

December2013

⚠ FINAL ⚠

This report has been submitted to the Board. Please do not edit further.

Incubator PMC report for December 2013

Timeline

Wed December 04	Podling reports due by end of day
Sun December 08	Shepherd reviews due by end of day
Sun December 08	Summary due by end of day
Tue December 10	Mentor signoff due by end of day
Wed December 11	Report submitted to Board
Wed December 18	Board meeting

Shepherd Assignments

Alan D. Cabrera	BatchEE
Alan D. Cabrera	Falcon
Andrei Savu	Sentry
Dave Fisher	MRQL
John Ament	Kalumet
John Ament	S4
Justin Mclean	NPanda y
Matthew Franklin	Drill
Matthew Franklin	Sirona
Raphael Bircher	Allura
Raphael Bircher	Storm
Roman Shaposhnik	Aurora
Roman Shaposhnik	Tajo
Ross Gardler	Wave
Suresh Marru	Streams
Suresh Marru	Twill

Report content

Incubator PMC report for December 2013

The Apache Incubator is the entry path into the ASF for projects and codebases wishing to become part of the Foundation's efforts.

There are 35 podlings currently under incubation.

* Community

New IPMC members:

Raphael Bircher
Dave Brondsema
Sergio Fernández
Lars Hofhansl
Steven Noels
Andrew Purtell
Billie Rinaldi
Sebastian Schelter
Michael Stack
Till Westmann

People who left the IPMC:

(None)

* New Podlings

- * Phoenix is an open source SQL query engine for Apache HBase, a NoSQL data store. It is accessed as a JDBC driver and enables querying and managing HBase tables using SQL.
- * Twill is an abstraction over Apache Hadoop YARN that reduces the complexity of developing distributed applications, allowing developers to focus more on their business logic
- * Log4cxx -- Logging for C++. This new podling is a reboot of an existing Apache project with a fresh group of contributors.

* Graduations

The board has motions for the following:

Helix

* Retirements

- * The Provisionr podling voted to retire and then followed through on completing the retirement guide (a rarity). Best wishes to the wonderful, wonderful Provisionr people in their future endeavors.

* Releases

The following releases were made since the last Incubator report:

Nov 11 Apache Allura (incubating) 1.0.1
Nov 13 Apache VXQuery 0.2-incubating
Nov 14 Apache Tajo 0.2.0-incubating
Nov 22 Apache Helix 0.6.2-incubating
Nov 22 Apache Helix 0.7.0-incubating
Dec 01 Apache Tez 0.2.0-incubating

It took 0-39 days for the third IPMC vote to arrive.

Release	VOTE start	Third IMPC +1	Days
Apache Allura (incubating) 1.0.1	Oct 14	Nov 10	27
Apache VXQuery 0.2-incubating	Oct 3	Nov 10	39
Apache Tajo 0.2.0-incubating	Nov 1	Nov 14	13
Apache Helix 0.6.2-incubating	Nov 15	Nov 22	7
Apache Helix 0.7.0-incubating	Nov 15	Nov 20	5
Apache Tez 0.2.0-incubating	Nov 26	Nov 26	0

* IP Clearance

- * Citrix has donated a plugin for Apache CloudStack to support RDP connectivity for virtual machines running within Hyper-V hypervisors.
- * IBM has donated Distributed UIMA Cluster Computing (DUCC), a cluster management system providing tooling, management, and scheduling facilities to automate the scale-out of applications using the UIMA

framework.

- * Adobe has donated the Sling replication module, which allows resources to be replicated between Sling instances.

* Infrastructure

- * The Incubator dist area is now fully transitioned to dist.apache.org. Infra is pleased.

* Miscellaneous

- * The Incubator experienced a rough cycle for releases. Several podlings experienced long wait times for IPMC votes. One release made it out despite IP clearance not having been completed. Another was improperly released with only two IPMC +1 votes.
- * The initiative to reform release voting continues. We achieved consensus for one particular approach and tried it. The approach seems not to have succeeded (though some argue it has not been given enough time). We are now negotiating a second approach. We have not yet achieved a breakthrough, but progress is being made and discussions are civil and proceeding at a measured pace rather than overheating as past efforts have.
- * We've started recruiting non-IPMC shepherds and the new volunteers have brought substantial energy and enthusiasm to the task.
- * The Chair finished coding up several new tools around reporting which will hopefully make it easier to sustain high report quality into the future. An attempt to recruit a "Report Manager" for this cycle did not yield any volunteers, though.

----- Summary of podling reports -----

* Still getting started at the Incubator

Twill
Usergrid

* Not yet ready to graduate

No release:

Aurora
Storm
Wave

Community growth:

Allura
Drill
Falcon
Kalumet
MRQL
NPanday
S4
Sentry
Streams
Tajo

* Ready to graduate

The Board has resolutions for the following:

Helix

* Did not report, expected next month

BatchEE
Sirona

Table of Contents

Allura
Aurora
BatchEE
Drill
Falcon
Kalumet
MRQL
NPanday
S4
Sentry
Sirona
Storm
Streams
Tajo
Twill
Usergrid
Wave

Allura

Forge software for the development of software projects, including source control systems, issue tracking, discussion, wiki, and other software project management tools.

Allura has been incubating since 2012-06-25.

Three most important issues to address in the move towards graduation:

1. Continue to grow community
2. Make another release
3. Move project development to ASF hardware

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

It was difficult to get responses from IPMC and our mentors, for our first release. It finally did get the needed responses and votes, and the IPMC has been discussing general improvements in this area, which is good.

How has the community developed since the last report?

Several new folks on the mailing list have been asking questions about setup and configuration of Allura. A few contributions by Rui Ferreira have been submitted through the tracker but not merged yet.

How has the project developed since the last report?

Our legal/notice files were finalized and we made our first release. Many fixes & features have been added, primarily regarding importing projects from other sites into Allura.

Date of last release:

2013-11-11

When were the last committers or PMC members elected?

2013-04-11

Signed-off-by:

[](allura) Ross Gardler
[](allura) Greg Stein
[X](allura) Jim Jagielski

Aurora

Aurora is a service scheduler used to schedule jobs onto Apache Mesos.

Aurora has been incubating since 2013-10-01.

Three most important issues to address in the move towards graduation:

1. First release.
2. Establish a consistent development process and release rhythm.
3. Expanding the community and adding new committers.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

- None at this time.

How has the community developed since the last report?

- Finish bootstrapping project, IP clearance, initial website.
- Announcement and live demo at first #MesosTownhall.

How has the project developed since the last report?

- Code has been approved and migrated to Apache, complete with (scrubbed) commit history.
- Development work has continued inside of Twitter, but will now begin moving to the open.

Date of last release:

- No releases as of yet. Targeting first release for mid-January.

When were the last committers or PMC members elected?

- No new committers; still stabilizing code and preparing for first release.

Signed-off-by:

[x](aurora) Jake Farrell
[x](aurora) Benjamin Hindman
[](aurora) Chris Mattmann
[x](aurora) Henry Saputra

Shepherd notes:

Jake Farrell (jfarrell):

For Aurora everything is progressing nicely. All initial setup and on boarding is completed and development has started using ASF infrastructure (jira/reviewboard/etc). No blockers or issues at this time.

BatchEE

FIXME

BatchEE has been incubating since 2013-10-03.

Shepherd notes:

Marvin Humphrey (marvin):

Podling did not report. Please report next month.

Also, please edit podlings.xml metadata and address all "FIXME" tags.

Apache: Project Drill

Apache Drill is a distributed system for interactive analysis of large-scale datasets that is based on Google's Dremel. Its goal is to efficiently process nested data, scale to 10,000 servers or more and to be able to process petabytes of data and trillions of records in seconds.

Drill has been incubating since 2012-08-11.

Three Issues to Address in Move to Graduation:

1. Continue to attract new developers and and early users with a variety of skills and viewpoints
2. Continue to develop deeper community skills and knowledge by building additional releases
3. Demonstrate community robustness by rotating project tasks among multiple project members.

The community has made significant progress on items 1 and 2.

Issues to Call to Attention of PMC or ASF Board:

none

How community has developed since last report:

Community awareness and participation were strengthened through a meeting of the Bay Area Apache Drill User Group with over 100 participants locally in San Jose and remotely via Cisco-hosted Webex. On site speakers included 3 Drill contributors, two from San Jose and one from Seattle
<http://www.meetup.com/Bay-Area-Apache-Drill-User-Group/>

Additional events include:

- * Code for 1st milestone release was posted and made available via project website; release was socialized via mailing list, Twitter, blogs and presentations
- * Several new full-time developers joined the project
- * Apache Drill received a Bossie award "Best open source big data tools" 17 September 2013 <http://s.apache.org/B3H> (infoworld.com)

Mailing list discussions:

Subscriptions to the Drill mailing lists have risen to 415. There has been active and increasing participation in discussions on the developer mailing list, including new participants and developers. Participation on the user list is growing although still small; mainly activity takes place on developer mailing list.

Activity summary for the user mailing list:

http://mail-archives.apache.org/mod_mbox/incubator-drill-user/
December to date 12/5/2013: 14
November 2013, 28
October 2013, 37
September 2013, 13

Topics in discussion on the user mailing list included but not limited to:

How to load data into Drill; use of distributed mode; direction to study the src

Activity summary for the dev mailing list:

http://mail-archives.apache.org/mod_mbox/incubator-drill-dev/
* December to date 12/6/2013: 35 (mainly jira; some discussion)
* November 2013, 347 (jira, focused discussions)
* October 2013, 299 (jira; focused discussions)
* September 2013, 659 (jira, focused discussions)

Recent topics on the dev mailing list have included:

- * Buffer allocation of cast into var length type.
- * Improve Parquet read performance.
- * Schema discovery tool for scanning raw files and generating optiq schema.
- * Discussions on CAST functionality.
- * sqlline connect to Remote Drill
- * Limit operator end-to-end
- * many focused discussions about 1st milestone release

Code

Total Commits from 01/09/2013 - 06/012/2013: 106

Detailed information regarding the commits is shown in this chart from GitHub: <http://s.apache.org/MnI> (github.com)

Nine contributors have participated in this GitHub code activity; there have been 101 forks of the Apache Drill project on GitHub which is a good indicator of strong interest outside the group of core contributors.

Code commits (during the period 01/09/2013 - 06/012/2013) include but not limited to following:

- * Spooling batch buffer
- * Fix over memory pre-allocation within ParquetRecordReader.
- * Implement simple metrics framework
- * prepare release drill-1.0.0-m1.
- * Implement builders for Scan, Sort, LogicalPlan and PlanProp.

Community Interactions

The weekly Drill hangout continues, conducted remotely through Google hangouts Tuesday mornings 9am Pacific Time to keep core developers in contact in realtime despite geographical separation. A Gdoc is being updated regarding the discussions during hangout.
<http://s.apache.org/4Gc> (docs.google.com)

The community stays in touch through @ApacheDrill Twitter ID, and by postings on various blogs including Apache Drill User <http://drill-user.org/> which has had several updates and through international presentations at conferences.

Viability of the community is also apparent as participants in the open source Apache Drill community came together on November 4th meet-up of the Bay Area Apache Drill User Group, with 391 members enrolled. The group looked at how Drill works now and what will be the next steps in the project. The event marked the recent first official release of the Apache Drill project.

Presentations

- * A talk on Apache Drill by Michael Hausenblas is scheduled on Tue 10 Dec "Query engine for heterogenous large scale datasets" at Decemberi Big Data Meetup in Budapest
<http://www.meetup.com/Big-Data-Meetup-Budapest/events/138089032/>
- * WebEx of talks by Jacques Nadeau, Tim Chen and Steven Phillips at Bay Area Drill User Group; play video:
<http://s.apache.org/eN> (cisco.webex.com)
- * Drill talk by Michael Hausenblas at JAX London "Large-scale, interactive ad-hoc queries over different data stores with Apache Drill" including demo, 29 October 2013
- * Podcast by Jacques Nadeau at All Things Hadoop, episode 17: "Using Apache Drill for Large Scale, Interactive, Real-Time Analytic Queries"
<http://s.apache.org/8ZM> (allthingshadoop.com)
- * Drill talk by Michael Hausenblas at the Stockholm HUG "Interactive analytics for large-scale data-sets"
<https://speakerdeck.com/mhausenblas/hug-stockholm-apache-drill>

Sample Articles/presentations (out of many):

- * How to use Apache Drill (inc distributed mode) Detailed description is available via Drill Github sandbox.
<https://github.com/mhausenblas/apache-drill-sandbox/tree/master/M1>
- * Interactive analytics: large scale data-set
<https://speakerdeck.com/mhausenblas/hug-stockholm-apache-drill>
- * Lifetime of a Query in Drill by Timothy Chen, includes link to his slides. <http://s.apache.org/INZ> (tnachen.wordpress.com)
- * Drill blog post by Ellen Friedman at Apache Drill User reports on meet-up of the Bay Area Apache Drill User Group
<http://s.apache.org/a8r> (drill-user.org)
- * Drill article posted by Ted Dunning on MapR Technologies blog site 8 November 2013: "Apache Drill Achieves 1st Milestone Release"
<http://www.mapr.com/blog?s=Apache+Drill>
- * Blog post by Yash Sharma on "How to Contribute to Apache Drill: Implementing Drill Math Functions"
<http://s.apache.org/od> (confusedcoders.com)

Social Networking

@ApacheDrill Twitter entity is active and has grown by ~44%, to 632 followers.

How project has developed since last report:

1. Approval of first release of Apache Drill - M1 achieved; code posted via project website
2. Query now works in distributed mode
3. Significant code drops have been checked in from a number of developers to achieve Milestone 2 release.
4. New developers are contributing.

Signed-off-by:

[x](drill) Ted Dunning
[x](drill) Grant Ingersoll
[x](drill) Sebastian Schelter

Shepherd notes:

Matt Franklin (mfranklin):

I have reviewed the community and have concluded that it is very active and healthy. IMO, #3 in the issues to address before graduation is not critical and can be successfully mitigated by documenting the tasks that are currently being executed by single individuals. My recommendation is that the Drill podling begin the graduation process after completion of the next release.

Marvin Humphrey (marvin):

The report is admirably thorough and must have taken a long time to prepare. Perhaps consider that a report with a reduced level of detail similar to the other reports may be more in line with Board expectations.

The report was filed too late to be incorporated into a review by the assigned shepherd. Please file on time.

Please wrap future reports at 77 columns, and please take more care with making indentation reflect logical hierarchy. Also, use only s.apache.org URL shorteners in future reports. Otherwise, someone downstream must clean up the report -- either someone from the IPMC (I took care of reformatting this month) or the Board -- before it is published in the official Board minutes.

Falcon

Falcon is a data processing and management solution for Hadoop designed for data motion, coordination of data pipelines, lifecycle management, and data discovery. Falcon enables end consumers to quickly onboard their data and

its associated processing and management tasks on Hadoop clusters.

Falcon has been incubating since 2013-03-27.

Three most important issues to address in the move towards graduation:

1. Add new and diverse committers
2. Build and grow community

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

- No

How has the community developed since the last report?

- * We have seen more contributions
- * We have had Google hangouts to meet with users and contributors on the project

How has the project developed since the last report?

- * We are preparing ourselves for the next release of Falcon (version 0.4-incubating). We are in the process of testing & verifying the release candidate.

Date of last release:

2013-08-15 (0.3-incubating)

When were the last committers or PMC members elected?

At time of incubation. No new committers or PMC members elected since.

Signed-off-by:

[](falcon) Arun Murthy
[X](falcon) Chris Douglas
[](falcon) Owen O'Malley
[X](falcon) Devaraj Das
[X](falcon) Alan Gates

Kalumet

Kalumet a complete environment manager and deployer including J2EE environments (application servers, applications, etc), softwares, and resources.

Kalumet has been incubating since 2011-09-20.

Three most important issues to address in the move towards graduation:

1. Prepare Kalumet 0.7.0-incubating release
2. Publish the documentation of the website
3. Refactore some modules

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

No

How has the community developed since the last report?

Unfortunately, we didn't have a lot of activity during the last two months.

We count on the first release to gather more people on the project (contributors and users).

How has the project developed since the last report?

We released Apache Kalumet 0.6-incubating release on November, 22, 2013.

We are preparing the publication of the release, publication of the documentation on the website, and announcement e-mails.

Date of last release:

2013/11/22

When were the last committers or PMC members elected?

None since the beginning of the podling.

Jim (Jagielski) decided to resign as mentor on 2013/07/01.

Signed-off-by:

[](kalumet) Henri Gomez
[X](kalumet) Jean-Baptiste Onofre
[](kalumet) Olivier Lamy

Shepherd notes:

John Ament (johndament):

Kalumet seems to be OK, but a little bit slow. They have been incubating for a little over 2 years now, I think if they can get together on another release shortly they should consider graduating, there was some confusion for a little while over IPMC vs PPMC votes which caused a delay in getting their 0.6 release out the door.

MRQL

MRQL is a query processing and optimization system for large-scale, distributed data analysis, built on top of Apache Hadoop, Hama, and Spark.

MRQL has been incubating since 2013-03-13.

Three most important issues to address in the move towards graduation:

1. recruit more users, developers, committers and PMCers
- 2.
- 3.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

none

How has the community developed since the last report?

We have released our first release.

How has the project developed since the last report?

Extended the interface between the query translator and the run-time evaluation engine to make query evaluation on Spark more efficient.

Date of last release:

2013-10-31

When were the last committers or PMC members elected?

2013-03-13

Signed-off-by:

[X](mrql) Alan Cabrera

[](mrql) Anthony Elder
[](mrql) Alex Karasulu
[](mrql) Mohammad Nour El-Din

Shepherd notes:

Dave Fisher (wave):

This podling achieved one benchmark towards successful incubation - a release.

When commits and mailing list activity are reviewed this looks like a tiny community with only a single developer and some administrative support.

I am not sure if this another example of a community that is mostly from a single corporation and is doing all of the work outside and contributing it through one person with off-list decision making and communication. If so then the mentors need to look into it and get the real decision making process in public.

There is a user ML, but the only email ever on that list was the release announcement.

NPanday

NPanday allows projects using the .NET framework to be built with Apache Maven.

NPanday has been incubating since 2010-08-13.

The podling continues to have a low amount of activity. However, there are still committers involved that wish to see it succeed. A recent prod from the Incubator PMC has reignited some interest, including from new contributors and existing users of the project.

Unfortunately, we have gradually lost mentors, and the last one resigned in November. Thankfully, Raphael Bircher has volunteered to step in. There are ASF members on the project, however we could still use other independent mentors willing to keep an eye on the project and keep prodding.

In the current state, we might find it difficult to get enough votes together for a release.

We have requested that JIRA notifications be sent to the dev@ list instead of commits@, so we can keep as much discussion in one place as possible.

With the recent activity, we have started to identify a list of things that need to be done to get activity moving again, and the next release out. This needs to be fleshed out, including details of who is willing to contribute to them.

The last release was made on 16 May 2011.

Signed-off-by:

[x](npanday) Raphael Bircher

Shepherd notes:

Justin Mclean (jmclean):

Last report was July 2013. At noted in the current report there are some issues to overcome for the project to move forward.

In the last month there been some increased activity on the mailing lists (thanks to the last shepherd) but as yet no fixes or patches have been submitted. The project currently has only one mentor (just appointed) and the PPMC seems to have gone missing.

Ripple

Ripple is a browser based mobile platform development and testing tool.

Three most important issues to address in the move towards graduation:

1. Fix licensing issues, specifically some of the images.
2. Build up the community.
3. Make a release under Apache.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

No

How has the community developed since the last report?

We are still seeing increased non-committer participation in threads, issues, etc.

However, aside from the odd contribution, most of the recent work has been committer based.

How has the project developed since the last report?

Intel has been doing a lot of work on the project (in a fork)- still waiting for a possible contribution back into the ASF project.
<http://software.intel.com/en-us/html5/tools>

Continuing Cordova 3.0 work to get to a stable, releasable product.

The Chrome Extension build target was removed altogether from the project, in an effort to simplify and streamline the building/compilation of the project into a usable application, and because the project has been continually moving towards a [more feature-able] NodeJS backend, capable of being run in any Browser.

Date of last release?

N/A

When were the last committers or PMC members elected?

When we were inducted into Apache.

Signed-off-by:

[](ripple) Jukka Zitting
[X](ripple) Christian Grobmeier
[](ripple) Andrew Savory

S4

S4 (Simple Scalable Streaming System) is a general-purpose, distributed, scalable, partially fault-tolerant, pluggable platform that allows programmers to easily develop applications for processing continuous, unbounded streams of data.

S4 has been incubating since 2011-09-26.

Three most important issues to address in the move towards graduation:

1. No development since last release
2. Limited activity on mailing lists
3. Need new committers/contributors

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

There was a discussion related to retirement but there was some interest from the community to continue the project. However there were no contributions since last release.

How has the community developed since the last report?

No activity on dev mailing list over last couple of months

How has the project developed since the last report?

No major commits since last release. One of the reason

Date of last release:

June 2013: S4 0.6.0 released.

When were the last committers or PMC members elected?

Was reported by March 2013 (Daniel Gómez Ferro).

Signed-off-by:

[X](s4) Patrick Hunt
[](s4) Arun Murthy

Shepherd notes:

John Ament (johndament):

The board report reflects my sentiments as well. S4 seems to be in a bit of rut. I tried kicking off some conversations on the dev mailing list, no luck. It seems like there are at best five active participants, between the users list and dev list. Considering that there hasn't been a commit since last board report, it doesn't come off as a good sign for me. I think retirement may be an option to start exploring.

Sentry

Sentry is a highly modular system for providing fine grained role based authorization to both data and metadata stored on an Apache Hadoop cluster.

Sentry has been incubating since 2013-08-08.

Three most important issues to address in the move towards graduation:

1. Build community and add new committers
2. Continue to release at regular intervals

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

None

How has the community developed since the last report?

Community development continues. Sentry held its first key signing meetup for committers. A number of committers participated.

How has the project developed since the last report?

A number of new JIRAs have been filed and are being actively worked on. Support for Solr index level authorization is complete and available on trunk. A new test framework to run hive integration tests on a real cluster is available as well.

Date of last release:

2013-09-25

When were the last committers or PMC members elected?

Last committers and PMC members were elected when the project entered incubation.

Signed-off-by:

[X](sentry) Arvind Prabhakar
[](sentry) Joe Brockmeier
[](sentry) David Nalley
[](sentry) Olivier Lamy
[X](sentry) Patrick Hunt
[](sentry) Thomas White

Sirona

Monitoring Solution.

Sirona has been incubating since 2013-10-15.

Shepherd notes:

Marvin Humphrey (marvin):

Podling did not report. Please report next month.

Storm

Storm is a distributed, fault-tolerant, and high-performance realtime computation system that provides strong guarantees on the processing of data.

Storm has been incubating since 2013-09-18.

Three most important issues to address in the move towards graduation:

1. Licensing and IP clearance.
2. First Apache release.
3. Growing the community.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

Storm has a number of dependencies that are forked/modified versions of other projects (including Apache Thrift). We have started a discussion on the dev mailing list regarding the best way to resolve these issues. As we have yet to do our first Apache release, we will be looking for additional guidance from our mentors in reaching that goal.

How has the community developed since the last report?

We continue to see interest/involvement from contributors in the form of github pull requests.

How has the project developed since the last report?

We have released Storm 0.9.0 from the non-Apache github process. This will be the last non-Apache release. Our next focus will be to work on our first Apache release and migrating to the Apache Way.

Date of last release:

N/A

When were the last committers or PMC members elected?

Upon entry into the Incubator.

Signed-off-by:

[X](storm) Ted Dunning
[X](storm) Arvind Prabhakar
[X](storm) Devaraj Das
[X](storm) Matt Franklin
[X](storm) Benjamin Hindman

Shepherd notes:

Raphael Bircher (rbircher)

Storm has started three months ago. It is quite active and looks ok, so far. They created a last release outside Apache and they work now at the IP. I noticed that <http://incubator.apache.org/storm/> still not exist. If you search at Google for storm, you only find the GitHub page. I find a minimal site on the Apache server important. so people can easely find ML and co.

Streams

Apache Streams is a lightweight server for ActivityStreams.

Streams has been incubating since 2012-11-20.

Three most important issues to address in the move towards graduation:

1. Growing the active developer community
2. Making our second incubating release
- 3.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

How has the community developed since the last report?

Received and incorporated community code contributions from Carol Hansen for Streams Web interface.

How has the project developed since the last report?

Switched to using Apache Cassandra for default persistent storage.
Very good community based architectural discussions, like about using Apache Camel vs Web Services vs SpringMVC, possible integration of Apache Storm, etc.
Evaluation of alternatives done in separate branch from trunk.

Date of last release:

2013-01-09

When were the last committers or PMC members elected?

2013-07-18 (Danny Sullivan)

Signed-off-by:

[X](streams) Matt Franklin
[X](streams) Ate Douma
[](streams) Craig McClanahan

Tajo

Tajo is a distributed data warehouse system for Hadoop.

Tajo has been incubating since 2013-03-07.

Three most important issues to address in the move towards graduation:

1. IP Clearance check
2. More committers
- 3.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

There are few of active mentors due to their personal reasons. We need more mentors.

How has the community developed since the last report?

- * We had the first meetup in US through Bay Area Hadoop User Group.
- * Mailing traffic grows significantly (e.g., Oct - 659 and Sep - 564)
- * More contributors are participating in Tajo project.
- * There is a usecase report in SK telecom, the biggest Korea telco.

How has the project developed since the last report?

- * Added JDBC driver
- * Added Hive-compatible RCFile with Text/Binary (De)Serializer
- * Added more SQL standards functions
- * Allowed more SQL features (inline views, between predicate, in predicate)
- * Enabled TajoCatalog to access Hive meta store.

Date of last release:

2013-11-20

When were the last committers or PMC members elected?

Not yet

Signed-off-by:

[](tajo) Chris Mattmann
[](tajo) Owen O'Malley

Twill

Twill is an abstraction over Apache Hadoop® YARN to reduce the complexity of developing distributed applications.

Twill has been incubating since 2013-11-14.

Top three items to resolve before graduation:

- More committers from different organizations.
- IP clearance.
- Regular Releases.

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

- None at this time.

How has the community developed since the last report?

- This is the first report since incubation.
- One new contributor since incubation.

How has the project developed since the last report?

- Bootstrapping is complete.
- Code has been renamed from weave to twill, migrated to org.apache packages, and imported into git.

Date of last release:

- No releases as of yet. Targeting first release for January.

What are the plans for the next period?

- Create the website
- Initial Release

When were the last committers or PMC members elected?

- No new committers since incubation.

Signed-off-by:

[] (twill) Vinod K
 [] (twill) Arun C Murthy
 [] (twill) Tom White
 [X] (twill) Patrick Hunt
 [] (twill) Andrei Savu

 Usergrid

Usergrid is an open-source Backend-as-a-Service ("BaaS" or "mBaaS") composed of an integrated distributed NoSQL database, application layer and client tier with SDKs for developers looking to rapidly build web and/or mobile applications.

Usergrid has been incubating since 2013-10-03.

Three most important issues to address in the move towards graduation:

1. Learning the Apache Way
2. Growing a diverse community
3. Getting project infrastructure and codebase setup at Apache

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

- None at this time

How has the community developed since the last report?

- We have established Contributor Workflow Policy for cutting releases from the Apache Git repo and accepting Pull Requests via GitHub (<<http://s.apache.org/QTQ>>)

How has the project developed since the last report?

- Our project infrastructure is now in place
- We have imported our source code to the Apache Git repo

Date of last release:

- No releases yet.

When were the last committers or PMC members elected?

- When we entered incubation

Signed-off-by:

[] (snoopdave) Dave Johnson
 [] (jfarrell) Jake Farrell
 [X] (jim) Jim Jagielski
 [] (lewismc) Lewis John McGibbney

[x](lresende) Luciano Resende

Wave

A wave is a hosted, live, concurrent data structure for rich communication. It can be used like email, chat, or a document.

Wave has been incubating since 2010-12-04.

Three most important issues to address in the move towards graduation:

1. Make the initial release
2. Extend the community

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of?

No

How has the community developed since the last report?

The community is in a relatively quiet state. There's is a certain amount of interest in the project but seems like steep learning curve along with tightly coupled client-server architecture poses obstacles for new developers to begin actively contributing.

How has the project developed since the last report?

Mainly worked on the next release candidate and sorting out related legal issues.

Date of last release:

Not yet

When were the last committers or PMC members elected?

The last 2 committers were elected around 30th June 2013

Signed-off-by:

[X](wave) Christian Grobmeier
[X](wave) Upayavira

Shepherd notes:

Christian Grobmeier (grobmeier):

It was discussed to end graduation because lack of momentum, which led to a lot of emails. So far there are a couple of people around but to less who do actual coding work. It was discussed briefly to push development via a Kickstarter project. So far there is no outcome of the discussions.