

January2016

/\ *FINAL* (!)

Incubator PMC report for January 2016

Timeline

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

Shepherd Assignments

Drew Farris	Apex		
John Ament	Conce rted		
John Ament	MADlib		
Justin Mclean	Kudu		
P. Taylor Goetz	Batch EE		
P. Taylor Goetz	DataFu		
P. Taylor Goetz	Sirona		
Timothy Chen	FreeM arker		
Timothy Chen	Mynewt		
Timothy Chen	Unomi		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="25ce4990-a7cb-4136-8308-a418e2ee44c9"><ac:plain-text-body><![CDATA[[none]	Eagle]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c36f0657-7b42-45e2-ac7e-cdfdc743c6f2"><ac:plain-text-body><![CDATA[[none]	Finera ct]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6d7caf47-8cc0-414a-9383-316f91fbc0e2"><ac:plain-text-body><![CDATA[[none]	Geode]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7875ec25-61f7-4861-a1da-d71b93dccb39"><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="3b3ddfb2-5ed6-45d9-8229-ea39641d4011"><ac:plain-text-body><![CDATA[[none]	HORN]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a897f3db-03b9-4d98-b0a5-c25f21824e3b"><ac:plain-text-body><![CDATA[[none]	Impala]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d1363146-048a-4510-9801-497d5d33fa41"><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="63171224-6196-4e88-9515-f44d5a9d65ca"><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="097807f8-fe28-4ac1-ae0c-9db1117a37b8"><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="a60adca8-1003-4a87-945f-ac0dd62c80c0"><ac:plain-text-body><![CDATA[[none]	OpenAz]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b80ee910-1988-42bc-aae9-3463b31981bb"><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="0c7608d0-fcae-4b37-9718-9035a33fcd5"><ac:plain-text-body><![CDATA[[none]	S2Gra ph]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="33f6e68e-a797-46f2-96b3-9d80f502bb39"><ac:plain-text-body><![CDATA[[none]	Syste mML]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="81647880-c1d3-464d-afa8-518f4090ab28"><ac:plain-text-body><![CDATA[[none]	log4cx x2]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for January 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 48 podlings currently undergoing incubation.

The end of the year was relatively quiet for the Incubator.

* Community

New IPMC members:

- James Taylor
- Nick Kew

* New Podlings

- Fineract
- Milagro

* Releases

The following releases entered distribution during the month of December:

- 2015-12-08 Apache Atlas 0.6-incubating
- 2015-12-09 Apache Myriad 0.1.0-incubating
- 2015-12-28 Apache Slider 0.90.2-incubating

Wave's release candidate has set a record for frustration. It was presented to general@incubator on November 3rd 2015, and still has only one IPMC +1 vote over two months later.

* Infrastructure

Despite the heroics of Infra, JIRA's temperamental imports have occasionally delayed migrations for new podlings over the years. SystemML is the latest podling to experience difficulties.

* Miscellaneous

- The Log4CXX podling has no active Mentors, but a volunteer has stepped forward. The situation is developing.

* Credits

- Report Manager: Marvin Humphrey

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

* Still getting started at the Incubator

- Fineract
- Impala
- Kudu
- Metron
- Milagro

* Not yet ready to graduate

No release:

- Eagle
- Geode
- HAWQ
- MADlib
- Mynewt
- Rya
- S2Graph
- SystemML
- Unomi

Community growth:

- Apex
- BatchEE
- FreeMarker
- Sirona

* Ready to graduate

- DataFu

* Did not report, expected next month

- Concerted
- HORN
- ODF Toolkit
- OpenAZ

* In crisis

- log4cxx2

Table of Contents

Apex
BatchEE
Concerted
DataFu
Eagle
Fineract
FreeMarker
Geode
HAWQ
Impala
Kudu
log4cxx2
MADlib
Metron
Milagro
Mynewt
Rya
S2Graph
Sirona
SystemML
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. Continue release Apex-Core, and Apex-Malhar on ongoing basis.

2. Continue grow contributors beyond the initial set.

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

No issues, thanks for your support.

How has the community developed since the last report?

The community is very engaged with the development of the project.

We see continuous uptick in mailing list participation (1392 messages on dev@ for December, 86+ subscribers).

We have seen new contributors and work to integrate with other Apache ecosystem projects such as Geode and NiFi.

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

Meetup membership grew to ~1000 in less than 4 months

How has the project developed since the last report?

Release Malhar 3.2.0-incubating on 2015-11-16
Roadmap published on web site:
<http://apex.incubator.apache.org/roadmap.html>
JIRA migration completed for Core and Malhar.

Various metrics for December are as follows:

+-----+			
Metric	Core	Malhar	
+-----+			
Non Merge Commits	17	18	
Contributors	11	10	
Jira New Issues	18	32	
Resolved Issues	25	14	
Pull Requests merged	8	16	
Pull Requests proposed	5	18	
+-----+			

We are anticipating releases for Apex Core and Malhar in January.

Date of last release:

2015-11-16 Apex (Malhar) 3.2.0-incubating
2015-10-31 Apex (Core) 3.2.0-incubating

When were the last committers or PMC members elected?

1 new PPMC member and 2 committers were added in December.
PPMC and committer list was updated on the status page.

Due to omission in the incubation proposal initial committers voted to adjust PPMC membership as per original intention. The vote result can be found here: <http://s.apache.org/jeJ>

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:

There is a very high level of activity across mailing lists and JIRA.

Development is clearly happening in an open, collaborative way.

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. increase the community
2. enhance the documentation
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?

Activity on the mailing lists did increase a bit.

How has the project developed since the last report?

Several enhancement in Mojo (making it easier to use with Gradle), several fixes and enhancements in components have been done.

Date of last release:

batchee-0.3-incubating on 2015-11-12

When were the last committers or PMC members elected?

Reinhard Standtner hss been added as committer on the 2nd dec.

Signed-off-by:

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

Shepherd/Mentor notes:

Concerted

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

Concerted has been incubating since 2015-10-14.

Shepherd/Mentor notes:

Jake Farrell (jfarrell):

There has been very little activity on the mailing lists, John D. Ament started a thread about the project's health.

Chris Nauroth (cnauroth):

+1 for concerns on lack of activity. Attempting to rally PPMC.

Julian Hyde (jhyde):

I have the same concerns about low activity level. Allegedly there is *private* activity (re-factoring code to ready it for submission) but I

am encouraging them to show *public* activity as well.

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?

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

How has the project developed since the last report?

- * Apache DataFu 1.3.0 source release completed, which is the first release since entering the Incubator. DataFu 1.3.0 was also released to Maven.
- * Website (<http://datafu.incubator.apache.org/>) has been updated with instructions on how to use the source release or artifacts from Maven.

Date of last release:

* 2015-11-14

When were the last committers or PMC members elected?

Signed-off-by:

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

Shepherd/Mentor notes:

Roman Shaposhnik (rvs):

The community appears to be in the final stretch before graduation, hopefully there's enough critical mass for it to happen.

Eagle

Eagle is a Monitoring solution for Hadoop to instantly identify access to sensitive data, recognize attacks, malicious activities and take actions in real time.

Eagle has been incubating since 2015-10-26.

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

1. Build the momentum with the community.
2. Create an apache release.
3. Work with the community on a roadmap and clarify the vision.

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

Nothing urgent at this moment.

How has the community developed since the last report?

1. We have submitted papers and applied to conferences. We worked with other open source projects like Siddhi, Apache Ranger during December.
2. Dev community showed increasing interests. In Dec, there have been 878 messages on dev@ mailing list. It is an increase of 30% mom. With more quality discussions around integration, features and release process.
3. There were increased participation in the dev mailing list pertaining to new feature development and roadmap.

There was a request from the community about the roadmap and clear picture of the vision. We have published the first roadmap to the community and are working on it.

How has the project developed since the last report?

1. Main features/improvements added include:
 - a) EAGLE-53 Finalized the docker image by fixing few bugs and scripts.
 - b) EAGLE-32 Worked on creating a generic monitoring framework.
 - c) EAGLE-47 Policy audit module
 - d) Bug fixes & documentation improvement
2. 47 new JIRAs filed, 30 resolved (In Dec 2015)
3. 30 code commits (In Dec 2015)

Date of last release:

None

When were the last committers or PMC members elected?

None

Signed-off-by:

[X](eagle) Owen O'Malley
[X](eagle) Henry Saputra
[X](eagle) Julian Hyde
[X](eagle) P. Taylor Goetz
[X](eagle) Amareshwari Sriramdasu

Shepherd/Mentor notes:

Julian Hyde (jhyde):

Activity on the dev list is healthy and trending upwards.

Fineract

Fineract is an open source system for core banking as a platform.

Fineract has been incubating since 2015-12-15.

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

1. Push the initial source code to Apache
2. Inform the non-technical part of the existing Mifos.org community of the change in project structure
3. Getting new committers up to speed on Apache processes and philosophy

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

None, currently.

How has the community developed since the last report?

The project just has started. All members of the initial community are taking their first steps into the Apache Way. The larger Mifos.org community is still operating as it was before, and is a mix of 'clients' who use the platform and services provided by Mifos.org and partners who leverage more on the platform and extend on it. Most partners have already been involved and informed about the plans to become an Apache project, most of whom are part of the current Fineract Apache community.

How has the project developed since the last report?

This is the first report.

Date of last release:

N/A

When were the last committers or PMC members elected?

Given the early state and missing code base, no additional committers or PMC members were elected.

Signed-off-by:

[](fineract) Ross Gardler
[X](fineract) Greg Stein
[X](fineract) Roman Shaposhnik

Shepherd/Mentor notes:

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.

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

1. Finding contributors
2. Do the first release from ASF

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

None

How has the community developed since the last report?

We have one new committer who works on the Chinese documentation.

How has the project developed since the last report?

The web site was moved over to the Apache infrastructure, and the freemarker.org domain was transferred to Apache too.

Preparing for the first release from ASF, issues reported by Rat were fixed, and the build scripts were reworked to give output that conforms the release policy and conventions.

There were a few smaller pull requests which were all reviewed and integrated.

As a side note, there's a yet unresolved legal/policy issue (INFRA-10787) regarding if we can use freemarker.org as the canonical domain (while it's served from the Apache infrastructure), or we must redirect it to freemarker.incubator.apache.org (as it is now).

Date of last release:

There was no release from the Incubator yet

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
[](freemarker) David E. Jones
[](freemarker) Ralph Goers
[X](freemarker) Sergio Fernández

Shepherd/Mentor notes:

Geode

Apache Geode is a distributed, in-memory database with strong data consistency, built to support transactional applications with low latency and high concurrency needs.

Geode has been incubating since 2015-04-27.

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

1. Have our first Apache (incubating) release
2. Expanding the community to include contributors and committers outside of Pivotal.
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?

- The only item pending for the Alpha-1 release is completion of the LICENSE and NOTICE file review (GEODE-610)
- The Geode Clubhouse hosted a meeting covering detailed aspects of OQL/Querying on Geode presented by Jason Huynh. Video to be published on YouTube (<https://www.youtube.com/channel/UCaY2q0U1WjAgEGL7uhCLs6A>)
- The new website is live and with good feedback from community, check it out!
- The breakdown of JIRA tickets is the following:

Q3 / 2015 - 65 created, 19 resolved.
Q4 / 2015 - 342 created, 227 resolved.
Q1 / 2016 - 11 created, 6 resolved.

- There was a total of 66 pull requests on GitHub with 7 still open.
- The breakdown of the mailing lists messages for October, November and December of 2015:

org.apache.geode.issues	2,473
org.apache.geode.commits	2,242
org.apache.geode.dev	1,247
org.apache.geode.user	169

- There are now 143 subscribers on the dev and 135 on the user mailing lists.
- Upcoming community events in January / 2016:
 - FOSDEM 2016 in Brussels:

- Taxi trip analysis (DEBS grand-challenge) with Apache Geode (incubating)
- Big Data meets Fast Data: an scalable hybrid real-time transactional and analytics solution

How has the project developed since the last report?

- GEODE-77 (about JGroups) got merged into develop making a release now possible. This was a huge effort and now JGroups is treated like any other dependency in the project allowing for easy updates when needed.
- There is Agile Dashboard tracking the release progress:
<http://s.apache.org/s0v>
- Builds now runs Apache RAT for license audit thanks to a Gradle plugin. (GEODE-608)
- Jenkins nightly build now produces source and binary artifacts. Namely `apache-geode-src-1.0.0-incubating-SNAPSHOT.tar.gz` and `apache-geode-1.0.0-incubating-SNAPSHOT.tar.gz`
- The code for HTTP session replication and Hibernate modules are now being merged into develop. (GEODE-14)
- The community did a squashathon (GEODE-409 branch) focusing to fix intermitent failures on the CI with great results, increasing the number of Unit and Integration tests. Current numbers are the following, by category:

```
UnitTest: 1,572 tests
IntegrationTest: 2,571 tests
DistributedTest: 7,166 tests
```

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Konstantin Boudnik (cos):

For the next report: for me as a mentor (and perhaps for others on IPMC) it is not very insightful to see the number of pull requests or JIRAs opened. It isn't what "project development" constitutes in my mind. I would like to see more info about how project is getting into Apache Way, what important discussions are held, what decisions are made through community consensus, etc.

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, increase dev list activity and adding new contributors
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?

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

How has the community developed since the last report?

1. One HAWQ Nest community meeting (topic: Data Federation with Apache HAWQ Using the PXF Extension Framework)
Video: <http://bit.ly/HAWQNest1217>, Slides: <http://bit.ly/HAWQNest1217ss>
2. The community shows increasing activities. In Dev, there have been 494 messages on dev@/user@, compared with 459 messages in Nov.

How has the project developed since the last report?

1. Main features/improvements added include:
 - a) Improve recovery performance during some exception cases (HAWQ-241, HAWQ-232)
 - b) Support writable external table executing on a specified number of virtual segment (HAWQ-293)
 - c) Improve HAWQ resource manager resource allocation algorithm and RPC framework (HAWQ-234)
 - d) Support fault injector in 2.0 (HAWQ-288)
 - e) Clean up some IP and build related issues (HAWQ-277, HAWQ-186, HAWQ-184, HAWQ-228)
 - f) Some key bug fixes for 2.0 Beta
2. 2.0 Beta RC2 was proposed
3. 98 new JIRAs filed, 59 resolved (In Dec 2015)
4. 82 code commits (In Dec 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
 [](hawq) Thejas Nair
 [x](hawq) Roman Shaposhnik

Shepherd/Mentor notes:

Justin Erenkrantz:

There has been good progress on improving the licensing situation (HAWQ-271) as well as conversations about what it means to be a contributor (yielding <https://cwiki.apache.org/confluence/display/HAWQ/Becoming+a+committer>). All good and healthy so far.

Konstantin Boudnik (cos):

Same comment as for Geode: the report is almost copy-cut from one another (if not in the letter, but in the spirit). Did you guys had anything worth mentioning besides of new features and bug fixes? Like Justin mentioned above - that what is very important for the project development.

Impala

Impala is a high-performance C++ and Java SQL query engine for data stored in Apache Hadoop-based clusters.

Impala has been incubating since 2015-12-03.

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

1. Resolve any issues around use of Gerrit as code-review tool.
2. Movement of existing JIRA / Git / wiki / e-mail resources to Apache equivalents
3. Initial release as incubating project.

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?

Slowly - Impala is still in the very early stages of incubation, and performing the mechanical tasks of code movement and infrastructure setup is our first priority. The holiday period in the United States has slowed this effort slightly, but we look forward to picking up pace in early 2016. There have been no additions to the committer or PMC lists since incubation began.

How has the project developed since the last report?

We have performed some of the basic initial tasks for incubation - establishing wiki pages, Git repositories and accounts for the initial committer set. Our next steps are:

1. Finalize the SGA from Cloudera
2. Move existing @cloudera.org e-mail aliases to their @impala.incubator.apache.org equivalents.
3. Move source code from Cloudera git repository to Apache git repo.
4. Improve out-of-box build and test experience so that community can easily evaluate release artifacts.
5. Migrate cloudera.org JIRA tickets to issues.apache.org.

Date of last release:

NA

When were the last committers or PMC members elected?

At the time of the Incubation vote, 2015-12-03.

Signed-off-by:

[X](impala) Tom White
[X](impala) Todd Lipcon
[](impala) Carl Steinbach
[](impala) Brock Noland

Shepherd/Mentor notes:

Kudu

Kudu is a distributed columnar storage engine built for the Apache Hadoop ecosystem.

Kudu has been incubating since 2015-12-03.

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

1. Move development to the Apache-hosted git repository, bug tracking, and code review.
2. Build a more diverse development and user community
3. Perform releases under the ASF

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

None so far.

How has the community developed since the last report?

- The development community has remained about the same since the initial proposal, with overall code contribution levels a bit lower than historical averages due to the end-of-year holidays.
- Xiaomi reported that they are now running Kudu in a production use case on ~70 nodes ingesting 10B rows/day. This is Kudu's first known production use case, and we expect continuing and increasing participation from developers at Xiaomi.
- To spread the word about Kudu, various community members have given a number of talks over the last 6 weeks, including: Strata/Hadoop World (Singapore), Big Data Tech Con (Beijing), Shanghai Streaming Meetup, Korea Big Data Think Tank (Seoul), San Francisco Spark Hackers, The Hive Big Data Think Tank (Palo Alto), Charlotte Hadoop User Group, Cloudera World Tokyo, Big Data Warehousing Meetup (NYC), and more.
- Members of the Kudu community met up with Apache Drill contributors and built an initial version of Kudu/Drill integration. We anticipate further collaboration with Drill and other ASF TLPs.

How has the project developed since the last report?

This is our first report since incubation. In the past month, we've accomplished a few milestones:

- filed an SGA as well as ICLAs for initial committers
- granted karma, created mailing lists, and subscribed mentors/committers
- updated our web site to comply with ASF branding
- filed INFRA requests to start working on git and JIRA migration (INFRA-10959, INFRA-10983)

Progress has slowed in the second half of December due to the holiday season, but we anticipate picking momentum back up in January.

Date of last release:

None under the ASF. Outside of the ASF, we released 0.6.0 on 11/24/2015.

When were the last committers or PMC members elected?

None since initial incubation.

Signed-off-by:

```
[X](kudu) Jake Farrell
[ ](kudu) Brock Noland
[x](kudu) Michael Stack
[ ](kudu) Jarek Jarcec Cecho
[x](kudu) Chris Mattmann
[ ](kudu) Julien Le Dem
[ ](kudu) Carl Steinbach
```

Shepherd/Mentor notes:

log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

1. Activate some community
2. Create a release
- 3.

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

Not sure if the project still has an official mentor assigned, after Christian Grobmeier is not the chair of the Apache Logging Services project anymore.

How has the community developed since the last report?

Not grown. However, the project aims to become a sub project of Apache Logging again. Community growth is not much a matter for this podling, as long as there are a few active contributors.

How has the project developed since the last report?

Not at all: There's no public development and it seems there's currently only one originally incubating community member left. Though some support is given and bugs are tracked.

Date of last release:

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

When were the last committers or PMC members elected?

Signed-off-by:

[](log4cxx2) Christian Grobmeier
[](log4cxx2) Scott Deboy

Shepherd/Mentor notes:

John D. Ament (johndament):

Project appears to be in bad shape. Even though there is activity on the ML there is no mentor engagement.

Marvin Humphrey (marvin):

Christian Grobmeier has resigned as Mentor, and attempts to reach Scott Deboy have not been successful. A thread was started on the general list and there remains some interest in the podling. Bill Rowe has volunteered to serve as Mentor provided that the contributors can rally.

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. Second community call held 12/18/15. On the call Roman suggested that the MADlib community do a 1.9 alpha release in the near term. There was general agreement and this is planned for January. A Release Manager has not yet been identified.
2. Material technical conversations on dev mailing lists and in the appropriate JIRAs. E.g., 51 emails on dev in Dec.
3. One new committer has been proposed and voting is under way on the private mailing list.
4. Two comprehensive proposals were posted to the dev mailing list from community members. One relates to the addition of geographically weighted regression (GWR) algorithms. The second involves Bayesian analysis of binomial response models for MPP databases, which makes extensive use of MADlib's new matrix operations. Both proposals are under active discussion on the mailing list currently.
5. We have been accepted to present a full talk at FOSDEM 2016/Brussels in the HPC, Big Data & Data Science Devroom on Jan 31. The title of the talk is: "MADlib: Distributed In-Database Machine Learning for Fun and Profit"

How has the project developed since the last report?

1. 5 JIRAs created and 4 resolved in last 30 days.
2. A SQL API guide has been added to the the MADlib wiki
<https://cwiki.apache.org/confluence/display/MADLIB/SQL+API+Guide>.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

One new committer has been proposed and voting is under way on the private mailing list.

Signed-off-by:

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

Shepherd/Mentor notes:

Metron

Metron is a project dedicated to providing an extensible and scalable advanced network security analytics tool. It has strong foundations in the Apache Hadoop ecosystem.

Metron has been incubating since 2015-12-06.

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

1. Building a diverse community of developers for Metron
2. Getting security practitioners to provide feedback on requirements
3. Make an 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?

- We added Debo Dutta as a committer

How has the project developed since the last report?

- We received a signed SGA from Cisco
- We picked our new logo
- We setup our new website
- We closed our first few Jiras

Date of last release:

No release yet.

When were the last committers or PMC members elected?

2016-01-08

Signed-off-by:

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

Shepherd/Mentor notes:

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. Do a first Apache release capable of producing a downloadable RTOS image with support for multiple peripherals and network connectivity.
2. 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. Mailing list activity since last report:
@dev 52 messages (1 Dec to 5 Jan), 1 new subscribers (20 total)
@commits 150+ messages
2. Repo is being cloned and looked at by several outside parties - need to get all interested parties on the dev mailing list.
3. Participation on the @dev mailing list covered a range of topics including feature implementation approaches, coding practices, website, and documentation rules discussions.

How has the project developed since the last report?

1. Feature support: 2 new JIRA issues were created to facilitate feature implementation and website discussions
2. Documentation support: Two new committers have started adding material to technical documentation.
3. Effort towards community growth: We have redesigned the project website to organize and present content more clearly and to highlight upcoming events where the project will have exposure.
4. Effort towards self governance: Consensus reached on how a potential

contributor is accepted over time as a committer based on quality of work.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

Signed-off-by:

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

Shepherd/Mentor notes:

Milagro

Distributed Cryptography; M-Pin protocol for Identity and Trust

Milagro has been incubating since 2015-12-21.
In view of the holiday season, it has only really started this week
(Jan 4th) and is currently bootstrapping.

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

1. Import the software.
2. Sign up the developers.
3. Audit the Intellectual Property.

All fairly meaningless until the Community is live at Apache.

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?

n/a.

Initial committers have been asked to sign up to the lists and submit ICLAs.

How has the project developed since the last report?

n/a.

Date of last release:

n/a.

When were the last committers or PMC members elected?

Initial Committers only.

Signed-off-by:

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

Shepherd/Mentor notes:

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

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 documented (on Confluence) and tested procedures for committers to review and incorporate Pull Requests into the incubator-rya Apache git repository.
- * The committers are becoming much more comfortable with the Apache infrastructure and processes. Committers have been documenting processes on Confluence, creating Jira tickets, tying these to Rya components and versions, and submitting and pulling in PRs to resolve these Jira tickets.
- * The Rya user community (i.e. non-committers) have begun submitting issues through Jira and the dev list. These have become Jira tickets, resolved by committers, and tested by the community. Specifically, the community submitted RYA-19, RYA-17, and RYA-6. Committers resolved and the community tested RYA-19 and RYA-17.
- * Non-committers have also begun submitting PRs that have been incorporated into the apache repo. Specifically, RYA-13 and most of RYA-7 was done by non-committers.
- * We've started logging "beginner" Jira tickets. These are issues that we've found in Rya that a new developer could take on as a good way to "get their feet wet" in the code base.

How has the project developed since the last report?

- * We've completed adding the Apache license headers and added license header verification to the project's build.
- * Resolved a handful of issues that users have encountered.
- * Committed features: Delete support for secondary indices, more efficient
- * Mongo Core and Geo indexing.
- * Still working on creating a website for the project.
- * Working on deploying artifacts onto Maven
- * Working on creating a quick-start VM for users to try out Rya.

Date of last release:

Not applicable

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:

Josh Elser (elserj):

Continuing to see growing activity, but not much collaboration between reviews among committers (saw some ignored comments on pull requests). Something to watch in the future.

S2Graph

S2Graph is a distributed and scalable OLTP graph database built on Apache HBase to support fast traversal of extremely large graphs.

S2Graph has been incubating since 2015-11-29.

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

1. Make a release
2. Attract users and contributors
3. Foster more and diverse committers

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?

- * Project infrastructures (i.e., jira, wiki, mailing list and git repository) setup has been completed, and the site and the source code are moved to the Apache infrastructure.
- * All members have submitted necessary CLAs.
- * SGA have submitted from Kakao.
- * There have been no new committers or PPMC members elected.
- * We are discussing the development workflow and the code review policy.

How has the project developed since the last report?

- * We are applying ALv2 header to all files and rat-plugin to audit them.

Date of last release:

No yet

When were the last committers or PMC members elected?

No

Signed-off-by:

[X](s2graph) Andrew Purtell
[x](s2graph) Seetharam Venkatesh
[X](s2graph) Sergio Fernández

Shepherd/Mentor notes:

Seeratham Venkatesh (venkatesh):

The community is bootstrapping and seeing a flurry of activities.

Sirona

Monitoring Solution.

Sirona has been incubating since 2013-10-15.

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

1. increase the community
2. do more releases
3. make the community more active

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?

No change

How has the project developed since the last report?

Several fixes have been done and enhancements in instrumentation side.

Date of last release:

2015-11-18

When were the last committers or PMC members elected?

Signed-off-by:

[x](sirona) Olivier Lamy
[](sirona) Henri Gomez
[X](sirona) Jean-Baptiste Onofre
[](sirona) Tammo van Lessen
[](sirona) Mark Struberg

Shepherd/Mentor notes:

SystemML

SystemML provides declarative large-scale machine learning (ML) that aims at flexible specification of ML algorithms and automatic generation of hybrid runtime plans ranging from single node, in-memory computations, to distributed computations such as Apache Hadoop MapReduce and Apache Spark.

SystemML has been incubating since 2015-11-02.

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

1. Grow SystemML community: increase mailing list activity, increase adoption of SystemML for scalable machine learning, encourage data scientists to adopt DML and PyDML algorithm scripts, respond to user feedback to ensure SystemML meets the requirements of real-world situations, write papers, and present talks about SystemML.
2. Core library improvements, including Apache Spark integration.
3. Produce a release

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

The community has been blocked by INFRA-10714 since November. Beginning Jan 6th, we are manually updating the missing fields in JIRA so that we can properly manage project issues and delegate issues to new users to grow our community.

How has the community developed since the last report?

Users have asked several excellent questions on the dev list, and existing committers are actively helping these users. The project is receiving pull requests from contributors, committers have discussed these pull requests with their contributors, and contributions have been merged into the project. Matthias Boehm has presented talks regarding the SystemML Optimizer at TU Dresden, HTW Dresden, and TU Berlin. Fred Reiss presented a talk on Dec 8th.

How has the project developed since the last report?

Numerous core library improvements have been made to project. Additional documentation has been created to help new users. The test suite has been refactored to increase maintainability and performance. Contributions have been made by non-IBM contributors. We are developing our 2016 roadmap, as seen in the "[DISCUSS] Project Roadmap" thread on the mailing list at <http://s.apache.org/9hK>.

Date of last release:

NONE

When were the last committers or PMC members elected?

NONE

Signed-off-by:

[x](systemml) Luciano Resende
[](systemml) Patrick Wendell
[](systemml) Reynold Xin
[](systemml) Rich Bowen

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.

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

1. Improve documentation on the website
2. Make some incubating releases
3. Grow the community size and diversity.

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 got feedbacks from users and interested people. A good sign is that the 1.0.0-incubating release has been voted by people outside of the initial committer set (even if the release vote didn't pass).

How has the project developed since the last report?

We did improvement on the website, and on the code. Our focus was on Unomi 1.0.0-incubating release. Unfortunately, a couple of issues (missing src assembly especially) blocked the vote. We are fixing those issues to submit a take 2 on 1.0.0-incubating release.

Date of last release:

Unomi 1.0.0-incubating has been proposed to vote, but due to a couple of issue it has been declined. We are working to submit a take 2 on 1.0.0-incubating.

When were the last committers or PMC members elected?

None

Signed-off-by:

☒ (unomi) Jean-Baptiste Onofré

☐ (unomi) Bertrand Delacretaz

☒ (unomi) Roman Shaposhnik

☒ (unomi) Chris Mattmann

Shepherd/Mentor notes: