

November2015

/!\ *FINAL* (!)

Incubator PMC report for November 2015

Timeline

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

Shepherd Assignments

Andrei Savu	Cotton
Andrei Savu	TinkerPop
Drew Farris	Twill
John Ament	Climate Model Diagnostic Analyzer
John Ament	HORN
John Ament	Unomi
Justin Mclean	HAWQ
Matthew Franklin	DataFu
Matthew Franklin	MADlib
Matthew Franklin	Mynewt
P. Taylor Goetz	AsterixDB
P. Taylor Goetz	Concerted
P. Taylor Goetz	REEF
Raphael Bircher	Apex
Raphael Bircher	Sirona
Ross Gardler	Ripple
Suresh Marru	Blur
Suresh Marru	Kylin
Suresh Marru	Slider
Timothy Chen	Rya
Timothy Chen	Tamaya

Report content

Incubator PMC report for November 2015

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

There are 47 podlings currently undergoing incubation.

With 702 messages to `general@incubator`, October was one of the heaviest months for email traffic in a long time.

Podlings nearing graduation received extra scrutiny this month, presumably as an indirect result of various recent proposals to rework graduation to include more structured review. Such evaluations of podling readiness, while illuminating, put stress on the Mentors and contributors to the podling under the microscope and raise the overall tension in the Incubator.

Another topic of discussion was disengaged Mentors. An initiative to engage Mentors privately when podlings do not report for two months or more yielded positive results, with some Mentors reaffirming their commitment and others officially moving on.

* Community

New IPMC members:

- Tom Barber (`magicaltrout`)
- Patrick Wendell (`pwendell`)
- Reynold Xin (`rxin`)
- Phil Sorber (`sorber`)
- Julien Le Dem (`julien`)
- Jacques Nadeau (`jacques`)

No one left the IPMC this month, but a handful of Mentors have stepped down from their posts with various podlings. Most of the resignations were from Mentors who had become inactive, so the Incubator's rolls have become a bit more accurate.

* New Podlings

- Eagle
- Mynewt
- Concerted
- SystemML

* Graduations

The board has motions for the following:

- Brooklyn
- Groovy
- Kylin
- REEF

* Releases

The following releases were made since the last Incubator report:

- 2015-10-05 Apache Johnzon 0.9.2-incubating
- 2015-10-07 Apache Singa 0.1.0
- 2015-10-14 Apache AsterixDB 0.8.7-incubating
- 2015-10-14 Apache REEF 0.13.0-incubating
- 2015-10-23 Apache Kylin 1.1-incubating
- 2015-10-28 Apache TinkerPop 3.0.2-incubating
- 2015-10-30 Apache Apex v3.2.0-incubating
- 2015-10-31 Apache Slider 0.81.1-incubating

* IP Clearance

- Taobao (China) Software Co., Ltd and several individuals donated Alibaba JStorm, "A fork of Apache Storm with Clojure code replaced with Java, and several other improvements".

A discussion calling into question what responsibilities -- if any -- the Incubator should have with regards to IP Clearance eventually went quiet, leaving the status quo intact.

* Miscellaneous

- Droids has retired.
- Kalumet has retired.
- The Corinthia community is voting on retirement.

* Credits

- Report Manager: Marvin Humphrey

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

* Still getting started at the Incubator

- Concerted
- Mynewt
- Rya

* Not yet ready to graduate

No release:

- AsterixDB
- Climate Model Diagnostic Analyzer
- Cotton
- MADlib
- DataFu
- HAWQ
- HORN

Community growth:

- Apex
- Blur
- Slider
- Tamaya
- TinkerPop
- Twill

* Ready to graduate

- Kylin
- REEF

* Did not report, expected next month

- Ripple (5 months late)
- Sirona
- Unomi

Table of Contents

Apex
AsterixDB
Blur
Climate Model Diagnostic Analyzer
Concerted
Cotton
DataFu
HAWQ
HORN
Kylin
MADlib
Mynewt
REEF
Ripple
Rya
Sirona
Slider
Tamaya

TinkerPop
Twill
Unomi

Apex

Apex is an enterprise grade native YARN big data-in-motion platform that unifies stream processing as well as batch processing.

Apex has been incubating since 2015-08-17.

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

1. Release Apex-Core, and Apex-Malhar on ongoing basis.
2. Grow contributors beyond the initial set.
3. Move to using Apache JIRA

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

The project is still using the original Atlassian instance of JIRA. INFRA-10144 has tracked the Apache JIRA migration request since August.

How has the community developed since the last report?

The community is very engaged with the development of the project. There have been 788 messages on dev@ for October.

The community has been active building meetup groups in various locations: <http://s.apache.org/jKT>

- Number of meetups: 6 in total so far: 4 in USA cities and 2 in Indian cities
- Total members in worldwide Apache Apex meetups: 495
- Events: 1 in Oct; 4 scheduled in Nov: CapitalOne presentation in Chicago; DataTorrent presentation in Pune; GE presentation in San Ramon; PubMatic presentation in San Jose;

Future: Planning to start meetups in more cities, and be able to organize 2-3 events a month.

How has the project developed since the last report?

The first release under incubation passed voting. Apex (Core) 3.2.0-incubating was released 2015-10-30. The community is currently discussing roadmap for upcoming releases.

Various metrics are as follows:

Metric	Core	Malhar
Non Merge Commits	159	39
Contributors	15	9
Jira New Issues	58	19
Resolved Issues	34	30

The next release target is Apex Malhar with release candidate expected first week of November.

Date of last release:

2015-10-30 Apex (Core) 3.2.0-incubating

When were the last committers or PMC members elected?

Committers and Mentors came about via incubation proposal.

Signed-off-by:

[X](apex) Chris Nauroth
[X](apex) Alan Gates
[X](apex) Hitesh Shah
[X](apex) Justin Mclean
[X](apex) P. Taylor Goetz
[X](apex) Ted Dunning

Shepherd/Mentor notes:

Chris Nauroth (cnauroth):

The community continues to show great, open development activity on the mailing list. The first release was handled well with careful attention to detail when IPMC members requested a second RC with some changes. I would like to request that the infra team prioritize completion of the JIRA migration tracked in INFRA-10144.

Alan Gates (gates):

There have been some questions regarding which committers in the project are also on the PPMC. Originally only 6 developers were proposed for initial committers, but as part of the entry to incubation this was expanded to the current 30+. The understanding of some in the project was that the PPMC still only consisted of those original 6, while the other ~24 devs were just committers. The mentors have explained that it does not work this way. The community is talking through how to approach this disconnect. The mail thread is at <http://s.apache.org/qsV>.

AsterixDB

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

AsterixDB has been incubating since 2015-02-28.

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

1. Do an Apache release that can also provide binary artifacts with correct LICENSEs and NOTICEs.
- 2.
- 3.

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

The discussion on the use of an external Gerrit instance for code reviews is considered to be resolved.

The current workflow ensures that

- a) The ASF git repository is the canonical repository for the project.
- b) Only committers can manually commit changes to the ASF git repository.
- c) All discussions on the Gerrit instance are sent to notifications@asterixdb.incubator.apache.org.

While it would be good to also have the Gerrit instance run on ASF hardware, this is neither necessary to document the provenance of the code nor to enable list members to follow and participate in the activities of the project. Also, there seems to be no infra capacity to provide a Gerrit service.

How has the community developed since the last report?

- A paper on AsterixDB was presented at VLDB'15.
- A deep-dive introduction to AsterixDB was given at the first workshop on Big Data Open Source Systems (BOSS'15).
- Chris Hillery was added as a committer on Oct 20.
- Increased traffic on the users list.

How has the project developed since the last report?

- All Google Code issues were migrated to the ASF JIRA and triaged.
- The discussion on the use of an external Gerrit instance for code reviews is considered to be resolved.
- AsterixDB 0.8.7 and the associated Hyracks 0.2.16 were released.

Date of last release:

2015-10-22 AsterixDB 0.8.7
2015-09-22 Hyracks 0.2.16

When were the last committers or PMC members elected?

2015-10-20 Chris Hillery

Signed-off-by:

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

Shepherd/Mentor notes:

Chris Mattmann (mattmann):

Great work!

Ted Dunning (tdunning):

The handling of the Gerrit question was particularly good.

Blur

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

Blur has been incubating since 2012-07-24.

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

1. Greater community involvement.
2. Produce releases.
- 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?

- Subscriptions: user@ - 55[0]; dev@ - 70[+3]

How has the project developed since the last report?

- A few additional features, and bug fixes.

Date of last release:

2014-07-29

When were the last committers or PMC members elected?

2014-07-28

Signed-off-by:

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

Shepherd/Mentor notes:

Climate Model Diagnostic Analyzer

CMDA provides web services for multi-aspect physics-based and phenomenon-oriented climate model performance evaluation and diagnosis through the comprehensive and synergistic use of multiple observational data, reanalysis data, and model outputs.

Climate Model Diagnostic Analyzer has been incubating since 2015-05-08.

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

1. Increase email and communication traffic
2. Build a CMDA website at Apache
3. Learn more about Apache way

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

The project is still in its nascence and the PPMC is still learning the Apache way and moving out of JPL/CMU centric development. The project team wants to do a good job they are just moving along slowly.

How has the community developed since the last report?

More commits and code are being done on the Apache Git repo, almost exclusively. They still come in bulk rather than incrementally. Also notes from meetings held between JPL and CMU project members are being flushed to the list to increase awareness and participation.

How has the project developed since the last report?

1. Deployed to the Amazon Cloud (AWS);
2. Students from the 2015 JPL center for climate sciences summer school used the Amazon Cloud deployment for class projects;
3. New design and implementation of front end web interface for flexibility and to facilitate provenance;
4. Acquired new ECMWF data;
5. Created a knowledge graph based on provenance accumulated from the 2015 Summer School usages;
6. New services to come soon: Random Forests and Empirical Orthogonal Function Analysis (EOF).

More code is coming. Need to do a website, and start to use issue tracking, etc.

Date of last release:

When were the last committers or PMC members elected?

Signed-off-by:

[](climatemodeldiagnosticanalyzer) James W. Carman
[X](climatemodeldiagnosticanalyzer) Chris Mattmann
[X](climatemodeldiagnosticanalyzer) Michael James Joyce
[](climatemodeldiagnosticanalyzer) Kim Whitehall
[X](climatemodeldiagnosticanalyzer) Gregory D. Reddin

Shepherd/Mentor notes:

John D. Ament (johndament):

There are a couple of concerns raised in this report, specifically that ASF infrastructure is not the canonical source of truth for the podling and limited external participation.

Concerted

Apache Concerted is a Do-It-Yourself toolkit for building in-memory data engines.

Concerted has been incubating since 2015-10-14.

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

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

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

- None at this time.

How has the community developed since the last report?

Latest Additions:

- * PMC addition: NA
- * Contributor addition: NA

Issue backlog status since last report:

- * Created: 20
- * Resolved: 8

Mailing list activity since last report:

- * @dev 133 messages
- * @issues 129 messages
- * @commits 8 messages

How has the project developed since the last report?

- Project has been focused on bootstrapping and getting setup
- Status page setup, mailing lists created, jira setup
- Grant filed and code imported, website in development
- Apache headers applied to imported source
- Working on scripts to automate the release process for our first Apache Concerted release candidate

Date of last release:

- Working towards our first release candidate.

Signed-off-by:

[X](concerted) Chris Nauroth
[](concerted) Daniel Dai
[X](concerted) Jake Farrell
[X](concerted) Julian Hyde
[](concerted) Lars Hofhansl

Shepherd/Mentor notes:

Julian Hyde (jhyde):

Good level of activity; beginning to engage with community beyond initial committers.

Cotton

Cotton is an Apache Mesos framework for running MySQL instances.

Cotton has been incubating since 2015-06.

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

1. Increase project communication on mailing lists, JIRA
2. Attract new project contributors and committers
3. Establish a first release at Apache

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

Development of Cotton has largely been inactive for the last two months, due to limited time from the original developer and community engagements including MesosCon Seattle and MesosCon Europe.

How has the community developed since the last report?

There has been limited participation on the mailing lists and JIRA, however a discussion on the private@ list was established about this issue and there's a desire to keep the project in the incubator and be more proactive in building community in the coming months.

How has the project developed since the last report?

- * Apache Cotton was presented at MesosCon, available for viewing on YouTube: <https://www.youtube.com/watch?v=BFS8-mt8alU>
- * Limited development has occurred since the last report.

Several JIRA issues outline some of the earliest outstanding work, including setting up a project website and renaming portions of the codebase to reflect the new Apache Cotton name (formerly Mysos).

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](cotton) Jake Farrell
[x](cotton) Dave Lester
[](cotton) Benjamin Hindman
[X](cotton) Henry Saputra

Shepherd/Mentor notes:

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. Do first release
2. Grow user and contributor base

3. Increased committer activity

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

- * Performing the initial release remains the most important milestone.

How has the community developed since the last report?

- * No new activity in the community since the last report.

How has the project developed since the last report?

- * The website documentation (<http://datafu.incubator.apache.org/>) has been updated and brought up to date with the current state of the project and build system, making it easier for newcomers to get started. This was the last major task blocking release.
- * All the release tasks filed for our first release have now been completed. A discussion has been opened in the dev mailing list on the topic of doing our first release. A vote will likely be held in the next few days.

Date of last release:

- * Not yet released. First release will likely happen within the coming weeks.

When were the last committers or PMC members elected?

- * November 2014

Signed-off-by:

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

Shepherd/Mentor notes:

HAWQ

HAWQ is an advanced enterprise SQL on Hadoop analytic engine built around a robust and high-performance massively-parallel processing (MPP) SQL framework evolved from Pivotal Greenplum Database.

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, increase dev list activity and adding new committers/pmc members
3. Infrastructure migration: create Jenkins projects that build HAWQ binary, source tarballs and docker images, 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?

We have only just started the incubation, everything seems to be smooth, nothing urgent at this time.

How has the community developed since the last report?

1. Most of the core contributors/committers have started working on Apache repository
2. One Meetup for HAWQ hosted by "Big Data Community" in Beijing.
3. The community is active with the development of the project. In Oct,

There have been 357 messages on dev@. and 21 messages on user@, compared with 95 messages in total last month.

4. Community shows interests to integrate HAWQ with systems in the ecosystem.

How has the project developed since the last report?

1. First release has been proposed and targeted in Nov (<http://s.apache.org/p4L6>)
2. 95 JIRAs filed, 46 open, 30 resolved, 19 closed (In Oct 2015)
3. 58 code commits (In Oct 2015)

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

Shepherd/Mentor notes:

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. Finalize Architecture and API design
2. Continue to gain and attract contributors
3. Produce a first Apache release

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

Nope

How has the community developed since the last report?

- We're receiving contributions from new contributors Elmulod and Zach :-)
- All PPMC members are in both private@ and dev@ and signed ICLA.
- All required infrastructures e.g., mailing lists, JIRA, Git repo are now available.
- dev@horn.incubator.apache.org: 32 subscribers.

How has the project developed since the last report?

- Development of project website has for the most part been completed.
- Initial code import from Apache Hama ML package has been completed.
- The System Architecture and Programming APIs design planning is currently in progress.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Kylin

Kylin is a distributed and scalable OLAP engine built on Hadoop to support extremely large datasets.

Kylin has been incubating since 2014-11-25.

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

1. Passed graduation vote, already sent proposal to next board meeting
- 2.
- 3.

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

* Nothing

How has the community developed since the last report?

1. In Oct 2015, there were 308 messages on dev@, 56 JIRA new issues be created, 111 issues have been resolved
2. Seshu Adunuthula and Qianhao Zhou has presented Apache Kylin at ApacheCon Big Data EU 2015 on Sep 28, 2015
3. Luke Han and Shaofeng Shi have presented two sessions at Apache Roadshow China 2015 in Beijing on Oct 24, 2015
4. Apache Kylin Meetup has been organized in eBay Shanghai office with 100+ attendees including presenters from Kylin, Tez and Zeppelin committers on Oct 10, 2015
5. Apache Kylin won InfoWorld Bossie Awards 2015 - The Best Open Source Big Data Tools.

How has the project developed since the last report?

1. Released 1.0 and 1.1, is preparing release 1.2 and 2.0
2. Introduced Fast Cubing (MR) and Spark Cubing engine
3. Support setting for HBase compression with Snappy or GZip
4. Support load data to separated HBase cluster
5. Apache Kylin Interpreter has been submitted to Apache Zeppelin main code base and be included in Zeppelin's release.
6. Fixed slowness with many IN() values
7. Fixed AVG not work issue
8. Upgraded Calcite to 1.4
9. Web UI refined and introduced new for 2.x branch
10. other Bug fixes

Date of last release:

2015-10-25 apache-kylin-v1.1-incubating
2015-09-06 apache-kylin-v1.0-incubating

When were the last committers or PMC members elected?

2015-08-18 1 new committer, Dayue Gao from meituan.com
2015-08-25 1 new committer, Hua Huang from MiningLAMP
2015-09-25 1 new committer, Yerui Sun from meituan.com

Signed-off-by:

[X](kylin) Owen O'Malley
[X](kylin) Ted Dunning
[X](kylin) Henry Saputra
[X](kylin) Julian Hyde
[X](kylin) P. Taylor Goetz

Shepherd/Mentor notes:

Ted Dunning (tdunning):

There was a bit of confusion about the graduation vote (which passed, but not unanimously). IPMC and Kylin project members are working to fix the issues to the satisfaction of the IPMC member who voted -1.

Marvin Humphrey (marvin):

The issue that was raised by the IPMC Member voting -1 was acknowledged as legitimate by multiple IPMC members, and a plan of action was proposed which was endorsed by one of the Kylin Mentors. So while there may have been differences of opinion about urgency, the matter has been guided steadily towards a constructive resolution.

(update) The -1 has been changed over to a +1 and the issue resolved.

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. Produce a first Apache (incubating) release.
2. Expand the community, increase dev list activity and add new committers/pmc members.
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. Meetup 10/1/15 @ Pivotal Labs, New York, NY: "MADlib and HAWQ for Advanced SQL Machine Learning on Hadoop" <http://s.apache.org/VbG>
2. Meetup 10/29/15 @ Pivotal Palo Alto, CA: "Data Science at Scale for IoT" <http://www.meetup.com/Pivotal-Open-Source-Hub/events/225426787/>

How has the project developed since the last report?

1. All known issues related to IP cleanliness described in <https://wiki.apache.org/incubator/MADlibProposal> have been fixed and pushed to the Apache repo.
2. All software activity tracking has migrated to Apache MADlib JIRA from previous tool. 18 JIRAs created and 2 resolved in last 30 days.
3. All commits and code are now being done on the Apache Git repo.
4. Three new quick start guides have been written: i) install, ii) user, and iii) developer. The goal is to make it easier to onboard new community members.
5. A new Greenplum DB sandbox VM with MADlib pre-installed has been created and made available publicly at <https://github.com/greenplum-db/gpdb-sandbox-tutorials>. The goal is to make it easier to onboard new community members - they can download and start trying MADlib right away with no install/setup.
6. A "catchup JIRA" was filed <https://issues.apache.org/jira/browse/MADLIB-912> in order to catch up

between the time of the code grand to Apache and bringing in dev work that was already in flight at the time. We apologize for any inconvenience in clubbing together these multiple items; it was a one-time operation.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

No new members added on top of the initial committer list.

Signed-off-by:

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

Shepherd/Mentor notes:

Konstantin Boudnik (cos):

I don't see much info on the community development. How many new contributors the project had gained? Were there any additions in the mailing lists? Please consider providing this information in the next report.

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. Produce a first Apache release of downloadable RTOS image with support for multiple peripherals and network connectivity.
2. Develop a roadmap and release schedule for the Mynewt Operating System
3. Expand the community, increase dev list activity, and add new committers/pmc members.

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

Just started incubation, nothing specific to report at this time.

How has the community developed since the last report?

We are approximately 2 weeks into the incubation process.

1. Core infrastructure has been set up. Core contributors have completed their ICLA's. Apache accounts, mailing lists, JIRA, and website are in place.
2. First drop of documentation is linked to the incubator website on AFS.
3. Participation on the @dev mailing list was seen with installation queries and troubleshooting questions from mentors and 1 new subscriber.

How has the project developed since the last report?

Early activity:

1. Initial code drop on Apache git repository for Mynewt RTOS for M4 and M3 architectures in debug environment
2. Initial code drop on Apache git repository for newt tool to build basic RTOS images
3. Core infrastructure is set up

Date of last release:

No release yet.

When were the last committers or PMC members elected?

Initial list of committers have been on-boarded.
No new members added on top of the initial committer list.

Signed-off-by:

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

Shepherd/Mentor notes:

REEF

REEF (Retainable Evaluator Execution Framework) is a scale-out computing fabric that eases the development of Big Data applications on top of resource managers such as Apache YARN and Mesos.

REEF has been incubating since 2014-08-12.

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

1. Vote to graduate REEF to a top level project on the PPMC dev@ list passed with 18 binding +1 votes
2. Discussion to graduate REEF to a top level project on the IPMC general@ list is in progress.

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

No.

How has the community developed since the last report?

- * We have seen significant activity from the community since last report
 - 206 JIRA issues created
 - 179 JIRA issues resolved
 - 1300+ emails on dev (including auto-generated messages)
- * New committers have been added to the team
 - Geon-Woo Kim from Seoul National University
 - Dhruv Mahajan from Microsoft
- * Markus Weimer visited multiple European universities and institutions to present REEF and discuss potential collaborations.

How has the project developed since the last report?

- * 0.13.0-incubating is released
 - release date: October 14, 2015
 - vote: 3 +1 binding votes and 0 -1 votes
 - 187 issues, tasks, features were resolved in this release.

Date of last release:

2015-10-14
(0.13.0-incubating)

When were the last committers or PMC members elected?

Geon-Woo Kim: 2015-09-15
Dhruv Mahajan: 2015-09-09

Signed-off-by:

[X](reef) Chris Douglas
[](reef) Chris Mattmann
[](reef) Ross Gardler
[](reef) Owen O'Malley

Shepherd/Mentor notes:

Chris Mattmann (mattmann):

REEF voting on graduation.

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. Populate the website for the project
2. Become familiar with the release process and have a first release as part of the Apache
3. Create a "how to" document to explain the ways new contributors can become involved in the project, so we increase the size of the community.

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

No

How has the community developed since the last report?

All initial committers submitted their ICLAs and are subscribed to the dev and private mailing lists. There are 32 subscribers to the dev list. The initial committers are learning to use the Apache infrastructure and processes. We gave a presentation on Rya at Duke University, trying to increase awareness about the project.

How has the project developed since the last report?

Infrastructure is now set up (mailing lists, repositories, JIRA, etc) Getting the code reviewed for final publication took about 3 weeks, which was longer than we expected, but the code is now imported into the Apache git repository. We are working on creating a website for the project.

Date of last release:

Not applicable - the code was just imported into Apache repository

When were the last committers or PMC members elected?

Not applicable

Signed-off-by:

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

Shepherd/Mentor notes:

Sean Busbey (busbey):

Getting things started took a bit, presumably due to the review needed to get the code out, but the fledging community has been responsive to feedback. The next month should give a better idea of how strong a hand the mentors will need to use.

Slider

Slider is a collection of tools and technologies to package, deploy, and manage long running applications on Apache Hadoop YARN clusters.

Slider has been incubating since 2014-04-29.

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

1. Getting more external users
2. Getting more diverse set of developers
3. Getting more diverse set of committers/PMC

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

No

How has the community developed since the last report?

We are seeing traction in the community with a number of new users requesting improvements in deployment and management of applications using Slider. Patches were submitted by a couple of new contributors. We still have to get those people into long term coding, and then bring them in to the committer group.

How has the project developed since the last report?

We released a slider-0.81.1-incubating bug fix release with 30 JIRAs resolved and have begun discussing the 0.90.0 release.

Date of last release:

2015-10-29 slider-0.81.1-incubating

When were the last committers or PMC members elected?

2015-07-07: Yu (Thomas) Liu

Signed-off-by:

[](slider) Arun C Murthy
[X](slider) Devaraj Das
[](slider) Jean-Baptiste Onofré
[X](slider) Mahadev Konar

Shepherd/Mentor notes:

Tamaya

Tamaya is a highly flexible configuration solution based on an modular, extensible and injectable key/value based design, which should provide a minimal but extendible 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. Make the project known in projects world-wide.

2. Release more often.
- 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?

We have new people joining the mailinglist and also starting discussions for possible future involvement. We had quite broad feedback in the Java EE space, since we added full blown CDI support to Tamaya. We are looking forward to gain new committers within next reporting period hopefully.
During ApacheCon Europe different projects found our approach very interesting. We also have a request from a JUG in Switzerland going to present the project as well several conferences that accept the topic as well.

How has the project developed since the last report?

The project's API has proven to work very well. Many, also complex extensions were written and we started to add support for very important eco-systems (JavaEE, Spring, OSGI).
I think we can release the next release before X-Mas the latest.

Date of last release:

2015-08-22

When were the last committers or PMC members elected?

See last report (no new committers or PMC members).

Signed-off-by:

[X](tamaya) John D. Ament
[X](tamaya) Mark Struberg
[](tamaya) Gerhard Petracek
[](tamaya) David Blevins

Shepherd/Mentor notes:

TinkerPop

TinkerPop is a graph computing framework written in Java

TinkerPop has been incubating since 2015-01-16.

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

1. Add more members to the PMC.
2. Add more members to the committer list.
3. N/A (unless the mentors have ideas?)

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

No. Recently, we have really appreciated Daniel Gruno's (mentor) efforts to get our development community more connected by setting up an Apache HipChat account for us and interacting with us in a more real-time fashion on various procedures/policies of Apache in a proactive manner. This has helped to explain to us (through doing) what is required of Apache.

How has the community developed since the last report?

We have 2 new committers: Jason Plurad and Matthew Frantz. We are in the

process of [DISCUSS] about two other potential committers.

How has the project developed since the last report?

We have since instantiated a 'review-then-commit' model. Daniel Gruno (mentor) worked with us to set the appropriate policy documentation and thus far, it has been going very well.

In terms of the adoption of TinkerPop, note that Amazon recently announced that they use Apache TinkerPop for their order fulfillment network (1 trillion edges). This is very public discussion of our technology (videos + blog posts) should help to attract development talent.

We have compiled a list of the other Apache technologies that TinkerPop works with: Apache Cassandra, Apache HBase, Apache Hadoop, Apache Spark, Apache Giraph, Apache Atlas, and Apache Falcon. We bring this up to identify the integration of Apache TinkerPop into the greater Apache ecosystem.

The following articles/presentations/blogposts were provided about TinkerPop from TinkerPop members since the last report:

- * Rodriguez, M.A., "The Gremlin Graph Traversal Machine and Language" <<http://arxiv.org/abs/1508.03843>>, ACM Database Programming Languages Conference Proceedings, October 2015.
- * Mallette, S.P., "What's New In Apache TinkerPop?" <<http://www.slideshare.net/StephenMallette/tinkerpopfinal>>, Cassandra Summit, September 2015.
- * Rodriguez, M.A., Kupptitz, D., "The Benefits of the Gremlin Graph Traversal Machine" <<http://s.apache.org/0vW>>, DataStax Engineering Blog, September 2015.

Date of last release:

2015-09-16

When were the last committers or PMC members elected?

Jason Plurad (9/30) and Matthew Frantz (7/10). Both committers.

Please see out Project Status page as we have been diligent to update it accordingly.

<http://incubator.apache.org/projects/tinkerpop.html>

Signed-off-by:

[x](tinkerpop) Rich Bowen
[x](tinkerpop) Daniel Gruno
[](tinkerpop) Hadrian Zbarcea
[](tinkerpop) Matt Franklin
[](tinkerpop) David Nalley

Shepherd/Mentor notes:

Daniel Gruno (humbedooh):

I will note that despite HipChat and other direct message platforms involved, the async messages, especially revolving around larger community/project issues have been on the rise over the past 3 months, up 71% compared with the 3 months before that, as evident from <<http://s.apache.org/Atz>> (this is discounting git and jira messages). I see this as a positive signal that the podling is working towards a more open and inclusive discussion pattern. New committers are being on-boarded, and I hope to see more PPMC members join the ranks soon.

Twill

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

Twill has been incubating since 2013-11-14.

Top three items to resolve before graduation:

- More engagement from the community.
- Improve documentations and examples.
- 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?

- 7 new JIRA issues filed since last report
- 8 JIRA issues resolved since last report
- 2 new contributors submitted patches

How has the project developed since the last report?

- Version 0.7.0-incubating is being worked upon
- Fixed issues to support Microsoft Azure

Date of last release:

- 2015-07-24: 0.6.0-incubating

What are the plans for the next period?

- Encourage contributions from active users
- Identify potential committers
- Engage more on social channels (IRC and Twitter)

When were the last committers or PMC members elected?

- August 4, 2015 : Henry Saputra

Signed-off-by:

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

Shepherd/Mentor notes:

Unomi

Unomi is a reference implementation of the OASIS Context Server specification currently being worked on by the OASIS Context Server Technical Committee. It provides a high-performance user profile and event tracking server.

Unomi has been incubating since 2015-10-05.

Shepherd/Mentor notes:

Chris Mattmann (mattmann):

Missing report.

John D. Ament (johndament):

Missing podling status file, no website yet. Really new podling.

