

April2015

/!\ *FINAL* (!)

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

Incubator PMC report for April 2015

Timeline

| | |
|--------------|------------------------------------|
| Wed April 01 | Podling reports due by end of day |
| Sun April 05 | Shepherd reviews due by end of day |
| Sun April 05 | Summary due by end of day |
| Tue April 07 | Mentor signoff due by end of day |
| Wed April 08 | Report submitted to Board |
| Wed April 15 | Board meeting |

Shepherd Assignments

| | |
|--------------------|-------------|
| Andrei Savu | DataFu |
| Andrei Savu | Lens |
| Drew Farris | Usergrid |
| Justin Mclean | ODF Toolkit |
| Konstantin Boudnik | OpenAz |
| Matthew Franklin | Parquet |
| P. Taylor Goetz | Myriad |
| Raphael Bircher | CommonsRDF |
| Ross Gardler | Ripple |
| Suresh Marru | TinkerPop |
| Timothy Chen | Calcite |
| John D. Ament | Groovy |

Report content

Incubator PMC report for April 2015

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 40 podlings currently in incubation. There is another climate project in discussion for incubation which is an interesting continuation of the nascent trend of academic projects coming to Apache. This may be a singleton due to Chris Mattman's efforts, but it would be an interesting development if it gathers momentum.

There is a new ASFIncubator twitter entity. It has little activity so far, but has had a lot of followers jump on after Sally's announcement on @ASF.

Datafu and ODF Toolkit are showing very low activity. This isn't necessarily a serious problem yet, but may be a leading indicator. Wave has previously had a warning, but may be doing a little better.

* Community

New IPMC members:

Andrew Bayer
Gary Gregory

People who left the IPMC:

None

* New Podlings

Groovy
SINGA
(Myriad was vote result was sent on March 1st, but the project is just barely beginning)

* Graduations

None

* Releases

The following releases were made since the last Incubator report:

2015-03-03 apache-lens-2.0.1-beta-incubating
2015-03-16 nifi-0.0.2-incubating
2015-03-23 apache-calcite-1.1.0-incubating
2015-03-24 incubator-ignite-1.0.0-rc3
2015-03-27 slider-0.70.1-incubating

* IP Clearance

Usergrid is working feverishly to finish up their IP Clearance.
Updated SGA sent to secretary@a.o

CommonsRDF SGA filed 2015-03-27.

Zeppelin SGA submitted and acknowledged on 2015-03-20.

* Legal / Trademarks

The old debate about how to submit SGAs for ALv2 projects with huge numbers of copyright holders reappeared with Groovy's SGA. In the end, the SGA was signed by a single community member.

Ambiguity around Legal Policy language on when optional GPL library dependencies are permitted caused confusion.

* Infrastructure

The "Marvin" report reminders have not run in a timely manner for podlings for 4 months in a row. A dedicated Incubator-specific replacement is being contemplated.

* Miscellaneous

An official Twitter account for the Incubator was initialized but has not yet been used. Access will be maintained via TweetDeck.

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

* Still getting started at the Incubator

AsterixDB
CommonsRDF
Groovy
Myriad
OpenAZ

SINGA

* Not yet ready to graduate

No release:

Parquet (some components released but not others)
Ripple
TinkerPop
Zeppelin

Community growth:

Calcite
Lens
Usergrid
Wave

* Did not report, expected next month

DataFu
ODF Toolkit

Table of Contents

AsterixDB
Calcite
CommonsRDF
DataFu
Groovy
Lens
Myriad
ODF Toolkit
OpenAz
Parquet
Ripple
SINGA
TinkerPop
Usergrid
Wave
Zeppelin

AsterixDB

Apache AsterixDB is a scalable big data management system (BDMS) that provides storage, management, and query capabilities for large collections of semi-structured data.

AsterixDB has been incubating since 2015-02-28.

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

1. Get SGA signed from UCI and others.
2. Code drop and start actively using lists.
3. Do an Apache release.

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

No, still getting started.

How has the community developed since the last report?

Account setup is mostly done and code repositories are in place. Right now working mostly on the SGA to do the code drop.

How has the project developed since the last report?

Account and other infrastructure setup mostly done.
SGA and code drop still open.

Date of last release:

No Releases yet.

When were the last committers or PMC members elected?

No elected PMC and/or committers

Signed-off-by:

[X] (asterixdb) Chris Mattmann
[X] (asterixdb) Henry Saputra
[] (asterixdb) Ted Dunning
[X] (asterixdb) Ate Douma
[] (asterixdb) Jochen Wiedmann

Calcite

Calcite is a highly customizable engine for parsing and planning queries on data in a wide variety of formats. It allows database-like access, and in particular a SQL interface and advanced query optimization, for data not residing in a traditional database. (Renamed from Optiq on 2014-09-30.)

Calcite has been incubating since 2014-05-19.

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

1. Continue to build community
2. Add more committers and PPMC members
3. Demonstrate community robustness by rotating project tasks among multiple project members

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?

In March, there were 159 messages on dev@, 41 commits, 52 issues created, 37 issues resolved by 8 contributors.

Julian Hyde gave a talk "Why you care about relational algebra (even though you didn't know it)" about Calcite at Enterprise Data World in March. He also spoke at a Phoenix meet-up in San Francisco about Phoenix-Calcite integration.

Collaborations are under way with other Apache projects:

- * Phoenix is going to use Calcite for SQL parsing and planning; and a separate project is using Calcite's Avatica component to build a remote JDBC driver.
- * Added support for parsing/validating streaming SQL and collaborating with the Samza project on adding streaming SQL to Samza.
- * Collaboration with Hive and Drill continues.

How has the project developed since the last report?

We completed our milestone 1.0 release, continued our monthly release tempo with release 1.1, and are preparing 1.2.

Date of last release:

2015-03-23

When were the last committers or PMC members elected?

December 2014.

Signed-off-by:

[x](calcite) Ted Dunning
[X](calcite) Alan Gates
[X](calcite) Steven Noels
[X](IPMC) Chris Mattmann

Shepherd/Mentor notes:

Chris Mattmann:

* Calcite being discussed by Tom Barber at ApacheCon NA 2015.

CommonsRDF

Commons RDF is a set of interfaces and classes for RDF 1.1 concepts and behaviours. The commons-rdf-api module defines interfaces and testing harness. The commons-rdf-simple module provides a basic reference implementation to exercise the test harness and clarify API contracts.

CommonsRDF has been incubating since 2015-03-06.

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

1. Grow the CommonsRDF Community such that we can progress towards graduation
2. Further engage the podling with the nuances of the Apache Incubation process
3. Drive towards the first incubating release

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

A SGA was recently filed with secretary@ and the CommonsRDF codebase now resides at the ASF with the canonical source residing at <https://git-wip-us.apache.org/repos/asf?p=incubator-commonsrdf.git>

Project is not yet decided on goal of TLP or Commons component - but still aiming at the Commons component.

How has the community developed since the last report?

There have been no new committers, however the overwhelming majority of project-related development conversation is already taking place on the Apache lists.

How has the project developed since the last report?

CommonsRDF very recently joined the Apache Incubator. Bootstrapping and migration have quickly progressed (mailing lists, git, Jira) and the initial committer community is transitioning very well into the Apache incubator ecosystem.

Web site remains to be sorted - evaluating technologies (considering INFRA's proposed CMS deprecation).

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

N/A Commons RDF has only just joined the Apache Incubator.

Signed-off-by:

[X](commonsrdf) Rob Vesse
[X](commonsrdf) John D Ament
[](commonsrdf) Gary Gregory
[X](commonsrdf) Lewis John McGibbney
[X](IPMC) Chris Mattmann

Shepherd/Mentor notes:

Chris Mattmann:

* template dates can be removed.

Groovy

Apache Groovy is a multi-faceted programming language for the Java platform. Groovy is a powerful, optionally typed and dynamic language, with static-typing and static compilation capabilities, aimed at multiplying developers' productivity thanks to a concise, familiar and easy to learn syntax. It integrates smoothly with any Java program, and immediately delivers to your application powerful features, including scripting capabilities, Domain-Specific Language authoring, runtime and compile-time meta-programming and functional programming.

Groovy has been incubating since 2015-03-24

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

1. All the ICLA of all the PMC / Committers have been sent. Most Apache accounts have already been created.
2. On behalf of the whole Groovy development team and the community, Guillaume Laforge signed and scanned the SGA, and sent it to Secretary on April 1st.
3. We've communicated to the historical mailing-lists about the new lists, encouraging subscribers to join the new ones. The JIRA migration should be taking place over the week-end, and we'll also communicate with our community about that change. And last but not least, we need the Git repository to be created and import our current sources from Github.

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

No, just getting started.

How has the community developed since the last report?

Currently, the pretty large Groovy community is happy about the move into the ASF, but continues to be centered around our old infrastructure. That's why we're going to start attracting the community towards our new lists, to get them involved in this new era for the project.

How has the project developed since the last report?

No last report, as this is the first!

Date of last release:

We did a last non-Apache release as the vote for incubation was still ongoing, but we haven't yet planned the next release dates.

When were the last committers or PMC members elected?

No new elected PMC or committers.

Signed-off-by:

[X](groovy) Bertrand Delacretaz
[X](groovy) Emmanuel Lecharny
[X](groovy) Jim Jagielski
[X](groovy) Roman Shaposhnik
[](groovy) Andrew Bayer
[X](groovy) Konstantin Boudnik

Shepherd/Mentor notes:

Bertrand Delacretaz:

Thanks Guillaume for taking the lead on this!

Emmanuel Lecharny:

Quite a good start for a podling... Even the user mailing list starts to be active ! Waiting for the JIRA migration this week-end, and the code migration asap, then the podling will be almost fully setup, as soon as the web site will be available. Great work by the Groovy fellows, and the Infra peeps !

John D. Ament:

Project seems to be going well initially, doing the expected up front house keeping required considering the largish size of the code base. I would expect email volume to pick up quickly as the Codehaus services begin to shutdown.

Lens

Lens is a platform that enables multi-dimensional queries in a unified way over datasets stored in multiple warehouses. Lens integrates Apache Hive with other data warehouses by tiering them together to form logical data cubes.

Lens has been incubating since 2014-10-10.

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

1. Grow user and developer community and add more committers
2. Document all features and do one more release

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?

- Apache Lens talk accepted at ApacheCon and GIDS, 2015
- We saw user queries coming up on user mailing list regarding usage
- Establish whether "Apache Lens" is a suitable name : CLOSED through PODLINGNAMESEARCH-63
- Subscriptions:
 - dev@: 37
 - user@: 29
 - commits: 17

How has the project developed since the last report?

- The first release: 2.0.1-beta-incubating has been published
- Added improvements in the project to reduce time taken for rest api response
- Add features for metering, partition timeline cache, namespace level jars
- Mailing list activity :
 - * Commits@ : 266 in feb, 61 in March
 - * dev@ : 1512 in feb, 965 in March
 - * user@ : 18 in feb, 1 in March
- Issues :
 - * 30 Day Summary (as on 2nd March) 116 created and 81 resolved
 - * 30 day Summary (as on 29th March) 98 created and 78 resolved

Date of last release:

2.0.1-beta-incubating on 2015-03-04

When were the last committers or PMC members elected?

2014-12-24

Signed-off-by:

(lens) Christopher Douglas
 (lens) Jakob Glen Homan
 (lens) Jean-Baptiste Onofre

Myriad

Myriad enables co-existence of Apache Hadoop YARN and Apache Mesos together on the same cluster and allows dynamic resource allocations across both Hadoop and other applications running on the same physical data center infrastructure.

Myriad has been incubating since 2015-03-01.

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

1. Get ebay to donate the code to ASF. Ebay is willing and has done this before. The process is started, but will take time.
2. Get ICLAs from all committers and give them Apache accounts and karma.
3. Integrate with Apache infrastructure and make our first release.

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

None. Just getting started.

How has the community developed since the last report?

First report. Now we have mailing lists and discussions have started on dev@. In addition to committers/contributors from the 3 initial organizations (ebay, MapR, Mesosphere), we have 5 more contributors from at least two different organizations.

Gave presentations at Strata/Hadoop World and the March Bay Area Mesos Users Group and London MUG meetups. Upcoming presentations at ApacheCon, Data Science MD, and East Bay Big Data Group.

How has the project developed since the last report?

First report. Recent commits/PRs include:

1. Mesos framework authentication
2. Gradle builds, checkstyle, and Spock unit testing
3. In progress: Remote binary distribution
4. In progress: High availability for the Scheduler/RM
5. In progress: Dockerizing the Scheduler and Executor

Date of last release:

No releases yet. Still at 0.0.1 in the source.

When were the last committers or PMC members elected?

Still working on getting ICLAs for initial committers. Then we will vote in new committers from the existing developer community.

Signed-off-by:

(myriad) Benjamin Hindman
 (myriad) Danese Cooper
 (myriad) Ted Dunning
 (myriad) Luciano Resende

Shepherd/Mentor notes:

P. Taylor Goetz (ptgoetz):

The Myriad podling is still ramping up.

ODF Toolkit

Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents

Shepherd/Mentor notes:

Justin Mclean (jmclean):

Did not report. Low activity on the mailing list but code still being worked on, some discussion about graduation but community seems to be lacking motivation to do so. May need some help.

Chris Mattmann:

* report missing?

OpenAz

OpenAz is a project to create tools and libraries to enable the development of Attribute-based Access Control (ABAC) Systems in a variety of languages. In general the work is at least consistent with or actually conforming to the OASIS XACML Standard.

OpenAZ has been incubating since 01-2015

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

1. Get ICLA's from all the committers. So far AT&T and JP Morgan have done, but Oracle is taking more time.
2. Merging the source code from AT&T, JP Morgan and Oracle and creating POM's. AT&T and JP Morgan source code has been merged and POM's created, it needs uploading into the Apache repository.
3. Get our first release done.

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

None. Just getting started.

How has the community developed since the last report?

First report. Now we have mailing lists and discussions have started on dev@. Moving all communications onto the new mailing lists.

How has the project developed since the last report?

First report. Will be committing initial codebase by the end of the week April 10th.

Date of last release:

No releases yet. Still at 0.0.1 in the source.

When were the last committers or PMC members elected?

Still working on getting ICLAs for initial committers. Then we will vote in new committers from the existing developer community.

Signed-off-by:

[](openaz) Paul Fremantle
[](openaz) Emmanuel Lecharny
[](openaz) Colm O Heigeartaigh
[](openaz) Hadrian Zbarcea

Shepherd/Mentor notes:

Konstantin Boudnik:

The podling is in the initial stages: nothing to shepherd on just yet.

Marvin Humphrey:
Report was filed late.

Parquet

Parquet is a columnar storage format for Hadoop.

Parquet has been incubating since 2014-05-20.

Three most important issues

- 1st releases toward org.apache Parquet 1.6.0 GA
- Expanding the community and adding new committers
- Ensuring timely code reviews by committers, developing reviewers

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

- None at this time

Latest Additions:

- * PMC addition: None
- * Committer addition: Dan Weeks and Alex Levenson (from last report)

Issue backlog status since last report:

- * Created: 34
- * Resolved: 50

Mailing list activity since last report:

- * dev 560 messages: 111 in Jan, 136 in Feb, and 313 in Mar

How has the project developed since the last report?

- Preparing last commits for the first parquet-mr release candidate
- Planned parquet-mr 1.6.0 release schedule
- ASF required changes to parquet-mr are finished
- Released parquet-format 2.3.0, with org.apache packages
- Parquet presentation at Strata 2015 San Jose and the Presto meetup

Date of last release:

- parquet-format 2.3.0 released 19 Feb
- Not yet released: parquet-mr and parquet-cpp

Signed-off-by:

- [](parquet) Todd Lipcon
- [X](parquet) Jake Farrell
- [X](parquet) Chris Mattmann
- [X](parquet) Roman Shaposhnik
- [](parquet) Tom White

Ripple

Ripple is a browser based mobile phone emulator designed to aid in the development of HTML5 based mobile applications. Ripple is a cross platform and cross runtime testing/debugging tool. It currently supports such runtimes as Cordova, WebWorks and the Mobile Web.

Ripple has been incubating since 2012-10-16.

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

1. Making a new release of Ripple

2. Contributors and community engagement

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

Working on making a release for Ripple

How has the community developed since the last report?

Added 3 new committers

- Ilya
- Tim Barham
- Parashuram N

How has the project developed since the last report?

Working on a new release. We started a new vote, but it was reverted due to issues. A new package has been generated and we are in [DISCUSS] mode

Date of last release:

N/A

When were the last committers or PMC members elected?

March 2015

Signed-off-by:

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

Shepherd/Mentor notes:

Christian Grobmeier:

The project faces more activity recently.

SINGA

SINGA is a distributed platform for deep learning that is easy to use, extendible and scalable.

SINGA has been incubating since 2015-03-17.

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

1. Grow the community of developers and users.
2. Finish moving the project to Apache infrastructure.
3. Add examples, documents, tests and benchmarks for the first release.

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

N/A.

The project Incubation Status file is NOT yet up to date

The current project URL is: www.comp.nus.edu.sg/~dbsystem/singa

What has been done since the last report?

N/A. This is the first report.

Other project activities:

1. The Github repository now has 12 forks as developers are contributing features and doing testing.

2. SINGA is being used for several scientific papers being under submission.

What are the plans and expectations for the next period?

1. Finish moving the project to the Apache's ecosystem.
2. Focus on the scalability feature, documentation and testing.

Date of last release:

N/A.

Signed-off-by:

- (SINGA) Daniel Dai
- (SINGA) Alan Gates
- (SINGA) Ted Dunning
- (SINGA) Thejas Nair

TinkerPop

TinkerPop is a graph computing framework written in Java

TinkerPop has been incubating since 2015-01-16.

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

1. Releasing TinkerPop 3.0.0.M8 out under Apache Incubator.
2. Getting more active committers on the project.
3. Growing the homepage and documentation presence.

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 added Daniel Kuppitz as a Committer. A long time Gremlin designer and promoter.

How has the project developed since the last report?

We have fully migrated to Apache infrastructure and are very close to our first Apache release.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

2 weeks ago.

Signed-off-by:

- (tinkerpop) Rich Bowen
- (tinkerpop) Daniel Gruno
- (tinkerpop) Hadrian Zbarcea
- (tinkerpop) Matt Franklin
- (tinkerpop) David Nalley

Shepherd/Mentor notes:

Daniel Gruno:

While I am happy that a new committer has been invited to the group, it seems that the project is struggling with the how-to of the ASF. We need to further educate the PPMC on our procedures.

Usergrid

Usergrid is Backend-as-a-Service (BaaS) composed of an integrated database (Cassandra), application layer and client tier with SDKs for developers.

Usergrid has been incubating since 2013-10-03.

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

1. Continue to grow the Usergrid Incubating Community
2. Clarify the SGA for Apache Usergrid Incubating codebase
3. Drive towards another Usergrid incubating release

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

The Usergrid community has been working to clarify the status of the SGA for the Usergrid codebase. The original SGA listed a GitHub repo that did not include all of the source code that Apigee intended to grant. We are preparing a new SGA that will cover all code that is currently part of the project.

How has the community developed since the last report?

In short, the Usergrid incubating community is healthy with loads of mailing list and development traffic from a growing community. Usergrid continues to see new community members making themselves known on the mailing lists.

How has the project developed since the last report?

Development is very dynamic. The codebase is being actively developed with correspondence taking place on the Usergrid JIRA instance. The processes and workflows for development are pretty solid now.

Date of last release:

2015-01-21

When were the last committers or PMC members elected?

John D. Ament (johndament) 05 Jan 2015

Signed-off-by:

[X](usergrid) Dave Johnson
[X](usergrid) Jake Farrell
[X](usergrid) Jim Jagielski
[X](usergrid) John D. Ament
[X](usergrid) Lewis John Mcgibbney
[](usergrid) 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. Wave is very close to making its first release. This is the top priority.
2. The project needs to determine what the next few technical goals are.
3. The project needs to attract more developers and users.

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?

The community has been slow to develop. There are still many people downloading and using the software, we have seen several bug reports and questions, but there has not been much contribution by new members.

How has the project developed since the last report?

The code for the release has largely been finalized. We believe that major legal and technical issues have been overcome. Only a few final steps remain. The project believes that we are finally ready to make a release.

Date of last release:

N/A

When were the last committers or PMC members elected?

July 2014

Shepherd/Mentor notes:

Christian Grobmeier:

The mentors have missed to sign this report last time. For that reason I am adding it again. Since last month there was some discussion on how to prevent the project going to the attic.

Upayavira:

We have received a GSoC application, and a committer has accepted the role of mentor. The task involves simplification of the codebase which will make it more accessible to others, hopefully facilitating greater involvement, as it seems the greatest impediment to Wave is the (over)complexity of its codebase.

Signed-off-by:

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

Zeppelin

A collaborative data analytics and visualisation tool for distributed, general-purpose data processing systems such as Apache Spark, Apache Flink, etc.

Zeppelin has been incubating since 2014-12-23.

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

1. Complete PODLINGNAMESEARCH-64.
2. Finish issues migration to ASF JIRA.
3. First Apache release.

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?

First vote on a new committer happened this month.

At least 2 new contributors actively participating on the mailing lists\sending patches regular. There is also a fair amount of discussion @dev on the feature project direction.

How has the project developed since the last report?

SGA was submitted and acknowledged.
Source code have been migrated to ASF infrastructure (with GH mirror).

Work on running interpreters in a separate processes is finished.
Python interpreter was implemented with a pyspark support.
Scatter chart visualisation added.
Support of Spark 1.3.0.
Support of Debian (.deb) package.

Date of last release:

None

When were the last committers or PMC members elected?

2015-03-27 first new commiter\PPMC vote passed.

Signed-off-by:

(zeppelin) Konstantin Boudnik
 (zeppelin) Henry Saputra
 (zeppelin) Roman Shaposhnik
 (zeppelin) Ted Dunning
 (zeppelin) Hyunsik Choi