

October2016

/\ FINAL (!)

Incubator PMC report for October 2016

Timeline

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

Shepherd Assignments

Drew Farris	SensSoft		
Drew Farris	Traffic Control		
John Ament	Distribute dLog		
John Ament	MADlib		
Justin Mclean	Geode		
P. Taylor Goetz	AriaTosca		
P. Taylor Goetz	HORN		
P. Taylor Goetz	Juneau		
Timothy Chen	FreeMarker		
Timothy Chen	HAWQ		
Timothy Chen	Mynewt		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="13b33d1c-7e55-4b47-86da-266d59eb5a34"><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="a7643285-efe9-484d-88fd-e76fd0bd462e"><ac:plain-text-body><![CDATA[[none]	Annotator]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="87a332f1-e4c8-4a7a-8503-871e3eb0478c"><ac:plain-text-body><![CDATA[[none]	BatchEE]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="32b04989-6a73-4dd6-a330-eb77086530f9"><ac:plain-text-body><![CDATA[[none]	DataFu]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="559b1f00-6dc0-4196-92e1-a3d4786d32f1"><ac:plain-text-body><![CDATA[[none]	Gossip]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ff76f890-dd2d-4980-873b-cfd5ac98481d"><ac:plain-text-body><![CDATA[[none]	MRQL]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6cb77f00-fc72-41e2-98fd-a5e2e95e809b"><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="72e15d4d-24ed-440b-b386-d7ddf890caa4"><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="8785c8ce-2597-428e-af13-2a4561ff096f"><ac:plain-text-body><![CDATA[[none]	Streams]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f57c8359-359d-4199-acc8-4b85f0b8dcde"><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="4330ed74-b0ec-4e55-be12-1292bbf5c44c"><ac:plain-text-body><![CDATA[[none]	Quickstep]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="222d4ec3-e393-4e0e-bf78-e962a8f08c89"><ac:plain-text-body><![CDATA[[none]	Rya]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="af26c111-5330-44ff-83ac-04afce1a8810"><ac:plain-text-body><![CDATA[[none]	Tamaya]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5933fad5-1b7e-4a36-8bbd-9d0acd53b601"><ac:plain-text-body><![CDATA[[none]	Toree]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for October 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 currently 63 podlings undergoing incubation - new record from 61 last month!

* Community

New IPMC members:

- Felix Meschberger

People who left the IPMC:

N/A

* New Podlings

- Hivemall
- Spot
- NetBeans

Please note: none of the new podlings submitted reports. They will begin reporting next month.

* Graduations

The board has motions for the following:

None

* Releases

The following releases entered distribution during the month of September:

- Apache Singa 1.0.0-incubating 08 Sep 2016
- Apache Tephra 0.9.0-incubating 19 Sep 2016
- Apache Madlib 1.9.1-incubating 19 Sep 2016
- Apache Samoa 0.4.0-incubating 22 Sep 2016
- Apache Batchee 0.4-incubating 29 Sep 2016
- Apache Streams 0.3-incubating 29 Sep 2016

* IP Clearance

Apache Storm JMS Integration from P. Taylor Goetz
Apache Storm SQE from JWPlayer

A vote was started to import Aether from Eclipse Foundation as Maven Artifact Resolver, however it has not been closed as of the time of writing.

* Legal / Trademarks

Sally has begun work on helping improve the incubator brand. She has taken over the Twitter account and has proposed putting together a selection for a new logo.

* Infrastructure

Kudos to the infrastructure team for their hard work on the Netbeans proposal, getting all of the needed infrastructure assessment prior to the incubator accepting this podling.

Incubator usage of the new reporeq tool has been mixed. It would be great to see a blog post or something about how to submit git repository requests.

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- Annotator
- AriaTosca
- Juneau

* Not yet ready to graduate

No release:

- Airflow
- FreeMarker
- Quickstep

Community growth:

- Geode
- Gossip
- HAWQ
- MRQL
- Mynewt
- Rya

* Ready to graduate

- BatchEE

* Did not report, expected next month

- DataFu

Table of Contents

- Airflow
- Annotator
- AriaTosca
- BatchEE
- DistributedLog
- FreeMarker
- Geode
- Gossip
- HAWQ
- HORN
- Juneau
- MADlib
- Metron
- Milagro
- MRQL
- Mynewt
- ODF Toolkit
- Quickstep
- Rya
- SensSoft
- Streams
- Tamaya
- Toree
- Traffic Control

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.
- 3.

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

We discovered a private Gitter channel for committers that was created during the days when the Git project was owned by Airbnb. When the Github project's ownership moved to the ASF, so did that of all the Gitter channels. We killed the channel - all committers left the channel and can no longer even see/discover it.

How has the community developed since the last report?

1. We elected 2 new PPMC Members/Committers (Li Xuanji a.k.a. Zodiac & Sumit Maheshwari a.k.a. msumit) 2. Since our last podling report 3 months ago (i.e. between Jul 7 and Oct 4, inclusive), we grew our contributors from 160 to 191 3. Since our last podling report 3 months ago (i.e. between Jul 7 and Oct 4, inclusive), we merged >142 pull requests 4. Two meet-ups (One for HUG Warsaw, the other an Airflow meet-up @Stripe HQ in SF) and several public talks (e.g. DataDaySeattle, PyData SF) were held by the community 5. Since being accepted into the incubator, the number of companies using Apache Airflow has risen from 30 to 58

How has the project developed since the last report?

1. Podling name search approved : <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-110> 2. Very active discussion list, public Gitter channels and JIRA (more than 500 issues)

Date of last release:

None. First ASF release currently being discussed.

When were the last committers or PMC members elected?

As mentioned on <https://cwiki.apache.org/confluence/display/AIRFLOW/Announcements>, 1. On Aug 16, Sumit Maheshwari (a.k.a. msumit) was elected to the PPMC and Committer group 2. On Aug 19, Li Xuanji (a.k.a. zodiac) was elected to the PPMC and Committer group

Signed-off-by:

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

Annotator

Annotator provides annotation enabling code for browsers, servers, and humans.

Annotator has been incubating since 2016-08-30.

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

1. Setup project infrastructure at the ASF
2. Migrate existing community to the ASF
3. Engage committers listed in our Proposal to plan next steps

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?

The community has established support for incubation and initial committers and mentors have been added to project email lists.

How has the project developed since the last report?

Project code identified in our Proposal remains under active development with some components having achieved new version milestones in the last month.

Committers are working to establish a web presence and code collaboration on ASF infrastructure.

Date of last release:

n/a

When were the last committers or PMC members elected?

Signed-off-by:

[x](annotator) Nick Kew
[](annotator) Brian McCallister
[](annotator) Daniel Gruno
[X](annotator) Jim Jagielski

AriaTosca

ARIA TOSCA project offers an easily consumable Software Development Kit(SDK) and a Command Line Interface(CLI) to implement TOSCA(Topology and Orchestration Specification of Cloud Applications) based solutions.

AriaTosca has been incubating since 2016-08-27.

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

- 1.Establish a formal release process and schedule, allowing for dependable release cycles in a manner consistent with the Apache development process and formulate a plan for fully functional release 1 of ARIA and populate tasks to ASF Jira
- 2.Migrate existing users to apache mailing lists, redirect the present external project links to Apache.
- 3.Grow the community around the project.

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

Project is still being established and presently onboarding the initial committers.

How has the community developed since the last report?

Received Software Grant Agreement (SGA) from Gigaspaces.

The mailing lists, Jira, Github Repo, Incubator podling, Confluence Wiki have been setup for the project.

The initial set of committers now have their apacheIds.

How has the project developed since the last report?

The project is presently being setup in Incubator. We expect the first code push to happen in the month of October.

Date of last release:

None

When were the last committers or PMC members elected?

Project is running with the initial set of committers.

Signed-off-by:

[X](ariatosca) Suneel Marthi
[X](ariatosca) John D. Ament
[](ariatosca) Jakob Homan

BatchEE

BatchEE projects aims to provide a JBatch implementation (aka JSR352) and a set of useful extensions for this specification.

BatchEE has been incubating since 2013-10-03.

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

1. fix last few cleanups on the licence/notice files
2. enhance the website and download page which is pretty poor at the moment
- 3.

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?

Some interactions on JIRA and twitter, Stian participated to the release vote!

How has the project developed since the last report?

Got some enhancements, fixes and a new release.

Date of last release:

2016-09-29

When were the last committers or PMC members elected?

2015-12-02

Signed-off-by:

[](batchee) Jean-Baptiste Onofré
[](batchee) Olivier Lamy
[X](batchee) Mark Struberg

DataFu

johndament:

Discussions on this podling seem to have stopped completely. There was a graduation discussion back in August, which seems to have dropped completely after some release content issues were identified.

DistributedLog

DistributedLog is a high-performance replicated log service. It offers durability, replication and strong consistency, which provides a fundamental building block for building reliable distributed systems.

DistributedLog has been incubating since 2016-06-24.

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

1. Continue to grow the community, especially people from outside Twitter
2. Improve documentation, especially admin and user guides
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?

1. Increase in contributions from community, more discussion of subproject ideas
2. Lots of engagement on documentation requests and etc
3. Improved responsiveness of committers to community members

How has the project developed since the last report?

1. Documentation has improved significantly with addition of clean, easy to understand and well designed website with landing page
2. First release and repackaging of the project under apache namespace are being discussed

Date of last release:

NA

When were the last committers or PMC members elected?

NA

Signed-off-by:

[X](distributedlog) Flavio Junqueira
[X](distributedlog) Chris Nauroth
[](distributedlog) Henry Saputra

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. Community growth; we need more developers.
2. Fill blanks left on the status page.

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?

There was a discussion about a graduation attempt, and the overall result (though opinions can slightly differ) was that we better see if we can attract more developers first. For that, we intend to start a branch that gets rid of the several backward compatibility constraints that makes contribution and addressing some interesting templating problems rather difficult or impossible, and which makes the code base more attractive in general. (Note that this is 15 year old code base, and the need for this is not a new idea.) This branch would be also the opportunity to switch to org.apache Java packages and org.apache Maven coordinates.

We have positively completed the verifications about possible name conflicts (PODLINGNAMESEARCH-107).

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

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:

1. Expanding the community to include contributors and committers outside of Pivotal.
2. Establish a steady release cadence.
3. Discuss graduation on the geode dev list.

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?

- Karen Miller was added as a new committer and PPMC member
- 3 geode clubhouses on
 - Domain Driven Design & Reactive Programming w/ Apache Geode
 - What's New with Spring Data Geode
 - Events & Continuous Query
- 3 presentations involving geode at SpringOne:
 - Where Does Apache Geode Fit in CQRS Architectures?
 - Spring Data and In-Memory Data Management in Action
 - Design Tradeoffs in Distributed Systems - How Southwest Airlines Uses Geode
- A geode user group meeting in Tokyo

How has the project developed since the last report?

- We renamed packages from com.gemstone to org.apache to move us closer to graduation
- The geode documentation was donated to the project
- We had our third public release 1.0.0-incubating.M3
- We have cut a release branch and getting close to our final 1.0.0 release, using the apache package names

Date of last release:

2016-08-23

When were the last committers or PMC members elected?

Karen Miller 2016-07-18

Signed-off-by:

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

Shepherd/Mentor notes:

(rvs) Honestly, I think the project is one release away from starting a formal graduation discussion.

Gossip

Gossip is an implementation of the Gossip Protocol.

Gossip has been incubating since 2016-04-28.

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

1. Produce a release
2. Update the site, especially how-to-contribute documentation
3. Develop specifications for protocol

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?

We have accepted features/patches from two contributors. We have opened up a discussion to add one of them as a project committer.

How has the project developed since the last report?

The ability to gossip user defined payloads (<https://issues.apache.org/jira/browse/GOSSIP-21>) has been added to the project.

This is a key feature that downstream tools needed to use the project.

Contributors have moved the project from Junit4 to Junit5 (<https://issues.apache.org/jira/browse/GOSSIP-11>), as well as redesign/re-factoring of system internals (<https://issues.apache.org/jira/browse/GOSSIP-15>)

Date of last release:

(None yet)

When were the last committers or PMC members elected?

No

Signed-off-by:

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

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 in past 3 months. The community becomes more open for discussion around the roadmap, various features that committers are working on, infrastructure enhancements for the project.

2. One talk at trust cloud computing summit

* Apache HAWQ: The leading SQL-on-Hadoop Query Engine.
(<http://www.cnii.com.cn/technology/img/4598.files/yicheng.html>)

3. Two meetups: Data Science at Scale with HAWQ and MADlib and Hadoop

* <https://www.meetup.com/Data-Engineers-Guild/events/233974633/>
* <https://www.meetup.com/futureofdata-siliconvalley/events/232976650/>

How has the project developed since the last report?

1. The release candidate (Apache HAWQ 2.0.0.0-incubating RC4) has been proposed and the voting process on the dev mailing list completed. The main target of the first release is to clear all IP related issues for HAWQ and this is a source code tarball only release. Full list of JIRAs fixed/related to the release: link
<<https://cwiki.apache.org/confluence/display/HAWQ/HAWQ+Release+2.0.0.0-incubating>>

2. New features added include

* snappy compression for AO tables
* HAWQ register feature for registering data into HAWQ native tables
* PXF - JSON support (<https://issues.apache.org/jira/browse/HAWQ-178>)

Date of last release:

We have not had a release yet.

When were the last committers or PMC members elected?

Add one committer in September: Kavinder Dhaliwal

Signed-off-by:

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

Shepherd/Mentor notes:

(rvs) Now that the first release of HAWQ has officially been published, I believe the community has unblocked itself and will be running in high gear. Looking forward to a cadence of releases and community expansion.

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?

* We received few 'pull request' contributions from government research institute.

How has the project developed since the last report?

* Improved CLI commands
* Developing recurrent neural network and LSTM
* Writing a paper about HORN project (<http://arxiv.org/abs/1608.00781>)

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

Signed-off-by:

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

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. Build up community with non-IBM and non-Salesforce contributors.
2. Release a version of the product to spark external interest in the project.
3. Finish building community architecture and guidelines.

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?

No new members in September. The hope is that a release of the product will spark interest.

How has the project developed since the last report?

Versioning and release guidelines have been created in Confluence. Currently in the process of creating our first release.

Date of last release:

None.

When were the last committers or PMC members elected?

Signed-off-by:

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

Shepherd/Mentor notes:

johndament:

The project has room for growth. Its currently a 1 man project, will be good to see a release go out to help build a larger community.

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. Need guidance from Incubator PMC on how to resolve the BSD licensing switch over to Apache License. What should be the content of the license headers for files that were previously BSD licensed and then granted to ASF? Related legal-discuss threads:
http://mail-archives.apache.org/mod_mbox/www-legal-discuss/201609.mbox/%3CALGG8z03zHhbFegXoi4fH+vXtF+9m7x6hak9RjKQjapuzi67gQ@mail.gmail.com%3E
http://mail-archives.apache.org/mod_mbox/www-legal-discuss/201603.mbox/%3CD1AF43C-370B-4E58-B0EF-2E29D242F50B%40jaguNET.com%3E
2. Continue to produce regular Apache (incubating) releases.
3. Continue to 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?

Yes-please see #1 above and provide guidance.

How has the community developed since the last report?

1. Two new committers added to the project:
 - * Orhan Kislal (9/7/16)
 - * Nandish Jayaram (9/7/16)
2. MADlib related events in Q3 2016:
 - * Jul 27 - MADLib community call. Topic: Open discussion on Apache MADlib project (hosted by Greg Chase, Frank McQuillan)
 - * Aug 19 - Presentation to Hortonworks. Topic: Apache MADlib, Apache HAWQ (incubating) and Apache Zeppelin (Rahul Iyer, Frank McQuillan)
 - * Sep 13 - MADLib community call. Topic: Deep dive on MADlib 1.9.1 release (hosted by Greg Chase, presentation by Frank McQuillan)
 - * Sep 21 - Meetup at Hortonworks San Francisco. Topic: Future of data - Apache MADlib and Apache HAWQ (Tushar Pednekar)
 - * Sep 22 - Meetup at Hortonworks Santa Clara. Topic: Future of data - Apache MADlib and Apache HAWQ (Tushar Pednekar)
3. Material technical conversations on dev mailing lists and in the appropriate JIRAs and pull requests.

How has the project developed since the last report?

1. 3rd ASF release MADlib v1.9.1 released on Sep 19, 2016. Features include: path functions (phase 2), 1-class support vector machines for novelty detection, prediction metrics, sessionization, pivoting.
2. Community has started active development on the v1.10 release.
3. 13 JIRAs created and 5 resolved in last 30 days.

Date of last release:

MADlib v1.9.1 on 9/19/16.

When were the last committers or PMC members elected:

Orhan Kislal on 9/7/16 and Nandish Jayaram on 9/7/16.

Signed-off-by:

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

Shepherd/Mentor notes:

tdunning:

This project seems to be ticking along pretty reasonably. The only worry I have about it is that it seems to be strongly centered around a few (or even just one) very strong contributors. That is a worry relative to longevity and community building. Overall, I don't think that the project is getting much marginal value from incubation.

johndament:

Its unclear what guidance from the IPMC is required if the podling is already reaching out to legal, which would be the main thing I can think of to recommend to them right now.

rvs:

@johndament: I think we need to formalize whatever decision by legal. I'll create a formal LEGAL JIRA soon.

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 12-08-2015

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 in the process of producing our third Apache release

* How has the community developed since the last report

- We received strong community contributions in both code, documentation, and testing from two additional community members and are hoping to make them committers shortly

- David Lyle and Nick Allen were promoted to PPMC
- Tweaked and re-voted on the community bylaws
- Introduced a by-weekly community demo and meeting

* How has the project developed since the last report

- We closed on 41 more Jiras and made our code base more compliant and introduced several important features. We also had lively discussion on the dev boards about how these features should be implemented.

Signed-off-by:

[x](metron) Billie Rinaldi
[](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. Make a first release
2. Continue to migrate to full Open Development.
3. Building the MILAGRO community - engaging developers and cryptographers beyond the core companies, raising awareness and helping to secure future of internet.

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?

We have elected two new committers. User questions are dealt with well on the mailinglist, although the volume remains low.

How has the project developed since the last report?

Commit activity has finally migrated from Github. The project is very keen to adopt full two-way mirroring with github as soon as possible. Mailinglist activity remains low.

Main development focus over last couple of months was on crypto library (MILAGRO crypto lib), which was re-factored and improved, in order to improve the code quality, readability, and test coverage.

Date of last release:

n/a

When were the last committers or PMC members elected?

Alessandro Budroni and Nicola Asuni were elected in July 2016.

Signed-off-by:

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

MRQL

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

MRQL has been incubating since 2013-03-13.

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

1. Increase the number of active committers
2. Increase adoption, expand user community, and increase user list activity

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?

There were no new developers or new committers since our last report.

How has the project developed since the last report?

There were 8 JIRA issues reported during this period, from which 7 have already been resolved. The main work during this period was on algebra-based and source-level debugging of MRQL queries. Basically, results from an MRQL query are annotated with lineage information, which links each result value with the input data that were used to calculate the value. Given a result value, our debugger not only can print the input values that contributed to the result value, but can also display the detailed workflow that was used to calculate this value (how-to provenance). In contrast other data-centric debuggers, such as Titian for Spark, our debugger supports on-line browsing and searching using GUIs and can work on all MRQL supported platforms (map-reduce, Spark, Flink, and Hama). It also allows to insert trace points in the query source to debug the query at the source level. We hope that this debugger will be a valuable tool for developing and understanding MRQL queries.

Date of last release:

2016-03-02

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

[X](mrql) Alan Cabrera
[X](mrql) Edward J. Yoon
[](mrql) Mohammad Nour El-Din

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 with diverse backgrounds, 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. 15 new subscribers on dev@ mailing list since last report. 2. Increased participation by 3rd parties through pull requests for new BSP support, features, and test cases. Outreach continues via conferences, exhibits, one on one meetings, beta testers, GSoC participation. 3. Vigorous discussions, feature proposals, and implementation suggestions on @dev mailing list.

How has the project developed since the last report?

1. Releases: No new releases since June 6th but work on several branches are being merged for first 1.0 beta release in October. The project has now closed 100+ pull requests. 2. Planning: Issues including bug reports, features, and wish list captured and tracked in ASF JIRA by members of the community. Work continues against 125 tickets opened for 1.0 releases (beta and GA). 3. Effort towards self governance: Voting successfully completed to grant committer status to one new candidate since last report.

Date of last release:

2016-06-06

When were the last committers or PMC members elected?

2016-10-04

Signed-off-by:

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

----- 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?

None.

How has the community developed since the last report?

Patches had been provided from LibreOffice developers for the ODF validator.

How has the project developed since the last report?

Svante Schubert is working on the major collaboration feature, likely to attract more developers. He aims to get the basic updates in before the next release based on JDK6. The release is aimed to be before the ApacheCon EU and one release afterwards on JDK 8 including the feature. He will attend the ApacheCon EU to get in contact with people from the POI project in order to evaluate possible ways to graduate.

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 (nick) - hopefully the discussions at ApacheCon will let us formulate a plan for where the project should go, to bring back to the list, and hopefully get ODF Toolkit out of the incubator!

Quickstep

Quickstep is a high-performance database engine.

Quickstep has been incubating since 2016-03-29.

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

1. Create and ASF release (current target is December 20, 2016).
2. Continue to build the developer community.
3. Acquire early adopters.

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?

1. We have two new community members (Tarun Bansal and Tianrun Li). 2. Working with a large company to fund three students for a year and allow them to contribute to Quickstep. Funding has been approved and in the process of finalizing the ICLA approvals.

How has the project developed since the last report?

1. Closed over 10 pull requests. 2. Started running end-to-end benchmarks/workloads to build towards a viable product. 3. Cleaned up the repository to have licensing and copyrights in line with ASF projects.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[x](quickstep) Konstantin Boudnik
[x](quickstep) Julian Hyde
[x](quickstep) Roman Shaposhnik

Shepherd/Mentor notes:

Julian Hyde (jhyde):

I am concerned that the project is evolving but is not doing it in the Apache way. For example, the "new community members" were welcomed on the mailing list but there was not a vote or discussion about making them committers. Also, there is still no release and the project does not seem to be making it a priority, and substantive discussions are not happening on the list (other than discussions about pull requests). Of course, it is great to see commit activity and a growing community.

rvs: I'm very much in agreement with Julian that the project seems to be somewhat in trouble with adopting itself to the "Apache Way". Honestly, it feels like an academic project hosted by ASF at this point. I believe at this point mentors need to show some "tough love" and push for having a real plan of how community thinks about addressing all the issues we've observed so far. Julian and I are working on making it a priority for the PPMC.

Rya

Rya (pronounced "ree-uh" /r/) 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. 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?

- * A new committer was invited and accepted to join the committers for Apache Rya (incubating).
- * Talk on Rya at the DC Graph Database meetup on August 30, 2016
- * Online presentation on Rya to a group of researchers and developers from Cray on September 29, 2016
- * Talk on Rya accepted at Accumulo Summit, October 11, 2016
- * The Rya "office hour" teleconference continues every other week, with few exceptions when people are too busy. 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 continue posting minutes and reference slides from the meetings on confluence to provide reference documentation to new contributors.
- * Multiple PRs from several non-committers which are integrated into the repository continue to be received

How has the project developed since the last report?

- * Worked on the first release. The first release candidate was proposed, but the vote did not pass on the dev list. Working on the next release candidate.
- * Resolved a handful of issues that users have encountered.
- * Committed features: added a benchmark tool for the pre-computed joins optimizer; added support for property chain inference; integrated Apache Fluo (incubating) with Rya so triples inserted in Rya are inserted into Apache Fluo (incubating) too in order to incrementally update any pre-computed joins registered with Fluo; added OPTIONAL support for pre-computed joins

Date of last release:

Not applicable. Initial vote on first release candidate did not pass on dev list. Working on the next release candidate.

When were the last committers or PMC members elected?

New committer Caleb Meier elected on August 31, 2016

Signed-off-by:

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

Shepherd/Mentor notes:

JE: Podling had a first release candidate staged, but it failed to numerous incompatibly-licensed transitive dependencies. Progress on release is stalled, but not forgotten.

SensSoft

SensSoft is a software tool usability testing platform

SensSoft has been incubating since 2016-07-13.

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

1. Moving towards the first Incubating release of the source code and other release artifacts.
2. Grow the Apache SensSoft (Incubating) community.

3. Complete the issues highlighted at the SensSoft Roadmap
<https://cwiki.apache.org/confluence/display/SENSOFT/Roadmap>

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?

The community has not grown over and above the list of initial committers. Draper Labs are doing a bunch of social media outreach and are also demo'ing SensSoft at a number of Human/Computer + UI/UX Events so we hope this will change.

How has the project developed since the last report?

The project has seen loads of activity on Jira and the mailing list. All code has been licensed under ALv2.0 with NOTICE's, DISCLAIMER's, etc now being committed across the SensSoft software suite. Jira is very busy. We are utilizing builds.apache.org for all projects builds... this is part of the aim to move towards CI for the SensSoft stack.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[](senssoft) Paul Ramirez
[X](senssoft) Lewis John McGibbney
[](senssoft) Chris Mattmann

Shepherd/Mentor notes:

Drew Farris (shepherd): One mentor active within the past month. Active JIRA/Developers Mailing Lis

Streams

Apache Streams (incubating) unifies a diverse world of digital profiles and online activities into common formats and vocabularies, and makes these datasets accessible across a variety of databases, devices, and platforms for streaming, browsing, search, sharing, and analytics use-cases.

Streams has been incubating since 2012-11-20.

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

1. Community growth and PMC maturity.
2. Demonstrate a consistent release schedule.
3. Participation of project community within related standards-bodies and Apache projects.

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

A vote to retire the project was initiated but cancelled after several individuals on the IPMC list expressed new interest in helping the project continue to grow and eventually graduate. The project is now on a monthly reporting schedule to monitor effective progress in growing the community and active participation.

How has the community developed since the last report?

dev@streams.incubator.apache.org
46 emails sent by 15 people, divided into 9 topics.

<http://streams.incubator.apache.org>
196 Sessions, 170 Users, 512 Pageviews

How has the project developed since the last report?

<https://github.com/apache/incubator-streams-master>
Excluding merges, 1 author has pushed 9 commits to master and 9 commits to all branches.
On master, 8 files have changed and there have been 114 additions and 193 deletions.

<https://github.com/apache/incubator-streams>
Excluding merges, 1 author has pushed 6 commits to master and 6 commits to all branches.
On master, 61 files have changed and there have been 85 additions and 66 deletions.

<https://github.com/apache/incubator-streams-examples>
Excluding merges, 1 author has pushed 12 commits to master and 15 commits to all branches.
On master, 67 files have changed and there have been 549 additions and 761 deletions.

92 Issues closed with release of 0.3-incubating
7 new Issues opened

Date of last release:

2016-10-03 : 03-incubating release

When were the last committers or PMC members elected?

2016-09-28: Suneel Marthi elected as mentor / PMC member

Signed-off-by:

[X](streams) Ate Douma
[](streams) Matt Franklin
[X](streams) Suneel Marthi

Shepherd/Mentor notes:

Ate Douma: I see promising improvements with open discussions about next steps and future plans, both technically and to broaden the community involvement. It's still wait and see however if and how the community does pick this up.

Tamaya

Tamaya is a highly flexible configuration solution based on a modular, extensible and injectable key/value based design, which should provide a minimal but extensible modern and functional API leveraging SE, ME and EE environments.

Tamaya has been incubating since 2014-11-14.

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

1. Extend Community
2. Improve Documentation
3. Stabilize existing code base

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

See following sections.

How has the community developed since the last report?

A number of existing committers have decided to leave the project, this includes the loss of two mentors and two other committers.

We have increasing feedback of use cases to be covered and also a few bug reports, so Tamaya is used by developers and companies. Additionally we are doing regular bi-weekly hangouts and have defined a well defined roadmap.

How has the project developed since the last report?

We have decided to improve the project setup, so we can release easier and faster in the future: separating the site from the project release, separating experimental modules and stable modules into separate repositories. This work is currently ongoing, but we make good progress and we plan to get out a release of the core parts around end of October.

Date of last release:

2016-04-06

When were the last committers or PMC members elected?

Phil Ottlinger at 24th April 2016.

Signed-off-by:

[X](tamaya) John D. Ament
[](tamaya) David Blevins

Shepherd/Mentor notes:

johndament:

The project has room to improve. With the loss of mentors and existing committers, need help mentoring this project.

Toree

Toree provides applications with a mechanism to interactively and remotely access Apache Spark. It enables interactive workloads between applications and a Spark cluster. As a Jupyter Notebook extension, it provides the user with a preconfigured environment for interacting with Spark using Scala, Python, R or SQL.

Toree has been incubating since 2015-12-02.

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

1. Resolve LGPL dependency: RESOLVED! Library is now released as MPL v2 and necessary changes made in Toree
2. Make a release: The community is working on RC from 0.1.x branch. Master has moved to start support for Spark 2.0
3. Grow a diverse community: We should put some emphasis on growing the community and making it diverse (the rule is at least three independent contributors) In progress. Project elected new PPMC member.

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

NONE, previous issue with LGPL dependency has been RESOLVED as the JeromQ dependency has now been released as MPL license.

How has the community developed since the last report?

1. The Apache Toree has elected a new PPMC member: Marius van Niekerk
1. Active communication in mailing list and gitter with early adopters

2. Still working on transitioning users from Spark Kernel project into Toree. Making significant progress here.
3. More external contributions being made. Spark 2 port done by outside contributor and merged into master

How has the project developed since the last report?

1. JeroMQ community has finalized work to transition to MPL. Release available and changes made in Toree src
2. Addressing issues opened by community

Date of last release:

None since incubation.

When were the last committers or PMC members elected?

No new additions since incubation

Signed-off-by:

[x](toree) Luciano Resende
[](toree) Reynold Xin
[x](toree) Hitesh Shah
[X](toree) Julien Le Dem

Traffic Control

Traffic Control allows you to build a large scale content delivery network using open source.

Traffic Control has been incubating since 2016-07-12.

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

1. Move code, issues, email lists and release process to ASF

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?

All CCLAs and ICLAs have been signed, and we are clear to proceed with moving the source from github.com to Apache. We are cleaning up the issues and open PRs at this time, and should be able to move the code before mid October.

How has the project developed since the last report?

- All legal hurdles for software grant have been cleared.

Date of last release:

We are trying to release 1.7 using the old process and git before the move in to Apache git.

When were the last committers or PMC members elected?

2016-07-12 - initial committer list.

Signed-off-by:

[x](trafficcontrol) Phil Sorber
[](trafficcontrol) Eric Covener
[](trafficcontrol) Daniel Gruno
[](trafficcontrol) J. Aaron Farr

Shepherd/Mentor notes:

Drew Farris (shepherd): Still getting off the ground, mailing lists are new. No observable mentor activity.

Phil Sorber (mentor): There were some struggles getting all the legal ducks in a row but that's all complete now so expect decent uptick in activity.