material developed by government employees is in the public domain and no U.S. copyright exists in works of the federal government. For the contractor developed material in the initial submission, the U.S. Government has sufficient authority per the ICLA from the copyright owner to contribute the Accumulo code to the incubator.

There has been some discussion regarding accepting contributions from US Government sources on <https://issues.apache.org/jira/browse/LEGAL-93>. We propose that the NSA will sign an ICLA/CCLA if that document could be slightly modified to explicitly address copyright in works of government employees. Specifically, we propose that the definition of “You” be modified to include “the copyright owner, the owner of a Contribution not subject to copyright, or legal entity authorized by the copyright owner that is making this Agreement.” In addition, section 2, the copyright license grant be modified after “You hereby grant” that either states “to the extent authorized by law” or “to the extent copyright exists in the Contribution.” These changes will permit US Government employee developed work to be included.

One proposed solution is to form a Collaborative Research and Development Agreement (CRADA) between the Apache Software Foundation and the US Government, but this will not solve the underlying problem that U.S. law does not grant copyright to works of government employees. At this time a CRADA is not necessary but should it be determined that a CRADA is necessary, we would like to work through that process during the incubation phase of Accumulo rather than before acceptance as this may take time to enter into an agreement.

External Dependencies

jetty (Apache and EPL), jline (BSD), jfreechart (LGPL), jcommon (LGPL), slf4j (MIT), junit (CPL)

Cryptography

none

Required Resources

- Mailing Lists
 - accumulo-private
 - accumulo-dev
 - accumulo-commits
 - accumulo-user
- Subversion Directory
 - <https://svn.apache.org/repos/asf/incubator/accumulo>

- Issue Tracking
 - JIRA Accumulo (ACCUMULO)
- Continuous Integration
 - Jenkins builds on <https://builds.apache.org/>
- Web
 - <http://incubator.apache.org/accumulo/>
 - wiki at <http://wiki.apache.org> or <http://cwiki.apache.org>

Initial Committers

- Aaron Cordova (aaron at cordovas dot org)
- Adam Fuchs (adam.p.fuchs at ugov dot gov)
- Eric Newton (ecn at swcomplete dot com)
- Billie Rinaldi (billie.j.rinaldi at ugov dot gov)
- Keith Turner (keith.turner at ptech-llc dot com)
- John Vines (john.w.vines at ugov dot gov)
- Chris Waring (christopher.a.waring at ugov dot gov)

Affiliations

- Aaron Cordova, The Interlective
- Adam Fuchs, National Security Agency
- Eric Newton, SW Complete Incorporated
- Billie Rinaldi, National Security Agency
- Keith Turner, Peterson Technology LLC
- John Vines, National Security Agency
- Chris Waring, National Security Agency

Sponsors

- Champion: Doug Cutting

Nominated Mentors

- Benson Margulies
- Alan Cabrera
- Bernd Fondermann
- Owen O'Malley

Sponsoring Entity

- Apache Incubator

AccumuloProposal (last edited 2011-09-09 15:26:16 by OwenOmalley)