

July2016

/\ *FINAL* (!)

Incubator PMC report for July 2016

Timeline

Wed July 06	Podling reports due by end of day
Sun July 10	Shepherd reviews due by end of day
Sun July 10	Summary due by end of day
Tue July 12	Mentor signoff due by end of day
Wed July 13	Report submitted to Board
Wed July 20	Board meeting

Shepherd Assignments

Drew Farris	Mynewt		
Drew Farris	iota		
John Ament	Gossip		
John Ament	Tephra		
Justin Mclean	Pony Mail		
P. Taylor Goetz	DataFu		
P. Taylor Goetz	HORN		
Timothy Chen	Climate Model Diagnostic Analyzer		
Timothy Chen	Geode		
Timothy Chen	Rya		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="dc3e76bc-9b2f-4921-99d2-d455411a7993"><ac:plain-text-body><![CDATA[[none]	Airflow]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d91ed8a3-47bc-432a-9732-ef979ae6ae46"><ac:plain-text-body><![CDATA[[none]	Batch EE]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d903a70b-3c6b-4de0-a57a-12a7056e42d1"><ac:plain-text-body><![CDATA[[none]	CarbonData]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="bbbfcfec7-e051-4270-b649-69f09d81286c"><ac:plain-text-body><![CDATA[[none]	Fluo]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a22dda30-8cdf-4cac-a6f6-586269fa9dee"><ac:plain-text-body><![CDATA[[none]	FreeMarker]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d39bc091-6e25-42dc-8955-23473bc14680"><ac:plain-text-body><![CDATA[[none]	Gearpump]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="342c7603-a89b-4ade-974e-561016dcd4a0"><ac:plain-text-body><![CDATA[[none]	HAWQ]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="453fdae9-c69f-44f2-bfde-294e1d0f9cf2"><ac:plain-text-body><![CDATA[[none]	Juneau]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="bb0d23c9-eac2-4163-9e8b-0000266efc15"><ac:plain-text-body><![CDATA[[none]	MADlib]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="452e016e-1d6a-4a92-b55a-a85d90676892"><ac:plain-text-body><![CDATA[[none]	Metron]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5d24de44-541e-4918-b91e-fa96aa8f97c4"><ac:plain-text-body><![CDATA[[none]	Milagro]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f8c5c605-eef2-4dcc-9eab-30c7da356b96"><ac:plain-text-body><![CDATA[[none]	ODF Toolkit]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="2c8a1ff5-3993-4dad-90b4-3b8308de6616"><ac:plain-text-body><![CDATA[[none]	Open Az]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="01c42831-3878-46e6-b027-446322f26158"><ac:plain-text-body><![CDATA[[none]	Pirk]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="27826033-0bc6-4b17-a3f0-94f5f9ae7af6"><ac:plain-text-body><![CDATA[[none]	PredictionIO]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ec12df5b-26ae-4863-90ed-ba5f1cb468c6"><ac:plain-text-body><![CDATA[[none]	Quickstep]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for July 2016

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 58 podlings currently undergoing incubation. We have one podling planning to graduate this month and added four podlings to the roster. June was a busy month, adding two IPMC member, and completing twelve releases, in addition to other activities.

* Community

New IPMC members:

- Jochen Wiedmann
- Tim Ellison

People who left the IPMC:

- N/A

* New Podlings

- CarbonData
- DistributedLog
- Juneau
- Pirk

* Graduations

The board has motions for the following:

- Kudu

* Releases

The following releases entered distribution during the month of June:

- 2016-06-01 Apache Range 0.5.3-incubating
- 2016-06-03 Apache Trafodion 2.0.0-incubating
- 2016-06-06 Apache Mynewt 0.9.0-incubating
- 2016-06-06 Apache Omid (incubating) 0.8.2.0
- 2016-06-09 Apache Kudu (incubating) 0.9.0
- 2016-06-10 Apache SystemML 0.10.0-incubating
- 2016-06-14 Apache Beam 0.1.0-incubating
- 2016-06-22 Apache Myriad 0.2.0-incubating
- 2016-06-26 Apache FreeMarker 2.3.25-incubating
- 2016-06-28 Apache Slider 0.91.0-incubating
- 2016-06-30 Apache Taverna 3.1.0-incubating
- 2016-06-30 Apache Kudu (incubating) 0.9.1

* IP Clearance

- N/A

* Legal / Trademarks

- An audit was conducted of podling websites to see how far off they are

compared to the Incubator's branding guide. Eleven podlings were found to be in complete violation and requests to correct have been issued. Most have already corrected the issue.

- Justin has provided a summary of building NOTICE and LICENSE files in the form of a screencast for podlings to use as reference.

* Infrastructure

- There are open questions around the usage of GitHub issues for issue tracking, it is unclear if it is a valid tool to use.

* Miscellaneous

- A vote was started to retire the CMDA podling. It was later canceled, giving them another three months to try to work through community issues. Public mailing list participation has increased as a part of this result.
- Discussions are resuming to retire the OpenAz podling, as only one of the three volunteers who stepped up to resume has availability to work on the podling.
- Batchee missed getting a report in. We will follow up in the next IPMC report.

* Credits

- Report Manager(s): Marvin Humphrey, John Ament, Ted Dunning

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

* Still getting started at the Incubator

- CarbonData
- DistributedLog
- Fluo
- Gossip
- Juneau
- Pirk
- Pony Mail
- PredictionIO

* Not yet ready to graduate

No release:

- Airflow
- HAWQ
- Milagro

Community growth:

- FreeMarker
- Geode
- MADlib
- Metron
- Mynewt
- Rya
- Tephra

* Ready to graduate

The Board has motions for the following:

- Kudu

* Did not report, expected next month

- CMDA
- OpenAz
- Quickstep

Table of Contents

Airflow
BatchEE
CarbonData
DataFu
Fluo
FreeMarker
Gearpump
Geode
Gossip
HAWQ
HORN
iota
Juneau
MADlib
Metron
Milagro
Mynewt
ODF Toolkit
Pirk
Pony Mail
PredictionIO
Rya
Tephra

Airflow

Airflow is a workflow automation and scheduling system that can be used to author and manage data pipelines.

Airflow has been incubating since 2016-03-31.

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

1. Getting an Apache release out.
2. Podling name search has to happen.
3. Verify distribution rights

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

We've done a lot of work on distribution rights (3), but we need to wrap it up.

We've been pushing out our release due to stability concerns and the need for a testing framework

How has the community developed since the last report?

- Since our last podling report 1 month ago, we grew our contributors from 148 to 160.
- In June 2016 we had 56 commits.
- We organized a meet up at WePay. We also had one before at Airbnb.
- We had several committers meetings.
- 1 new committer. 2 offers were given but 1 was declined.

How has the project developed since the last report?

- Migrated all files to contain Apache 2.0 license headers.
- Replaced charting library that was non-Apache 2.0 compliant (a large effort).
- Release Airflow incubator website, and added appropriate branding and disclaimers.

Date of last release:

None (had a non-Apache release on 2016-06-13)*

* Release was discussed and approved here:

https://mail-archives.apache.org/mod_mbox/incubator-general/201605.mbox/%3CEFD08AF-80DD-4614-BC74-BBCED9251BE1@gmail.com%3E

When were the last committers or PMC members elected?

2016-05-20

Signed-off-by:

[X](airflow) Chris Nauroth
[x](airflow) Hitesh Shah
[x](airflow) Jakob Homan

CarbonData

Apache CarbonData is a new Apache Hadoop native file format for faster interactive query using advanced columnar storage, index, compression and encoding techniques to improve computing efficiency, in turn it will help speedup queries an order of magnitude faster over PetaBytes of data.

CarbonData has been incubating since 2016-06-02.

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

1. Finalize code cleanup and code
2. Prepare releases
3. Grow up 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?

It's the first CarbonData report

How has the project developed since the last report?

We created the resources (git, github integration, Jira, ...).
The code donation has been done, and the first PR merged.

We are in the process of website creation (CWiki requested) and creating the Jenkins CI jobs.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](carbondata) Henry Saputra
[X](carbondata) Jean-Baptiste Onofre
[](carbondata) Uma Maheswara Rao G

Shepherd/Mentor notes:

Jean-Baptiste Onofre:

This report is the first CarbonData report. The activity is focused on the resources creation.

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. Grow user and contributor base
2. Increased committer activity
3. Continued releases

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?

- * Eyal Allweil was voted in as the newest committer and member of the PPMC.

How has the project developed since the last report?

- * A new UDF provided by Eyal was committed and another was submitted for review.
- * ASF-associated signing key committed in prep for next release, addressing feedback from previous release.

Date of last release:

2015-11-14

When were the last committers or PMC members elected?

July 2016 (Eyal Allweil)

Signed-off-by:

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

Fluo

Fluo is a distributed system for incrementally processing large data sets stored in Accumulo.

Fluo has been incubating since 2016-05-17.

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

1. Attract new contributors
2. Do a release
3. Conduct a podling name search

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 a new contributor to Fluo since the last report.

How has the project developed since the last report?

We have moved our repos (Fluo, Fluo Recipes, and Fluo Website) to Apache infrastructure. Our website is now hosted by Apache (<https://fluo.incubator.apache.org>) and was updated to reflect our incubation status. We are working on refactoring our package names, build tooling, and website, to reflect the Apache branding.

Date of last release:

Never

When were the last committers or PMC members elected?

Never

Signed-off-by:

[x](fluo) Billie Rinaldi
[x](fluo) Drew Farris
[x](fluo) Josh Elser

Shepherd/Mentor notes:

Josh Elser:

The transition to Apache has been slow but steady. In the past week, it seems like things are "returning to normal". Interactions with a new contributor have been good -- encouraging to community growth.

FreeMarker

FreeMarker is a template engine, i.e. a generic tool to generate text output based on templates. FreeMarker is implemented in Java as a class library for programmers.

FreeMarker has been incubating since 2015-07-01. The FreeMarker project exists and produces releases since around 2001.

Most important issues to address in the move towards graduation:

1. Discuss the project's graduation on the general list.
2. Fill blanks left on the status page.
3. Reaching a bigger community that will contribute code to the project, though it remains to be seen if it's required before graduation.

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?

There were no changes

How has the project developed since the last report?

The project has made a new public release, which come with new features, bug fixes, and documentation improvements. Discussion was started about the possibility of graduation. Some missing information on the status page was filled, and a formal "podling name search" was started. There was also progress related to tooling, though those projects are outside the ASF (Eclipse plugin at JBoss, and the Online Template Tester at Kenshoo).

Date of last release:

2016-06-26

When were the last committers or PMC members elected?

2015-10-21 Nan Lei, committer (non-PMC)

Signed-off-by:

[X](freemarker) Jacopo Cappellato
[](freemarker) Jean-Frederic Clere
[X](freemarker) David E. Jones
[X](freemarker) Ralph Goers
[X](freemarker) Sergio Fernández

Gearpump

Gearpump is a reactive real-time streaming engine based on the micro-service Actor model.

Gearpump has been incubating since 2016-03-08.

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

1. Make initial Apache branded release and release on a regular schedule.
2. Continue to evolve community interest and support.
3. Integrate within Apache Beam and akka-streams frameworks.

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?

Increased contribution from developers to provide integration with Redis, Storm and Kubernetes.

How has the project developed since the last report?

- Close to first Apache release.
- Improved code coverage.
- New logo selected.
- Continued integration within Apache Beam (Gearpump Runner). PR under review.
- Continued integration within akka-streams (Gearpump Materializer). First PR committed.
- 18 issues created and 14 resolved

Date of last release:

No release yet

When were the last committers or PMC members elected?

No new committers or PMC members elected

Signed-off-by:

[X](gearpump) Andrew Purtell
[](gearpump) Jarek Jarcec Cecho
[](gearpump) Reynold Xin
[](gearpump) Todd Lipcon
[](gearpump) Xuefu Zhang

Geode

Geode is a data management platform that provides real-time, consistent access to data-intensive applications throughout widely distributed cloud architectures.

Geode has been incubating since 2015-04-27.

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

- Rename packages to org.apache
- Expanding the community to include contributors and committers outside of Pivotal.
- Establish a release cadence.

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 continue to work on the "Three most important issues to move towards graduation" from our previous report.
- We had our second public release (1.0.0-incubating.M2).
- We added 1 new committer.
- We are about to have our third release (1.0.0-incubating.M3).
- The Geode Clubhouse hosted a meeting covering the following topics :
 - "Cluster Replication via WAN" by Mike Stolz
 - "Southwest Airlines' Experience with Geode & GemFire at Scale" by Brian Dunlap
- The breakdown of JIRA tickets is the following :
 - Q3/2015 306 created, 138 resolved
 - Q4/2015 342 created, 225 resolved
 - Q1/2016 438 created, 359 resolved
 - Q2/2016 451 created, 322 resolved [As of June 28th, 2016]
- There was a total of 169 pull requests on Github with 6 still open.
- The breakdown of the mailing lists messages for April, May, June of 2016:
 - org.apache.geode.issues 3,992
 - org.apache.geode.commits 3,572
 - org.apache.geode.dev 2,006
 - org.apache.geode.user 273
- There are now 155 subscribers on the dev and 165 on the user mailing list

How has the project developed since the last report?

- Had our second release, 1.0.0-incubating.M2, on April 22, 2016.
- Donation of native client components which included a C++ client and a .NET client.
- Inclusion of WAN and CQ components in the second release.
- Rotated release manager responsibilities for second and soon to be third release.
- Discussions about donation of documentation.
- Geode portfile was added to MacPorts.
- joetsimple source replaced with binary dependency.
- Active discussions on topics like integrated security for Geode and wording of the website.

Community events:

- June 2016
 - Midwest In Memory Data Grid Meet Up on using Apache Geode for real time transaction. (
 - <http://www.meetup.com/Chicago-In-Memory-Data-Grid-Meetup/events/230810905/>)
 - Apache Geode meetup in Tokyo
 - (<http://pivotal-japan.connpass.com/event/33297/>)

Date of last release:

2016-04-22

When were the last committers or PMC members elected?

Nabarun Nag (2016-06-14)

Signed-off-by:

[](geode) Konstantin Boudnik
[](geode) Chip Childers
[](geode) Justin Erenkrantz
[](geode) Jan Iversen
[X](geode) Chris Mattmann
[](geode) William A. Rowe Jr.
[X](geode) Roman Shaposhnik

Shepherd/Mentor notes:

Roman Shaposhnik:

An extremely vibrant community on the verge of getting to graduation soon.

Chris Mattmann:

The package name change to o.a - what is the hold-up there? Seems that there was a release (or two) already so just wondering. Seems like great progress.

Gossip

Gossip is a system to form peer-to-peer networks using the gossip protocol.

Gossip has been incubating since 2016-04-28.

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

1. Establish an Apache website for Gossip that makes it easy for new users and contributors to get started.
2. Facilitate discussions and build the protocol specification and implementation
3. Produce a usable 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 have had activity in the community from two people who have engaged in working on smaller tasks. We are discussion electing them as project committers.

How has the project developed since the last report?

We have created and resolved a number of tickets: 7 created, 3 resolved.

We have had some discussion around made some traction around implementing the protocol via transports other than UDP.

We are currently working on the web page which should attract more interest.

Date of last release:

Never

When were the last committers or PMC members elected?

Never

Signed-off-by:

[x](gossip) P. Taylor Goetz
[x](gossip) Josh Elser

[x](gossip) Sean Busbey

Shepherd/Mentor notes:

Josh Elser:

Podling is still finding its legs. Community processes are in the works and I think they're doing well so far.

HAWQ

Apache HAWQ is a Hadoop native SQL query engine that combines the key technological advantages of MPP database with the scalability and convenience of Hadoop. HAWQ reads data from and writes data to HDFS natively. HAWQ delivers industry-leading performance and linear scalability. It provides users the tools to confidently and successfully interact with petabyte range data sets. HAWQ provides users with a complete, standards compliant SQL interface.

HAWQ has been incubating since 2015-09-04.

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

1. Produce our first Apache Release
2. Expand the community, by adding new contributors and focusing on making sure that there's a much more robust level of conversations and discussions happening around roadmaps and feature development on the public dev mailing list
3. Infrastructure migration: create Jenkins projects that build HAWQ binary, source tarballs, and run feature tests including at least installcheck-good tests for each commit (HAWQ-127).

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

Everything seems to be smooth, nothing urgent at this time.

How has the community developed since the last report?

1. we have seen significant increase of communication and collaboration across the whole community from last several weeks. The community becomes more open for discussion around the roadmap, various features that committers are working on, and infrastructure enhancements for the project.
2. One meetup: Zeppelin meets MADlib & HAWQ
(<https://www.youtube.com/watch?v=p8V10IukBes&list=PL62pIycqXx-Qf6EXu5FDxUgXW23BH0tcQ&index=6>)
3. The PPMC was made aware of an erroneous use of Apache HAWQ (incubating) trademark by a vendor doing research with an older, eponymous Pivotal product. The PPMC is working with a vendor and ASF's VP of Branding to resolve the issue.

How has the project developed since the last report?

1. The first release candidate has been proposed and it is under the voting process on the dev mailing list
2. New features added include snappy compression for AO tables, HAWQ register feature for registering data into HAWQ native tables
3. The build system has been simplified a lot by making some components as optional (for example, pl/r, orca et al)
4. The community is consolidating the testing frameworks including installcheck tests and gtest into gtest based framework.
5. Including the incubating disclaimer on Apache HAWQ website based on the podling branding guidelines

Date of last release:

We have not had a release yet.

When were the last committers or PMC members elected?

No new committers/members from initial.

Signed-off-by:

[X](hawq) Alan Gates
[](hawq) Konstantin Boudnik
[](hawq) Justin Erenkrantz
[](hawq) Thejas Nair
[X](hawq) Roman Shaposhnik

Shepherd/Mentor notes:

Roman Shaposhnik:

I'm happy to report that the trademark violation incident was taken care of by the offending party quite to the delight of HAWQ's PMC.

HORN

HORN is a neuron-centric programming APIs and execution framework for large-scale deep learning, built on top of Apache Hama.

HORN has been incubating since 2015-09-04.

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

1. Make an initial release
2. Keep up development velocity.
3. Continue to grow community.

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

N/A

How has the community developed since the last report?

dev@ list size is continuously increasing: 41.

How has the project developed since the last report?

We implemented dropout network and MNIST classifier example based on neuron centric model.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

November

Signed-off-by:

[](horn) Luciano Resende
[](horn) Robin Anil
[x](horn) Edward J. Yoon

Shepherd/Mentor notes:

John Ament:

There has been a large drop in mail volume. No real mailing list communication.

iota

Open source system that enables the orchestration of IoT devices.

iota has been incubating since 2016-01-20.

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

1. Building a community now that we have completed source code upload of one for the main components of the system.
2. Moving towards an initial release before the end of the year
3. Working on IP clearance for the 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?

Good input from Anatole Tresch. We have three new contributors Barbara Gomes and two other individuals all have made a contribution this month.

How has the project developed since the last report?

We have uploaded source code for a major component of iota (the Fey engine). We have also uploaded a few basic components that can be used with Fey. We will be adding more components and we hope that individuals who download Fey can also begin the development of their own Fey components (called Performers). We hope they will share their Performers with the Community. We will be creating requests for the community to develop specific performers. E.g. an ActiveMQ performer.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

At the start of the project.

Signed-off-by:

[](iota) Daniel Gruno
[](iota) Sterling Hughes
[X](iota) Justin Mclean
[](iota) Hadrian Zbarcea

Shepherd/Mentor notes:

Justin Mclean:

Finally progress! Good to see but I would of expected to see some explanation of how the discussion has finally moved to the discussion list in issues.

Drew Farris:

It is great to see some progress! One active mentor. Active discussion on the mailing lists. Early discussions about releases.

Juneau

Apache Juneau is a toolkit for marshalling POJOs to a wide variety of content types using a common framework, and for creating sophisticated self-documenting REST interfaces and microservices using VERY little code.

Juneau has been incubating since 2016-06-24.

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

1. Finish up legal paperwork with IBM legal.
2. Build up community with non-IBM and non-Salesforce contributors.

3. Create website and SCM repositories.

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

There are some questions as to what our options are for bug reporting. It's not clear if Jira is our only 'practical' solution or if GitHub Issues is also an option.

How has the community developed since the last report?

This is the first report.

How has the project developed since the last report?

Still in the process of setting up the website. Git repo is on hold until legal paperwork is complete.

Date of last release:

N/A

When were the last committers or PMC members elected?

Signed-off-by:

[X](juneau) Craig Russell
[](juneau) Jochen Wiedmann
[X](juneau) John D. Ament

MADlib

Big Data Machine Learning in SQL for Data Scientists.

MADlib has been incubating since 2015-09-15.

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

1. Continue to produce regular Apache (incubating) releases.
2. Expand the community, increase dev list activity and add new contributors.
3. Execute and manage the project according to governance model required by the "Apache Way".

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?

1. MADlib related events in Q2 2016:
 - * April 19 - Joint community call MADlib - Greenplum Database.
Topic: MADlib v1.9 new features (Nandish Jayaram, Ivan Novick, Cesar Rojas, Frank McQuillan)
 - * May 5 - MADLib community call. Topic: Detailed review of the MADlib v1.9 release (Xiaocheng Tang, Frank McQuillan)
 - * June 21 - MADLib community call. Topic: Apache Zeppelin meets Apache MADlib (incubating) and Apache HAWQ (incubating) (Moon soo Lee, Rahul Iyer, Frank McQuillan)
 - * June 21 - Data Engineers Guild meetup in Palo Alto. Topic: The Analytics and Science Behind Connected Transportation (Srivatsan Ramanujam, Esther Vasiete, Ralph Rabbat, Frank McQuillan)
2. Material technical conversations on dev mailing lists and in the appropriate JIRAs and pull requests.
3. We are seeing some PostgreSQL experts chipping on SQL coding and making good suggestions in the pull requests.

How has the project developed since the last report?

1. 2nd ASF release MADlib v1.9 released on April 6, 2016. The goal of this 2nd release was general availability of MADlib v1.9 for community use.
2. 3rd ASF release MADlib v1.9.1 anticipated this summer depending on community input. Features include: path functions (phase 2), 1-class support vector machines for novelty detection, prediction metrics, sessionization, pivoting.
3. 2 JIRAs created and 14 resolved in last 30 days.

Date of last release:

MADlib v1.9 on 4/6/16.

When were the last committers or PMC members elected?

Xiaocheng Tang on 1/14/16.

Signed-off-by:

[](madlib) Konstantin Boudnik
[x](madlib) Ted Dunning
[x](madlib) Roman Shaposhnik

Metron

Metron integrates a variety of open source big data technologies in order to offer a centralized tool for security monitoring and analysis. Metron provides capabilities for log aggregation, full packet capture indexing, storage, advanced behavioral analytics and data enrichment, while applying the most current threat-intelligence information to security telemetry within a single platform.

Metron has been incubating since 2015-12-06.

Three most important issues to address in the move towards Graduation.

- Building a diverse community of developers for Metron
- Getting security practitioners to provide feedback on requirements
- Make an Apache release

Any issues that the Incubator PMC or ASF Board might wish/need to be aware of

- We are currently trying to produce our second Apache release of Metron

How has the community developed since the last report

- In the process of adding one additional committer, Michael Miklavcic
- We accepted numerous contributions from non-committers interested in Metron, including additional parsers and sample telemetries
- We host a weekly UI working session and were able to source requirements for a new UI prototype
- We hosted a Metron meetup, which received a lot of interest from the community
- We saw an uptick in interest from the community on the dev boards which corresponds to people trying the product prior to an official release, which is a good sign.
- We are receiving community JIRAs for bug reports and security issues

How has the project developed since the last report

- We closed over 100 Jiras and are in the latest stages of releasing our second Apache release.
- We added additional installation automation and documentation
- We added additional features around parsing and a rules engine through a DSL called Stellar
- We are adding infrastructure called "Model as a Service" to uptake ML

models

Date of last release:

April 2016

When were the last committers or PMC members elected?

David Lyle on 22 March 2016

Signed-off-by:

[x](metron) Billie Rinaldi
[x](metron) Chris Mattmann
[](metron) Owen O'Malley
[x](metron) P. Taylor Goetz
[](metron) Vinod Kumar Vavilapalli

Milagro

Distributed Cryptography; M-Pin protocol for Identity and Trust

Milagro has been incubating since 2015-12-21.

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

1. Bring regular day-to-day activity to Apache resources.
2. Make a first Apache release.
3. Attract community.

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

The project has resources outside Apache, including a prospective main documentation website milagro.io. This is not branded as Apache, but is considered to be owned by Milagro.

How has the community developed since the last report?

Conference presentations from key team members, including at ApacheCon/Vancouver and forthcoming in Tokyo, seek to build awareness among Apache and open source communities.

There appears also to be recent new community interest around the pre-Apache project at Github. This should probably be redirected to Apache, but we need to check.

How has the project developed since the last report?

Several existing repos formerly developed by Miracl at Github have been imported to Apache, and are mirrored to Github. These were packaged to be near release-ready prior to import, and we are looking to review readiness for a release.

The majority of work that has taken place at Apache has been the development of a Milagro web site, and project documentation at milagro.io. Code development appears to have continued at github, apparently due to confusion. This was discovered just today, so we plan to follow up tomorrow.

Date of last release:

N/A

When were the last committers or PMC members elected?

No new members elected during the period. NTT colleagues named in the proposal have joined.

Signed-off-by:

[](milagro) Sterling Hughes
[](milagro) Jan Willem Janssen
[x](milagro) Nick Kew

Shepherd/Mentor notes:

Nick Kew:

I am currently investigating whether recent community activity in the pre-Apache github resources relates to code that has been contributed to Milagro. If I find it is, I shall attempt to encourage that activity to move to Apache.

Mynewt

Mynewt is a real-time operating system for constrained embedded systems like wearables, lightbulbs, locks and doorbells. It works on a variety of 32-bit MCUs (microcontrollers), including ARM Cortex-M and MIPS architectures.

Mynewt has been incubating since 2015-10-20.

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

1. Make further point releases capable of producing a downloadable RTOS image with support for multiple peripherals and network connectivity with the goals of the first major (1.0) release in the last quarter of 2016 and continued demonstration of thorough understanding, repeatability, and maturity of process.
2. Continue to develop and execute policies that enable project contributors to achieve self-governance
3. Expand community - attract new project contributors and users, grow committer base

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?

1. Active mailing lists with increasing numbers of subscribers. The subscriber count more than doubled since the April report.
2. Continued outside interest by 3rd parties as evidenced by use in multiple new devices, bug reporting and fixing, feature additions. Outreach via conferences, exhibits, one on one meetings, beta testers, GSoC participation.
3. Vigorous discussions and implementation decisions via proposals and voting on @dev mailing list.

How has the project developed since the last report?

1. Releases: Four releases so far, two since the last report. The latest was version 0.9 on June 9th.
 - Regular process established and followed to check for license issues esp. with third party software
 - Tagging process and release candidate naming process tested and used
 - Voting process and timelines established and used
2. Documentation support: In addition to committers, several non-committers added documentation through pull requests on github mirror according to defined process.
 - Support for versioning of documents added - to be used starting next release.
3. Effort towards self governance:
 - Four releases completed successfully.
 - Voting successfully completed to grant committer status to two candidates after meeting threshold of patch submissions. Several new committer candidates are in the pipeline.

Date of last release:

2016-06-06

When were the last committers or PMC members elected?

At the start

Signed-off-by:

[](mynewt) Sterling Hughes
[X](mynewt) Jim Jagielski
[X](mynewt) Justin Mclean
[](mynewt) Greg Stein
[X](mynewt) P. Taylor Goetz

Shepherd/Mentor notes:

Drew Farris:

Three mentors active on mailing lists. High activity on mailing lists.

ODF Toolkit

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

ODF Toolkit has been incubating since 2011-08-01.

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

1. Attract more developers
2. Have frequent releases
3. Decide on possible ways to graduate

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

Not enough developer time for a new release in the summer.

How has the community developed since the last report?

Patches were being provided.

How has the project developed since the last report?

Several patches had been reviewed and applied, a new release is ongoing to update the JDK version and embrace new libraries.

Date of last release:

2014-06-02

When were the last committers or PMC members elected?

2012-10-29

Signed-off-by:

[](odftoolkit) Sam Ruby
[X](odftoolkit) Nick Burch
[](odftoolkit) Yegor Kozlov

Shepherd/Mentor notes:

Nick Burch:

Project is in a bit of a limbo state at the moment. Really needs another release soon, then depending on interest that gets, either find the project a home (eg POI or Commons), or give up...

Pirk

Pirk is a framework for scalable Private Information Retrieval (PIR).

Pirk has been incubating since 2016-06-17.

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

1. Initial Code Preparation: Ensure all dependencies are compliant with Apache License version 2.0 and that all code and documentation artifacts have the correct Apache licensing markings and notice.
2. Release Process: Establish a formal release process and schedule, allowing for dependable release cycles in a manner consistent with the Apache development process.
3. Community Building: Grow the community to establish diversity of background and expertise.

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?

This is our first report; we look forward to developing the community once our code, website, etc. is pushed.

How has the project developed since the last report?

This is our first report; we have been working to set up all of our infrastructure and anticipate pushing our code and website very soon.

Date of last release:

No releases thus far.

When were the last committers or PMC members elected?

Upon acceptance into the incubator, 2016.06.17.

Signed-off-by:

[x](pirk) Billie Rinaldi
[x](pirk) Joe Witt
[x](pirk) Josh Elser
[X](pirk) Suneel Marthi
[X](pirk) Tim Ellison

Pony Mail

Pony Mail is a mail-archiving, archive viewing, and interaction service, that can be integrated with many email platforms.

Pony Mail has been incubating since 2016-05-27.

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

1. Grow the community and diversity
2. Work towards the first ASF release
3. Get people interested in providing patches (work on the code)

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?

Many people have stepped up with issues/bug reports and other suggestions, which suggests an interest in the project. We are trying to get people interested in also providing patches for said issues, so as to grow the developer community.

How has the project developed since the last report?

A few bug fixes, repositories have been set up, mailing lists created, web site deployed. We are also seeing moderate activity on the IRC channel.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

No new committers have been elected since incubation started.

Signed-off-by:

[](ponymail) Andrew Bayer
[X](ponymail) John D. Ament

PredictionIO

PredictionIO is an open source Machine Learning Server built on top of state- of-the-art open source stack, that enables developers to manage and deploy production-ready predictive services for various kinds of machine learning tasks.

PredictionIO has been incubating since 2016-05-26.

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

1. Unblock open issues and code merges and create the first Apache branded release.
2. Migrating the existing community to Apache infrastructure, and growing the contributor base beyond initial committers.
3. Implement template gallery in a more scalable, community-driven portal, since PredictionIO relies on templates to function.

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

Apache JIRA and GitHub integration does not seem to be working 100%.

How has the community developed since the last report?

We have been working to set up infrastructure on ASF, and are looking to migrate the community from old to new infrastructure to allow further growth.

How has the project developed since the last report?

An initial set of JIRA tickets based on outstanding issues before incubation have been created. Initial committers / contributors are ready to work on them to march toward the first Apache branded release.

Date of last release:

No releases yet

When were the last committers or PMC members elected?

2016-06-06 (initial committers for incubation)

Signed-off-by:

[X](predictionio) Andrew Purtell
[](predictionio) James Taylor

[](predictionio) Lars Hofhansl
[](predictionio) Luciano Resende
[](predictionio) Xiangrui Meng
[X](predictionio) Suneel Marthi

Rya

Rya (pronounced "ree-uh" /re?/) is a cloud-based RDF triple store that supports SPARQL queries. Rya is a scalable RDF data management system built on top of Accumulo. Rya uses novel storage methods, indexing schemes, and query processing techniques that scale to billions of triples across multiple nodes. Rya provides fast and easy access to the data through SPARQL, a conventional query mechanism for RDF data.

Rya has been incubating since 2015-09-18.

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

1. Become familiar with the release process and have a first release as part of the Apache Foundation
2. Expand the documentation to be more formalized and more representative of the current codebase.
3. Add new committers to the project.

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 a Rya "office hour" teleconference every other week. Users and developers can ask questions, receive answers, discuss ideas for future developments. The minutes are sent to the dev@ list. Attendance at these meetings has grown beyond existing committers of the project to include new contributors. New contributors have led the discussion for new features that they have helped develop. We have also begun posting minutes and reference slides from the meetings on confluence to provide reference documentation to new contributors.
- * We have more PRs from non-committers which are integrated into the repository
- * Developed a more extensive Rya website and confluence page for documentation for new contributors

How has the project developed since the last report?

- * Integrated Rya with Apache build process -- Rya is currently being built by Jenkins on Apache servers
- * Resolved a handful of issues that users have encountered.
- * Upgrades for versions for a number of dependencies
- * Committed features: advanced OWL forward chaining reasoning tool for Accumulo based Rya, extensions to pre-computed join capabilities for more complex SPARQL queries, additions of timestamps for Mongo backed Rya.

Date of last release:

Not applicable

When were the last committers or PMC members elected?

Not applicable. The PPMC is currently discussing inviting a couple of active contributors to become committers.

Signed-off-by:

[x](rya) Josh Elser
[](rya) Edward J. Yoon
[x](rya) Sean Busbey
[x](rya) Venkatesh Seetharam

Shepherd/Mentor notes:

Josh Elser:

Podling activity is steady and committers are always responsive.
Discussions around inviting some contributors to join the project are on-going which is good. Prodding to make a release continues from me.

Tephra

Tephra is a system for providing globally consistent transactions on top of Apache HBase and other storage engines.

Tephra has been incubating since 2016-03-07.

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

1. Regular releases
2. Improve community engagement
3. 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?

1. Started to see mailing list and JIRA activity from non-PMC members.

How has the project developed since the last report?

1. tephra.incubator.apache.org website setup is complete.
2. Apache Tephra JIRA setup is complete.

Date of last release:

2016-05-31

When were the last committers or PMC members elected?

- None since coming to incubation

Signed-off-by:

[X](tephra) Alan Gates
[](tephra) Andrew Purtell
[X](tephra) Henry Saputra
[](tephra) James Taylor
[](tephra) Lars Hofhansl