TrafodionProposal

Trafodion Apache Incubator Proposal

Abstract

Trafodion is a webscale SQL-on-Hadoop solution enabling transactional or operational workloads on Hadoop.

Proposal

Apache Trafodion builds on the scalability, elasticity, and flexibility of Hadoop. Trafodion extends Hadoop to provide guaranteed transactional integrity, enabling new kinds of big data applications to run on Hadoop. Key features of Apache Trafodion include:

- Full-functioned ANSI SQL language support
- JDBC/ODBC connectivity for Linux/Windows clients
- Distributed ACID transaction protection across multiple statements, tables and rows
- Performance improvements for OLTP workloads with compile-time and run-time optimizations
- Support for large data sets using a parallel-aware query optimizer
- ANSI SQL security and data integrity constraints including referential integrity

Hewlett-Packard Company submits this proposal to donate its Apache License, Version 2.0 open source project known as Trafodion, its source code, documentation, and web site content to the Apache Software Foundation in order to build an open source community

Background

Trafodion is an open source project sponsored by HP, incubated at HP Labs and HP-IT, to develop an enterprise-class SQL-on-Hadoop solution targeting big data transactional or operational workloads. HP publically announced the open source project and uploaded the source code to GitHub in June 2014.

The SQL compiler, optimizer and executor components of Trafodion have a rich heritage. Under development since 1993, they were released as commercial closed source software in various flavors such as HP NonStop SQL/MX and HP Neoview. NonStop SQL/MX was designed for online transaction processing on HP's NonStop (formerly Tandem) fault-tolerant servers and is known for its high availability, scalability, and performance. Hundreds of companies and thousands of servers are running mission-critical applications today on NonStop SQL/MX. In addition, much of these components today are running internal to HP as the core of its Enterprise Data Warehouse (EDW), managing over a PB of data.

Starting in 2013, the software was modified to run on HBase and a new distributed transaction manager was written to run as an HBase co-processor.

Unlike most NOSQL and other SQL-on-Hadoop open source projects, Trafodion provides comprehensive ANSI SQL language support including full-functioned data definition (DDL), data manipulation (DML), transaction control (TCL) and database utility support.

Trafodion provides comprehensive and standard SQL data manipulation support including SELECT, INSERT, UPDATE, DELETE, and UPSERT/MERGE syntax with language options including join variants, unions, where predicates, aggregations (group by and having), sort ordering, sampling, correlated and nested sub-queries, cursors, and many SQL functions.

Utilities are provided for updating table statistics used by the optimizer for costing (i.e. selectivity/cardinality estimates) plan alternatives, for displaying the chosen SQL execution plan, plan shaping, backup and restoring the database, data loading and unloading, and a command line utility for interfacing with the database engine.

Explicit control statements are provided to allow applications to define transaction boundaries and to abort transactions when warranted, including BEGIN WORK, COMMIT WORK, ROLLBACK WORK and SET TRANSACTION.

Trafodion supports ANSI's grant/revoke semantics to define user and role privileges in terms of managing and accessing the database objects.

Rationale

The name "Trafodion" (the Welsh word for transactions, pronounced "Tra-vod-eee-on") was chosen specifically to emphasize the differentiation that Trafodion provides in closing a critical gap in the Hadoop ecosystem. Trafodion builds on the scalability, elasticity, and flexibility of Hadoop. Trafodion extends Hadoop to provide guaranteed transactional integrity, enabling new kinds of big data applications to run on Hadoop.

Current Status

HP released the Trafodion code under the Apache License, Version 2, in June of 2014. Since that time, we have had one major release in January 2015 and one minor release in April 2015. The focus of these releases has been in getting our base functionality, including security, working on top of Apache HBase, as well as improving performance, availability and scalability, and integrating better with HBase.

Meritocracy

We want to build a diverse developer community, based on the Apache Way, around Trafodion. To help developers become contributors, we have documentation on the wiki about the architecture, the source tree structure, and an example enhancement. We plan to publish our project backlog to the community, specifically highlighting areas where developers new to Trafodion may best start contributing, such as extending the database functionality with User Defined Routines (UDRs) and integrating with other Apache projects in the Hadoop ecosystem.

Community

We have already begun building a community but at this time the community consists only of Trafodion developers – all HP employees – and prospective users. We have participated in and hosted HBase Meetups and intend to ramp up our community building efforts.

The Trafodion project has seen interest in China, where HP has conducted proof-of-concepts with multiple companies and expects to see some of its first commercial deployments. To help recruit contributors and users in China, members of the team are translating Trafodion wiki content into Mandarin.

Core Developers

The core developers are very experienced in database and transaction monitor technology, with many having spent more than 20 years working in this space.

Alignment

Apache Trafodion relies on Apache HBase as its storage engine. The development team has collaborated with and gained valuable advice from working with the Apache HBase core developers. Apache Trafodion has federation capabilities as well, and can query Trafodion tables stored in HBase, native HBase tables, and Apache Hive tables.

Known Risks

Orphaned Products

HP Labs and HP-IT have been incubating Trafodion development for almost two years. This is part of HP's strategy to leverage its investment in database software and bring software to market as open source and is similar to HP's efforts with OpenStack. Trafodion builds on HP's equity investment in the Hadoop ecosystem and its efforts to monetize Hadoop through hardware, software, and services. HP wants Trafodion to be successful, as HP will offer a commercially supported distribution of Trafodion.

Inexperience with Open Source

We have been working with open source software in building closed source software for well over two decades. To help transition to doing open source development, the development team received guidance and best practices from HP developers working on OpenStack open source projects, many of whom have experience working on Apache and other open source projects as well. Since releasing Trafodion as an open source project in June of 2014, the committers and contributors have moved forward using open source development processes and tools for bug tracking and design blueprints and Jenkins for continuous integration. As part of the incubation process, we recognize we may need to change some of our development processes/tools and conduct our discussions using Apache email dlists.

Homogenous Developers

Since the initial development of Trafodion has been supported by HP, all of the current developers are HP employees. Through the support of the Apache incubation project, we aim to expand the list of developers and gain contributors from related SQL-on-Hadoop projects and the Apache HBase project. Trafodion developers are experienced with distributed development processes, being primarily based in Palo Alto, CA; Austin, TX; and Shanghai, China. Trafodion is written in C++ and Java.

Reliance on Salaried Developers

Currently all of the developers working on the project are paid by their employer to work on the project. These developers will work on the open source project as well as work on the commercially supported distribution of Trafodion that HP will offer.

Relationship with Other Apache Products

Trafodion is built upon Apache HBase and extends it to support ACID transactions with HBase co-processors for distributed transaction management and recovery. Trafodion envisions future collaborations with the Apache HBase project on performance optimizations, such as in the areas of mixed workload support, High Availability, etc. It also provides transactional support and querying from native HBase tables as well.

Trafodion uses Apache Zookeeper to coordinate and manage the distribution of connection services across the cluster for load-balancing and high availability reconnection purposes in the event a Trafodion process should fail.

Trafodion also envisions working with the Apache Ambari project on enabling better Trafodion manageability. While Ambari focuses on system and component level performance metrics, Trafodion manageability will focus in a complimentary way on database workload monitoring and performance analytics with capabilities more geared towards database administrators.

There are alternative open source projects that are providing SQL-on-Hadoop capabilities, such as Apache Hive, Apache Drill, and Apache Phoenix. These are more focused on reporting and analytics across data structures supported on HDFS. In comparison to all of these technologies Trafodion provides a very complete implementation of ANSI SQL, one of the most sophisticated optimizers for such workloads, a completely parallel data flow architecture that does not materialize intermediate results unless necessary, full ACID transactional support, ANSI GRANT/REVOKE security, and other capabilities that would take decades to build in these products. On the other hand currently Trafodion is just focused on HBase and querying Hive, whereas Hive and Drill provide access to other data formats in HDFS.

An Excessive Fascination with the Apache Brand

We understand the reputation and value of the Apache brand, and no doubt believe that it will help us attract contributors and users. Our primary goal is to follow a proven, open source development and community building model that will make Trafodion successful and enable better collaboration with other Apache projects in the Hadoop ecosystem. We also understand the rules and guidelines about the use of the Apache brand and intend to follow them.

Documentation

Documentation and technical details on Trafodion can be found at: http://www.trafodion.org/

Initial Source

The source is available today in a public github repository: https://github.com/trafodion/trafodion.

Source and Intellectual Property Submission Plan

The source code has already been released under the Apache License, Version 2. The manuals have been released in Adobe PDF format. As part of the submission process, the source for the manuals will be converted from a proprietary DocBook XML format to AsciiDoc.

External Dependencies

Two dependencies do not have Apache compatible licenses and will be addressed as we enter incubation. One dependency is log4cpp, which is licensed under the LGPL. A compatible alternative might be Apache incubator project log4cxx. The other dependency is unixodbc, which is used as the ODBC driver manager. We will look into how Apache Hive manages being able to use this incompatible software and do similar. All other dependencies have Apache compatible licenses, including Apache 2.0, MIT/X11, MIT, and BSD.

Cryptography

Trafodion does not contain any cryptographic code. It does call cryptographic libraries: OpenSSL for C++ code and Java Cryptography Extension (JCE) for Java code.

Required Resources

Mailing Lists

private@trafodion.incubator.apache.org dev@trafodion.incubator.apache.org commits@trafodion.incubator.apache.org

Git Repository

https://git-wip-us.apache.org/repos/afs/incubator-trafodion.git

Issue Tracking

JIRA: JIRA Trafodion (Trafodion)

Initial Committers and Affiliation

- Dave Birdsall, Hewlett-Packard Company, Dave.Birdsall<AT>hp<DOT>com
- Matt Brown, Hewlett-Packard Company, mattbrown<AT>hp<DOT>com
- Tharak Capirala, Hewlett-Packard Company, Tharak.Capirala<AT>hp<DOT>com
- Alice Chen, Hewlett-Packard Company, Alice.Chen<AT>hp<DOT>com
- John DeRoo, Hewlett-Packard Company, John.Deroo<AT>hp<DOT>com
- Roberta Marton, Hewlett-Packard Company, Roberta.Marton<AT>hp<DOT>com
- Amanda Moran, Hewlett-Packard Company, Amanda.Kay.Moran<AT>hp<DOT>com
- Suresh Subbiah, Hewlett-Packard Company, Suresh.Subbiah<AT>hp<DOT>com
- Sandyha Sundaresan, Hewlett-Packard Company,
- Sandhya.Sundaresan<AT>hp<DOT>com

Sponsors

Champion

Michael Stack, Stack<AT>apache<DOT>org

Nominated Mentors

- Andrew Purtell apurtell<AT>apache<DOT>org
- Devaraj Das, ddas<AT>apache<DOT>org
- Enis Söztutar, Enis<AT>apache<DOT>org
- Lars Hofhansl, larsh<AT>apache<DOT>org
- Michael Stack, Stack<AT>apache<DOT>org
- Roman Shaposhnik, rshaposhnik<AT>pivotal<DOT>io

We are seeking additional mentors.

Sponsoring Entity

Apache Incubator PMC