

February2014

⚠️ FINAL ⚠️

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

Incubator PMC report for February 2014

Timeline

Wed February 05	Podling reports due by end of day
Sun February 09	Shepherd reviews due by end of day
Sun February 09	Summary due by end of day
Tue February 11	Mentor signoff due by end of day
Wed February 12	Report submitted to Board
Wed February 19	Board meeting

Shepherd Assignments

Alan D. Cabrera	Sirona
Alan D. Cabrera	log4cxx2
Andrei Savu	Knox
Andrei Savu	NPanday
Dave Fisher	DataFu
Dave Fisher	Hadoop Development Tools
John Ament	Droids
Justin Mclean	Open Climate Workbench
Matthew Franklin	Phoenix
Matthew Franklin	Tez
Raphael Bircher	Usergrid
Roman Shaposhnik	Twill
Ross Gardler	DeviceMap
Suresh Marru	BatchEE
Suresh Marru	Blur

Report content

Incubator PMC report for February 2014

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 36 podlings currently under incubation.

* Community

New IPMC members:

Suresh Srinivas

Carl Steinbach

People who left the IPMC:

Dennis Lundberg

* New Podlings

(None)

* Graduations

The board has motions for the following:

Open Climate Workbench

At time of writing, graduation votes were underway for the following:

Knox
Spark

* Releases

The following releases were made since the last Incubator report:

Jan 21 Apache Metamodel 4.0.0-incubating

It took 9 days for the third IPMC vote to arrive.

Release	RC VOTE start	Third IMPC +1	Days
Apache Metamodel 4.0.0-incubating	Dec 12	Dec 21	9

* IP Clearance

* Salesforce donated the code for Phoenix, 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.

* Miscellaneous

* The issue of how best to integrate with GitHub is a recurring topic of discussion. Infra added a new feature -- forwarding of comments on pull requests to dev lists -- which was well received. However, as GitHub's interface is an external commercial service rather than something we run on our own hardware, it is difficult to ensure that communications will be archived, requiring additional vigilance from PMCs.

* A proposal for Hoya, an application to deploy and manage existing distributed applications in a YARN cluster, was discussed at length but has not yet been voted on.

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

* Still getting started at the Incubator

DataFu

* Not yet ready to graduate

No release:

Hadoop Development Tools
log4cxx2
Phoenix
Usergrid

Community growth:

Blur
Tez

Twill

Low activity:

Droids

* Ready to graduate

Open Climate Workbench
Knox

* Did not report, expected next month

BatchEE
DeviceMap
NPanday
Sirona

Table of Contents

Blur
DataFu
DeviceMap
Droids
Hadoop Development Tools
Knox
log4cxx2
Open Climate Workbench
Phoenix
Tez
Twill
Usergrid

Blur

Blur is a search platform capable of searching massive amounts of data in a cloud computing environment.

Blur has been incubating since 2012-07-24.

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

1. We received the Trademark assign document from trademarks@ and are pursuing getting it signed by Near Infinity.
2. Another Release
3. Another Release

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 continue to be small but active.
- Subscriptions: user@ - 54[+3]; dev@ - 61[+1]

How has the project developed since the last report?

The majority of effort has been around bug fixes that were needed. The code is now solid and tested at scale now so we anticipate a bug fix release soon.

Date of last release:

2013-10-09

When were the last committers or PMC members elected?

2013-08-30

Signed-off-by:

[](blur) Doug Cutting
[X](blur) Patrick Hunt
[X](blur) Tim Williams

DataFu

DataFu provides a collection of Hadoop MapReduce jobs and functions in higher level languages based on it to perform data analysis. It provides functions for common statistics tasks (e.g. quantiles, sampling), PageRank, stream sessionization, and set and bag operations. DataFu also provides Hadoop jobs for incremental data processing in MapReduce.

DataFu has been incubating since 2014-01-05.

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

1. Building ASF community
2. Release
3. Remaining incubator paperwork

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?

Since initial incubation, have received contributions from two new contributors.

How has the project developed since the last report?

First report. Have obtained all the necessary infra (git/jira/wiki,etc). Thirty JIRAs have been opened, 14 have been closed. Active discussion on mailing list as to community development, etc.

Date of last release:

None. First month of incubation.

When were the last committers or PMC members elected?

None. First month of incubation.

Signed-off-by:

[](datafu) Ashutosh Chauhan
[X](datafu) Roman Shaposhnik
[](datafu) Ted Dunning

Shepherd/Mentor notes:

Dave Fisher (wave):

New community to the incubator just getting started. Good guidance from Mentors. Needs Apache trademark attribution on site. Should have links to Mailing lists on the site.

DeviceMap

Shepherd/Mentor notes:

Marvin Humphrey (marvin):

A draft report was posted to the DeviceMap mailing list on Friday February 7th for review, but it was never added to the official report.

Droids

Droids aims to be an intelligent standalone robot framework that allows to create and extend existing droids (robots).

Droids has been incubating since 2008-10-09.

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

1. Activity

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

None, beside the real low traffic on mailing list and source code commits.

How has the community developed since the last report?

No changes to community.

How has the project developed since the last report?

The project has been very quiet / dormant for the past half year.

Date of last release:

2012-10-15

When were the last committers or PMC members elected?

2012-05-07

Signed-off-by:

[X](droids) Thorsten Scherler
[X](droids) Richard Frovarp

Shepherd/Mentor notes:

John Ament (johndament):

Droids is a long standing podling, with a robust code base but without much need for frequent releases. I think graduating into an existing TLP (not their own TLP) would be a good resolution, it helps avoid some of the over head of many binding votes to carry on while still maintaining an area to work in. The people working on droids were fast to respond and give feedback on what's happening, but like is mentioned the list is not very active (in the past 6 months, the main conversations going on were around status of the project).

Hadoop Development Tools

Eclipse based tools for developing applications on the Hadoop platform

Hadoop Development Tools has been incubating since 2012-11-09.

1. we have to finish the release of version 0.0.1, the main issue being resolving source and binary license and notice files.
2. we have to define and implement an effective approach to manage the "client connections" to multiple clusters.
3. Build Community

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?

- Srimanth Gunturi joined as committer and PPMC member

How has the project developed since the last report?

- Two additional 0.0.1 release candidates as we worked through issues
- Jenkins job setup

Date of last release:

- None

When were the last committers or PMC members elected?

- November 2013

Signed-off-by:

[](hadoopdevelopmenttools) Suresh Maru
[](hadoopdevelopmenttools) Chris Mattmann
[X](hadoopdevelopmenttools) Roman Shaposhnik

Shepherd/Mentor notes:

Dave Fisher (wave):

I don't see any on dev list activity since their release in November was held up due to dependency issues and LICENSE issues.

Are these guys talking somewhere else?

Knox

Knox Gateway is a system that provides a single point of secure access for Apache Hadoop clusters.

Knox has been incubating since 2013-02-22.

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

1. Expand community to include more diverse committers.
2. Review and complete all graduation readiness items.
3. Begin the process!

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

1. None

How has the community developed since the last report?

1. Two new committers added.
2. More user questions on user list trickling in.

How has the project developed since the last report?

1. Internally agreed to being graduation processes.
2. Working toward a 0.4.0 release.
3. A few IP issues found and resolved.
4. Resolved 142(+42) of 242(+52) total issues currently in JIRA.

Date of last release:

0.3.0: 2013-10-13

When were the last committers or PMC members elected?

2013-12-19: Added Vladimir Tkhir as new committer
2013-12-19: Added Maksim Kononenko as new committer

Signed-off-by:

[](knox) Owen O'Malley
[X](knox) Chris Douglas
[](knox) Mahadev Konar
[X](knox) Alan Gates
[](knox) Devaraj Das
[](knox) Chris Mattmann
[X](knox) Tom White

log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

1. We probably need to decide what's an acceptable platform/compiler list, sort out any remaining hiccups, and then move towards a release.
2. We need to publish the project site with updated content to reflect the new status and fix documentation bugs.
3. We need to make a broader use of the ASF infrastructure (notably the CI, as well as improved management of the Jira site - notably with triaging), and establish a roadmap for the next releases.

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

We are stuck on INFRA-7209 and how to properly publish the project site.

How has the community developed since the last report?

The committers are still getting acquainted with each other and the foundation. There have been some discussions about how to develop within the team of committers already, like using a CI, how to deal with issues regarding a changelog, which fixes should be covered by our first incubation release etc.

Moderation of the list seems to work, there were some support requests which were answered and problems introduced by commits on trunk discussed.

How has the project developed since the last report?

Some known bugs were fixed already and available patches applied, which is work in progress for the new release. It has been suggested to schedule most of the remaining available patches for an upcoming release of 0.11.0 to include as much content which has been provided over the years as possible.

Date of last release:

2008-04-03 was the official, pre-incubation 0.10.0

Many post-0.10.0 commits exist on trunk which we intend to get out as 0.11.x.

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](log4cxx2) Christian Grobmeier

[](log4cxx2) Scott Deboy

Open Climate Workbench

Shepherd/Mentor notes:

Justin Mclean (jmclean):

Report not submitted. Project is up for graduation at board meeting.

Phoenix

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.

Phoenix has been incubating since 12/11/2013.

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

- 1) Perform initial release out of apache incubator followed up quickly by two new releases
- 2) Make our customers successful and prove value of Phoenix+HBase over other alternatives
- 3) Attract new committers to the project

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

The below is informational for the IPMC (no need to forward the board):

We've been waiting for our Github issues to be imported into our Apache JIRA but it is taking a while. We're considering doing the import ourselves, as we're starting to see duplicate JIRAs and continue to have to point users back to our old Github issues.
[[<https://issues.apache.org/jira/browse/INFRA-7179>|INFRA-7179]]

How has the community developed since the last report?

We've cut over from our former Github-based open source project to our new Apache home and dev and user lists are busy.

No new committers.

How has the project developed since the last report?

- * Apache Brand Management has ok'd our name of Apache Phoenix [1]
- * IP Clearance is nearly complete: software grant has been submitted and acknowledged. Vote is underway and will close on 2/8. IP-CLEARANCE note has been posted on general.
- * Reviewing proposed release on dev list and if all looks good, will start a vote for our first release next week.
- * Working in parallel toward a new 3.0.0 release as well.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

No new committers or PMC members.

1. <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-44?filter=-2>

Signed-off-by:

[x](phoenix) Lars Hofhansl
[X](phoenix) Andrew Purtell

[X](phoenix) Devaraj Das
[](phoenix) Enis Soztutar
[X](phoenix) Steven Noels

Tez

Tez is an effort to develop a generic application framework which can be used to process arbitrarily complex data-processing tasks and also a re-usable set of data-processing primitives which can be used by other projects.

Tez has been incubating since 2013-02-24.

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

1. Growing the community
- 2.
- 3.

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?

No new PPMC members or committers added since the last report. We have had interest from both new users and contributors via mailing lists and JIRA. The Hive project recently merged its native support for Tez into their trunk branch. The Pig community is also working on integrating Tez and we have seen various JIRAs filed as well as some patches from Pig developers.

How has the project developed since the last report?

Since the first week of November, 2013 (last report), we have had around 180+ jiras filed and 130+ jiras resolved. We made the first release of Tez (version 0.2.0) in the first week of December, 2013. There is currently ongoing work to make Tez work on a secure Hadoop cluster, after which we plan to do a 0.3.0 release.

Date of last release:

2013-12-03

When were the last committers or PMC members elected?

At project incubation. No new PPMC members or committers added since the last report.

Signed-off-by:

[](tez) Alan Gates
[](tez) Arun Murthy
[X](tez) Chris Douglas
[](tez) Chris Mattmann
[x](tez) Jakob Homan
[](tez) 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.
- Regular Releases.

- Increase adoption.

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?

- Activity on the dev list up from 141 messages in December to 173 in January.
- Subscribers to the dev list up from 24 to 31.
- One new contributor in January.

How has the project developed since the last report?

- IP clearance was completed.
- Web site is up and running.
- First release is currently under vote on the incubator list.

Date of last release:

- First release is under vote and expected to happen 2/6/14.

What are the plans for the next period?

- Establish a cadence of monthly releases
- Engage more members of the community to contribute actively
- Improve documentation and website

When were the last committers or PMC members elected?

- No new committers since incubation.

Signed-off-by:

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

Shepherd/Mentor notes:

Roman Shaposhnik (rvs):

A brand new community off to a good start

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?

- A couple of new names have appeared on dev@usergrid which is positive.

How has the project developed since the last report?

- Have imported JIRA issues from pre-Apache project JIRA
- Work is continuing on a new "Core Persistence" module
- Almost done with Java package renaming to org.apache.usergrid

Date of last release:

- No releases yet.

When were the last committers or PMC members elected?

- Alex Karasulu, January 13, 2014

Signed-off-by:

- [] (snoopdave) Dave Johnson
- [X] (jfarrell) Jake Farrell
- [X] (jim) Jim Jagielski
- [X] (lewismc) Lewis John Mcgibbney
- [] (lresende) Luciano Resende