

April 2016

/\ *FINAL* (!)

Incubator PMC report for April 2016

Timeline

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

Shepherd Assignments

Drew Farris		iota
John Ament		Fineract
John Ament		Joshua
Justin Mclean		HAWQ
P. Taylor Goetz		BatchE E
P. Taylor Goetz		Quarks
P. Taylor Goetz		Rya
Timothy Chen		FreeMarker
Timothy Chen		Tephra
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="638da9d6-af43-4bd4-bbfc-1b2ae88b0665"><ac:plain-text-body><![CDATA[[none]	Apex]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="da56008e-0fdb-4aa5-b7c3-238698dc2681"><ac:plain-text-body><![CDATA[[none]	Concerted]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="27b12f4d-2b9c-400c-9707-446d0ccb99d3"><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="b08bf53a-ab08-4298-ae4a-ace4eb5562f9"><ac:plain-text-body><![CDATA[[none]	Gearpump]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f90c859d-2ef5-4f46-b87f-1d370d0efc9f"><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="823eec4b-e9c0-407c-bde2-8c0ccfc7dc06"><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="13c61231-ee4d-4a39-9cbe-7d4522ab5473"><ac:plain-text-body><![CDATA[[none]	HTrace]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f36009b7-79c7-4725-8e1a-5bebfeabcd7a"><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="16470eba-518f-4f33-8b53-95fc60a9fe78"><ac:plain-text-body><![CDATA[[none]	Johnzon]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="54d621e1-59a0-43b1-ab8f-c1418f7d09b9"><ac:plain-text-body><![CDATA[[none]	MADlib]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="1d7592ba-f167-4aaa-9171-a7b9bc40b071"><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="bfc6cb1a-747b-450c-9a72-09872b35e053"><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="4f0ed9f7-d085-4ce3-b8da-9e8c91372476"><ac:plain-text-body><![CDATA[[none]	Mnemonic]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a363e0cb-103d-46a2-8e24-599a9008b848"><ac:plain-text-body><![CDATA[[none]	Mynewt]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f110c58a-46ad-463e-933a-de41d19e82ad"><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="503e376e-40b7-4513-aedd-f537415e2afa"><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="9b67c950-7de2-487d-a705-2451bf682733"><ac:plain-text-body><![CDATA[[none]	Ranger]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="159ee136-e38b-4be8-b0ff-f771a5337aa6"><ac:plain-text-body><![CDATA[[none]	Sirona]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c63de27d-f1d8-4a4d-87d6-1b890e7c4e0b"><ac:plain-text-body><![CDATA[[none]	Toree]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for April 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 56 podlings currently undergoing incubation. We have four podlings planning to graduate this month and added six podlings to the roster. As is typical in the months following the annual members meeting, we have seen a jump in the number of IPMC members join.

* Community

New IPMC members:

- Marvin Humphrey
- Suneel Marthi
- Uma Gangumalla
- Henri Yandell

People who left the IPMC:

- Jukka Zitting

* New Podlings

- Airflow
- Gearpump
- Mnemonic
- Omid
- Tephra
- Quickstep

* Graduations

The board has motions for the following:

- Apex
- Johnzon
- TinkerPop
- AsterixDB

* Releases

The following releases entered distribution during the month of March:

- 2016-03-02 Apache MRQL 0.9.6-incubating
- 2016-03-03 Apache Apex Malhar 3.3.1-incubating
- 2016-03-03 Apache Htrace 4.1.0-incubating
- 2016-03-04 Apache Mynewt-0.8.0-b1-incubating
- 2016-03-09 Apache Kudu 0.7.1-incubating
- 2016-03-10 Apache Ranger 0.5.2-incubating
- 2016-03-11 Apache Taverna-language 0.15.1-incubating
- 2016-03-11 Apache Taverna-osgi 0.2.1-incubating
- 2016-03-14 Apache MADlib v1.9alpha-rc2-incubating
- 2016-03-21 Apache Unomi 1.0.0-incubating
- 2016-03-25 Apache Mynewt 0.8.0-b2-incubating
- 2016-03-27 Apache FreeMarker 2.3.24-incubating

* IP Clearance

- Apache Sling Dynamic Includes

* Infrastructure

We continue to struggle with podling report reminders.
A suite of manual reminders were sent, some were received
by the podlings and some were not.

* Miscellaneous

Based on group understanding, it appears that the Concerted podling
will retire soon. Same is likely for OpenAz.

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- Gearpump
- iota
- Joshua
- Metron
- Milagro
- Mnemonic
- Quarks

* Not yet ready to graduate

No release:

- Hawq
- Horn
- Impala
- Rya
- Toree

Community growth:

- Fineract
- Geode
- MADlib
- Ranger
- Wave

* Ready to graduate

The Board has motions for the following:

- Apex
- Johnzon
- TinkerPop
- AsterixDB

* Did not report, expected next month

- Concerted
- OpenAz (missed 4 months)
- Tephra

Table of Contents

Apex
BatchEE
Concerted

DataFu
Fineract
FreeMarker
Gearpump
Geode
HAWQ
HORN
HTrace
Impala
iota
Johnzon
Joshua
MADlib
Metron
Milagro
Mnemonic
Mynewt
ODF Toolkit
Quarks
Ranger
Rya
Sirona
Toree

Apex

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

Apex has been incubating since 2015-08-17.

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

None. Incubator graduation VOTE passed: <https://s.apache.org/ghTf>

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

Resolution to establish Apache Apex as TLP submitted to board.

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 (2180 messages on dev@ for March, 122+ subscribers).

The community discussed graduation, addressed outstanding issues and completed VOTE.

The community has been active building additional meetup groups: <http://s.apache.org/jKT>
Many presentations and uptick in engagement.

How has the project developed since the last report?

Release 3.3.1-incubating of Malhar on 2016-03-02

Various metrics for March are as follows:

Metric	Core	Malhar
Non Merge Commits	52	22
Contributors	10	11
Jira New Issues	40	31
Resolved Issues	43	20
Pull Requests merged	33	12
Pull Requests proposed	2	11

We are anticipating patch release 3.2.1 in April.

Date of last release:

2016-03-02 Apex (Malhar) 3.3.1-incubating

When were the last committers or PMC members elected?

2 PPMC members and 1 committer were added since December.
PPMC and committer list was updated on the status page.

Signed-off-by:

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

BatchEE

Apache BatchEE 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:

The PPMC feels that the project is ready to graduate soon.

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?

Bugfixing continues in BatchEE and we are actively working with the JBatch community towards a new JBatch specification

How has the project developed since the last report?

BatchEE contributes to the JBatch specification to make the project moving forward in the right direction.

Date of last release:

2015-12-08

When were the last committers or PMC members elected?

2015-12-11

Signed-off-by:

(batchee) Jean-Baptiste Onofré
 (batchee) Olivier Lamy
 (batchee) Mark Struberg

Shepherd/Mentor notes:

John D. Ament (johndament):

It seems that the community posted one report, but then a mentor posted a largely reworded report. No evidence of a discussion seems to have been done on public mailing lists for either report.

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:

Julian Hyde (jhyde):

No report, despite prodding from Chris, and no activity on lists.
I think this one's dead.

Jake Farrell (jfarrell):

No activity so far in 2016.

Chris Nauroth (cnauroth):

There has been no activity on the project for several months, and we don't see a viable way to salvage it. We are discussing closing it down.

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?

- * A new contributor opened several JIRAs regarding improvements and contributed patches. Two have been committed so far.

How has the project developed since the last report?

- * Improved instructions on loading projects in Eclipse based on discussion in JIRA.
- * Added checks in build system to catch issues using wrong JDK version.
- * Some UDFs were improved to be more efficient.
- * A new UDF is pending review.

Date of last release:

2015-11-14

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

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. Finalising the initial release
2. Improve the communication around the differences between the previous MifosX project and the new Fineract project, especially in the MifosX community, website etc.
3. With the first initial release, encourage the community to use the proper infrastructure (mailing lists, issue tracker) for the ongoing collaboration within the community.
4. Change management towards less 'key-man' dependency on previous tech/community leaders from MifosX, more towards community driven consensus.
5. Reduce clutter on mailinglist from JIRA updates etc.

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

On the mailinglist there have been discussions around the lack of traffic on it. This is part of the switchover from the current MifosX community and mailinglists towards the new Fineract lists. At the same time the recent meet-up of Mifosx and Fineract devs and people interested in it (see community development), has also triggered a lot of offline interaction.

How has the community developed since the last report?

While activity on the mailinglist has been pretty low, we've had a great meet-up in Amsterdam where a nice mix of existing MifosX community members and a group of new interested people was present. This has boosted interest in Fineract.

We have a high level of interest from new contributors throughout Africa - we are working to properly engage them and guide them to the correct collaboration channels in our Fineract community.

How has the project developed since the last report?

Individual members of the leading partner organizations building solutions using the Fineract platform have begun to effectively use the Fineract issue tracker to communicate and track the requirements and enhancements they're building and contributing to Fineract. These individuals are setting a good example that other individuals from our partner community should follow.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

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 more active contributors

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?

FreeMarker had two public releases: A Release Candidate (so that users can test it), and one final release. Apart from the new features and fixes, we have adjusted the source code and build process to follow Apache best practices more closely, and to be more appealing for contributors (switching to Java 5, fixing formatting where it didn't fit the modern Java conventions).

The legal/policy issue with the web site domain (INFRA-10787) was settled, and we are using our historical domain again (freemarker.org) until graduation.

Date of last release:

2016-03-28

When were the last committers or PMC members elected?

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

Signed-off-by:

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

Gearpump

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

Gearpump has been incubating since 2016-03-08.

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

1. Make initial Apache branded release
2. Initial community process definition and practice enforcement by following Apache policies
3. Build the first Apache branded informative website to make Gearpump contributor and end-user friendly.

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

The rights holder of the Gearpump copyright filed a CCLA including a Schedule B granting the Gearpump codebase to the Foundation. We are awaiting assistance from Infrastructure on INFRA-11435 to perform the

import.

How has the community developed since the last report?

All of the initial committers/PPMC have set up with Apache accounts and Apache JIRA accounts.

Discussion among developers has started on dev@gearpump.incubator.apache.org

How has the project developed since the last report?

The JIRA for the podling is active at <https://issues.apache.org/jira/browse/GEARPUMP> and seeing activity.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

No new committers or PMC members elected yet.

Signed-off-by:

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

Geode

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

Geode has been incubating since 2015-04-27.

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

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

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

None

How has the community developed since the last report?

- We have been able to address the "Three most important issues to move towards graduation" from the previous report
 1. We had our first public release
 2. We Added 4 new committers, one of which is outside Pivotal
 3. We are about to have a second release, establishing a regular release cadence
- We had a GeodeSummit <http://geodesummit.com/>
 - 22 sessions (including Keynote)
 - Sessions included current use cases as well as integrations with other Apache projects
 - Over 100 attendees
 - Received project proposal after the summit.
- The Geode Clubhouse hosted a meeting covering a walk through the different approaches for caching data in a Spring based application with Apache Geode by Luke Shannon
- The breakdown of JIRA tickets is the following:
Q3 / 2015 - 65 created, 19 resolved.

Q4 / 2015 - 342 created, 227 resolved.

Q1 / 2016 - 438 created, 362 resolved.

Q2 / 2016 - 8 created, 12 resolved (As of Apr 4th)

- There was a total of 123 pull requests on GitHub with 6 still open.
- The breakdown of the mailing lists messages for January, Feb, March of 2016:

org.apache.geode.issues 4300

org.apache.geode.commits 3731

org.apache.geode.dev 2154

org.apache.geode.user 294

- There are now 153 subscribers on the dev and 145 on the user mailing lists.
- Community events:
 - 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
 - February 2016:
 - Seattle Java User's Group (Talk about Introduction to Apache Geode)
 - March 2016:
 - Geode Summit in Palo Alto, California
 - Geode meetup in Palo Alto (<http://www.meetup.com/Pivotal-Open-Source-Hub/events/228983898/>)
 - Apache Apex (incubating) meetup featuring Geode integration (<http://www.meetup.com/Apex-Bay-Area-Chapter/events/228593080/>)

How has the project developed since the last report?

- Had first release, close to the second release (with 134 out of 138 issues resolved)
- We elected a release manager for the first release and have rotated that responsibility for the second release
- Had a discussion about how to number the releases
- Voted on the Release candidates, rejected the first but the second release candidate was voted as the final release
- The community discussed the criteria for becoming a committer and learned the correct way of nominating new committers.
- JIRA components were refined after a discussion on the mailing list.
- Consensus was achieved for the scope for the releases on the mailing list.
- Decision was made not to rename the command line tool on the mailing list.
- Each commit and pull request is run through CI using Travis CI
- Removed @author tags
- First release published to maven central
- Javadocs hosted on the project website.

Date of last release:

2016-02-06

When were the last committers or PMC members elected?

- Jinmei Liao (2016-01-26)
- Sai Boorlagadda (2016-02-19)
- Nitin Lamba (2016-03-08)
- Dave Barnes (2016-03-30)

Signed-off-by:

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

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. One HAWQ Nest community meeting (topic: GPORCA: Query Optimization as a Service) Video: <http://bit.ly/HAWQNest0223yt>
2. The community shows pretty high activities. In Dev, there have been 816 messages on dev@/user@ in March 2016.
3. Got around 29 pull requests from contributors that are not in the initial committer team.
4. GSoc 2016: 6 topics created for it.
5. FOSDEM 2016 event:
https://fosdem.org/2016/schedule/event/hpc_bigdata_hadoopsql/
<https://s.apache.org/BYMR>

How has the project developed since the last report?

1. Main features/improvements added include:
 - a) HAWQ-473 Add fault reason and history into hawq catalog
 - b) HAWQ-404 Add sort during INSERT of append only row oriented partition tables
 - c) HAWQ-558 Improve performance of truncate on small tables
 - d) HAWQ-532 HAWQ-580 HAWQ-587 Enhance the virtual segment number selection method
 - e) HAWQ-551 Enhance explain analyze with more information on data locality and resource manager
 - f) HAWQ-546 Enhance \d for HCatalog tables
 - g) HAWQ-144 Support HAWQ build on MAC
 - h) HAWQ-597 Add Travis CI(OSX) for MAC
 - i) Critical bug fixes.
2. 135 new JIRAs filed, 116 resolved (In Mar 2016)
3. 134 code commits (In Mar 2016)

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:

[(hawq) Alan Gates

(hawq) Konstantin Boudnik
 (hawq) Justin Erenkrantz
 (hawq) Thejas Nair
 (hawq) Roman Shaposhnik

Shepherd/Mentor notes:

Justin Mclean (jmclean):

A couple on minor issues around communication on list and releases but mentors are on top on it.

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. Increase development velocity
2. Make a 0.1 release
3. Refocus the direction we want to take this 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?

Number of subscribers has increased slightly, development has however been slow.

How has the project developed since the last report?

We have decided that we need to refocus on what we want HORN to look like

Date of last release:

None yet.

When were the last committers or PMC members elected?

October

Signed-off-by:

(horn) Luciano Resende
 (horn) Robin Anil
 (horn) Edward J. Yoon

HTrace

HTrace is a tracing framework intended for use with distributed systems written in java.

HTrace has been incubating since 2014-11-11.

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

1. Continue to grow the HTrace community
2. Continue to explore projects integratiing HTrace
3. Continue to develop and release stable HTrace incubating artifacts

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?

- * YCSB community are working on integrating HTrace.
- * Google Summer of Code 2016 projects for Apache HTrace are being planned. There are proposals such as
 - * distributed tracing for s3 and other alternative Hadoop FS
 - * Kudu span receiver
 - * integrating HTrace into Hadoop YARN
 - * HTrace'ing in Apache Nutch

How has the project developed since the last report?

- * +70 issues were fixed and released as part of 4.1.0-incubating.
- * Improved htraced a lot in Web-UI, robustness, metrics and logging
- * Added Docker support for development
- * Added Documentation for developers

Date of last release:

- * htrace-4.1.0-incubating on March 5th, 2016

When were the last committers or PMC members elected?

- * None this period. Last added committer: Abraham Elmahrek - 2.11.15

Signed-off-by:

[X](htrace) Jake Farrell
[](htrace) Todd Lipcon
[X](htrace) Lewis John Mcgibbney
[](htrace) Andrew Purtell
[X](htrace) Billie Rinaldi
[](htrace) Michael Stack

Shepherd/Mentor notes:

John D. Ament (johndament):

This podling's come pretty far. Maybe time to graduate soon?

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. Transition of development workflows to ASF (see <https://issues.cloudera.org/browse/IMPALA-3221>)
2. Initial release as incubating project.
3. Community growth

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 have been no additions to the committer or PMC lists since incubation began. We continue to see an uptick in external contributions, with two patches from new contributors this month. One contributor pleasingly reported it was "great to work closely with Impala community".

How has the project developed since the last report?

We have put together a list of tasks required to move development of Impala from Cloudera's infrastructure to the ASF. Since Impala was already a well-established project before the Incubator proposal, there is perhaps more decoupling required than for more nascent projects. The list is at <https://issues.cloudera.org/browse/IMPALA-3221>, and is being actively worked on. Note that this doesn't cover the standard podling tasks (like name search, etc).

Date of last release:

None since entering incubation.

When were the last committers or PMC members elected?

None since entering incubation.

Signed-off-by:

(impala) Tom White
 (impala) Todd Lipcon
 (impala) Carl Steinbach
 (impala) Brock Noland

iota

Open source system that enables the orchestration of IoT devices.

iota has been incubating since 2016-01-20.

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

1. Overcome current technical difficulties. You might be asking what is happening with Apache iota. The initial architecture that we wanted to upload was not as scalable as we required it to be. Rather than upload an architecture that we wanted to fundamentally change we have been developing a new backend architecture and code based that implements a significantly more scalable solution using Apache, Spark, Spark Streaming and Cassandra. We are still using Python and Django for our APIs and authentication (session and token based). Consequently, you haven't seen any activity on the Apache iota site and repository. This will be changing soon. Litbit (the contributing company whose business runs on iota) has been working around the clock to remake the backend using additional Apache components that makes the solution far more scalable.
2. Set up web site and import code. We will be setting up the Apache iota website to reflect the new architecture and uploading the initial source code as quickly as we can.
3. Promote Iota to a wider user base / grow the community. We are planning to promoting iota at an upcoming keynote being given by Scott Noteboom (Litbit's CEO). In addition we will be presenting iota at Apache Con in May 2016. Our goal in promoting iota will be to start building a community that will contribute to moving iota to it's first release as an Apache incubator project

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?

* This is the first submitted report

How has the project developed since the last report?

* This is the first submitted report

Date of last release:

* None

When were the last committers or PMC members elected?

None yet, just getting started.

Signed-off-by:

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

Shepherd/Mentor notes:

John D. Ament (johndament):

Its a bit concerning to see that the project is starting off in a non-ASF way. Even if the code isn't quite right for donation, may be better to start off writing new code in the ASF repos.

Justin Mclean (jmclean):

I too have concerns about why they are developing outside of the ASF. I've contacted them a few times with a please explain and have assurances that this will be fixed. I'd suggest placing the podling on monthly reports for a few more months.

Drew Farris (drew):

2 mentors active on the mailing list this month, but little activity as the team has not moved development to ASF. Concur with keeping podling on monthly reports.

Johnzon

Implementation of JSR-353 JavaTM API for JSON Processing.

Johnzon has been incubating since 2014-06-09.

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

1. Currently graduation, so nothing but issues necessary for graduation

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

* Currently graduating to a TLP

How has the community developed since the last report?

There are new people on the dev list and the project got contributions from non-committers. On the dev list there are approx. 20 msgs (on average) messages per month.

How has the project developed since the last report?

The project did its 10th incubator release. During the last three month 9 jira issues were reported and 6 of them fixed. Project has also started to implement the upcoming JSR-374 and JSR-367 specification. The project has also voted for a logo.

Date of last release:

2016-FEB-17 (0.9.3-incubating)

When were the last committers or PMC members elected?

Last committer was elected 2014-08-23, last PMC member was elected 2015-04-09 (Hendrik Saly)

Signed-off-by:

[X](johnzon) Justin Mclean
[](johnzon) Daniel Kulp

Shepherd/Mentor notes:

John D. Ament (johndament):

Podling is planning to graduate, but sign off received from only one mentor.

Joshua

Joshua is a statistical machine translation toolkit

Joshua has been incubating since 2016-02-13.

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

1. Migration of the entire project over to Apache infrastructure (nearly complete)
2. Building up community and attracting developers
3. Putting out regular Joshua releases

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

None that come to mind. We are occupied with the issues above, and things are proceeding nicely.

How has the community developed since the last report?

We have voted in two new committers and invitations have been extended to them.

How has the project developed since the last report?

A large pull request was accepted reporting large speed increases due to caching, and increasing the usability of Joshua as a library.

Date of last release:

2015-11-05 (prior to incubation)

When were the last committers or PMC members elected?

- Felix Hieber (2016-04-01)
- Kellen Sunderland (2016-04-01)

Signed-off-by:

[](joshua) Paul Ramirez
[X](joshua) Lewis John McGibbney
[](joshua) Chris Mattmann
[](joshua) Tom Barber
[](joshua) Henri Yandell

MADlib

Big Data Machine Learning in SQL for Data Scientists.

MADlib has been incubating since 2015-09-15.

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

1. Continue to produce regular Apache (incubating) releases.
2. Expand the community, increase dev list activity and add new

contributors.

3. Execute and manage the project according to governance model required by the "Apache Way".

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

None

How has the community developed since the last report?

1. Held three community calls in Q1 2016. Each call featured a different member of the Apache MADlib community presenting on a topic of interest to them:
 - * Jan 15th - Bayesian Analysis of Binomial Response Models on MPP Databases (Gautam Muralidhar)
 - * Feb 16th - An Overview of GWR Analysis of Spatial Data (Chenliang Wang)
 - * Mar 16th - MADlib on PostgreSQL and PGXN (AJ Welch)
2. One new committer has been added to the project (Xiaocheng Tang)
3. Presentation of Apache MADlib at FOSDEM'16 in Brussels (Frank McQuillan)
4. Material technical conversations on dev mailing lists and in the appropriate JIRAs (e.g., 111 emails on dev@ mailing list in Feb)
5. Several Google Summer of Code (GSoC) candidates have expressed interest in working on MADlib projects via dev@ mailing list.

How has the project developed since the last report?

1. 1st ASF release MADlib v1.9alpha on 3/11/16 which was intended to clear all potential IP issues in the code base and make it legally ready to be adopted by the community.
2. 2nd ASF release MADlib v1.9 is currently in IPMC voting as of this writing on 4/6/16. The goal of this 2nd release is general availability of MADlib v1.9 for community use.
3. Some features in the latest release: path functions, support vector machines, advanced matrix operations, covariance matrix, proportion of variance for PCA, support for Apache HAWQ (incubating) 2.0.
4. 15 JIRAs created and 27 resolved in last 30 days.

Date of last release:

Apache MADlib (incubating) v1.9alpha on 3/11/16.

When were the last committers or PMC members elected?

Xiaocheng Tang on 1/14/16.

Signed-off-by:

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

Metron

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

Metron has been incubating since 2015-12-06.

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

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

- Make an Apache release

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

- We are currently trying to produce the first Apache release of Metron

* How has the community developed since the last report

- We added one additional committer, David Lyle
- We accepted 2 separate contributions from non-committers interested in Metron
- We saw an uptick in interest from the community on the dev boards which corresponds to people trying the product prior to an official release, which is a good sign.

* How has the project developed since the last report

- We closed 20 more Jiras and are in the latest stages of releasing an Apache release.
- We added deployment automation for EC2 so trying out Metron in a multi-node environment is even easier
- We integrated the CI environment that infra set up for us at our request so that unit tests and apache rat license checks are run on every pull request submitted
- We added support for indexing enriched sensor capture data into Apache Solr indexing as well as Elasticsearch

Date of last release:

--

When were the last committers or PMC members elected?

David Lyle on 22 March 2016

Signed-off-by:

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

Milagro

Distributed Cryptography;
M-Pin and other pairing protocols for Identity and Trust

Milagro has been incubating since 2015-12-21.

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

1. Import the software: The source code is already prepared and placed in repositories on GitHub - soon to be imported to Apache repositories. This includes a full Milagro MFA implementation.
2. Awareness among developers: Milagro announcement - during ApacheCon North America - May 2016
3. General Awareness

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?

No substantial developments to report. The project is targeting community development following Apachecon presentations.

How has the project developed since the last report?

Project members are preparing to present the project at ApacheCon with a view to connecting with the wider Apache community and finding synergies with other projects.

Date of last release:

n/a

When were the last committers or PMC members elected?

February, as reported.

Signed-off-by:

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

Mnemonic

Mnemonic is a Java based non-volatile memory library for in-place structured data processing and computing.

Mnemonic has been incubating since 2016-03-03.
Mnemonic code was moved into ASF on 3/24

Various metrics for March are as follows:

```
+-----+
| Metric                | counts |
+-----+
| Non Merge Commits    | 82     |
| Contributors         | 10     |
| Jira New Issues      | 13     |
| Resolved Issues      | 4      |
| Pull Requests merged | 8      |
| Pull Requests proposed | 8     |
+-----+
```

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

1. Apache releases: we are preparing 0.1.0 release
2. Expand and build strong community, keep active dev list discussions, adding new contributors
3. Develop new features

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

Everything is on track, nothing urgent at this time.

How has the community developed since the last report?

The project is in early stage in setting up the infra. In addition to contributors listed in Mnemonic proposal, we had one new contributor joined us. We are actively discuss with Arrow community in dev list for the idea on integration. Apache Arrow community is very interested in Mnemonic integration for extending Arrow into persistent memory and storage.

How has the project developed since the last report?

We ported source code in Apache on 3/24 after SGA was filed on 3/23. We are doing several refactorings and defining coding guidelines to fit Apache project standard.

We are working on enabling Apache Arrow to use Mnemonic as underlying

pooling mechanism.

We are designing the native computing infrastructure, built native computing interface code.

Bug tracking JIRA setup was done, 13 issues were filed. 4 issues were solved.

Mailing list discussions are happening actively related to infrastructure, code standard, features for the project.

Date of last release:

No release yet, working on 0.1.0 release

When were the last committers or PMC members elected?

No election, since the project is just starting. As part of incubation process, initial committers and contributors were proposed and added.

Signed-off-by:

[X](mnemonic) Patrick Hunt
[X](mnemonic) Andrew Purtell
[X](mnemonic) James Taylor
[X](mnemonic) Henry Saputra

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 releases capable of producing a downloadable RTOS image with support for multiple peripherals and network connectivity. Do subsequent releases to indicate thorough understanding, repeatability, and maturity of process.
2. Continue to develop and execute policies that enable project contributors to achieve self-governance
3. Expand community - attract new project contributors and users, grow committer base

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

None

How has the community developed since the last report?

1. Healthy, active mailing lists with increasing numbers of subscribers
2. Continual outside interest by 3rd parties. Outreach via conferences, exhibits, one on one meetings, beta testers, GSoC participation.
3. Vigorous discussions and implementation decisions via proposals and voting on @dev mailing list.

How has the project developed since the last report?

1. Releases: Two releases so far - first on 22nd Feb, next on 18th March
 - Identified and fixed license issues, regular process established to check for license issues esp. with third party software
 - Tagging process and release candidate naming process tested and used
 - Voting process and timelines established and used
2. Documentation support: Several committers added documentation either through direct git pushes or pull requests on github mirror.
3. Effort towards self governance:
 - Two releases completed successfully.

- Two committer candidates going through voting for committer status after meeting threshold of patch submissions.

Date of last release:

2016-03-18

When were the last committers or PMC members elected?

At the start

Signed-off-by:

- [](mynewt) Sterling Hughes
- [X](mynewt) Jim Jagielski
- [X](mynewt) Justin Mclean
- [](mynewt) Greg Stein
- [](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?

Not many developers are active currently.

How has the community developed since the last report?

Several patches have been added to JIRA by community members. Issues have been reported.

How has the project developed since the last report?

Several patches have been applied, a new release is planned.

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 new release will generate a little more interest. I'd suggest we wait a few more months to see the effect of that, before deciding on what to do with the podling (staying in the incubator for ever isn't an option!)

Quarks

Apache Quarks is a programming model and micro-kernel style runtime that can be embedded in gateways and small footprint edge devices enabling local, real-time, analytics on the continuous streams of data coming from equipment, vehicles, systems, appliances, devices and sensors of all kinds (for example, Raspberry Pis or smart phones). Working in conjunction with centralized analytic systems, Apache Quarks provides efficient and timely analytics across the whole IoT ecosystem: from the center to the edge.

Quarks has been incubating since 2016-02-29.

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

1. Create and expand a diverse community of contributors and committers around Quarks project
2. Create the first Apache release of Quarks
3. Complete migration of website to host all content at Apache (javadoc is still on Github).

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

Nothing jumps to mind.

How has the community developed since the last report?

The Quarks community is keeping dev-related conversations on the mailing list. Because of this, it's allowed for the participation of independent contributors in mailing list discussions. Although our list of independent contributors is small, Quarks is still very new as an Apache project and we aim to actively grow our list in the near future.

How has the project developed since the last report?

- * The project has successfully migrated from Github-based development to using the Apache git repo, Jira, and the mailing lists.

Date of last release:

Quarks is still new, so we haven't created an Apache release as of yet.

When were the last committers or PMC members elected?

No changes from the initial committers.

Signed-off-by:

[](quarks) Daniel Debrunner
[](quarks) Luciano Resende
[X](quarks) Katherine Marsden
[X](quarks) Justin Mclean

Ranger

The Ranger project is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform.

Ranger has been incubating since 2014-07-24.

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

1. Increase the community participants for Apache Ranger
2. Re-organize documentation for easier access to end-users, new users and contributors.
3. Integration of Apache Ranger with other Apache Security Initiatives

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. Added new contributors:
(Bolke de Bruin,Deepak Sharma,Colm O hEigeartaigh,Tushar
Dudhatra,mugdha.varadkar,pradeep.agrawal,mehul,ankita.sinha)
Total Contributor Count: 23
2. Mailing list activity since last report (Dec-01-2015 to Mar-31-2016):
@dev 1534
@user 153
@commit 345
3. Issues (JIRAs) created/resolved since last report (after Dec-01-2015
before Mar-31-2016):
Created: 153
Resolved: 140

How has the project developed since the last report?

1. Released two minor releases :
0.5.1 release on 01/26/2016
0.5.2 release on 02/29/2016
2. Members have been working on a next major release of 0.6.0
 - a. Working with new contributors to add HSM support for Apache Ranger KMS.
 - b. Working to provide both Column-level-masking and Row-level-filtering support for data security
 - c. Working to add Tag-Based authorization policies

Date of last release:

Feb-29-2016

When were the last committers or PMC members elected?

1. Gautam Borad has been added as committer on June-29-2015

Signed-off-by:

[](ranger) Alan Gates
[](ranger) Daniel Gruno
[](ranger) Devaraj Das
[X](ranger) Jakob Homan
[X](ranger) Owen O'Malley

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. Become familiar with the release process and have a first release as part of the Apache Foundation
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. Add new committers to the project.

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

No

How has the community developed since the last report?

- * We have a Rya "office hour" teleconference every other week. We had some technical difficulties with screen sharing, but we had a working solution for the last meeting. Users and developers can ask questions, receive answers, discuss ideas for future developments. The minutes are sent to the dev@ list.
- * We have more PRs from non-committers which are integrated into the repository. We expect some of these non-committers to continue to submit PRs and become committers.
- * There were several ideas from Rya for Google Summer of Code, and a few students expressed interest; one of them submitted a proposal

How has the project developed since the last report?

- * Created a quick-start VM for users to try out Rya.
- * Created a website for the project.
- * Resolved a handful of issues that users have encountered.
- * Committed features: incremental update for precomputed join indices, extensions to query planning to support more complex precomputed join indices, free text search support for MongoDB, additional Geo indexing for MongoDB, extensions of examples for both Accumulo and MongoDB backed 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
[x](rya) Edward J. Yoon
[x](rya) Sean Busbey
[x](rya) Venkatesh Seetharam

Shepherd/Mentor notes:

Sean Busbey (busbey):

Project activity was a little slow, due to other time commitments for the dev community, but folks had a great discussion around next steps when it came time for a draft report.

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 but some more user feedbacks.

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
[](sirona) Jean-Baptiste Onofre
[](sirona) Tammo van Lessen
[](sirona) Mark Struberg
[X](sirona) Romain Manni-Bucau

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: This is in progress, and hopefully we will have the dependency license changed to a license that is compatible with AL2.
2. Make a release: nearing completion
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)

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

TOREE-262 - Progress on removal of LGPL dependency

How has the community developed since the last report?

1. Active communication in mailing list and gitter with early adopters
2. Still working on transitioning users from Spark Kernel project into Toree.
3. First external contribution made by user @dongjoon-hyun
4. Apache Toree was featured at SF Data Science meetup
<https://github.com/asimjalis/apache-toree-quickstart>

Date of last release:

None since incubation.

When were the last committers or PMC members elected?

No new additions since incubation

Signed-off-by:

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

Wave

A wave is a hosted, live, concurrent data structure for rich communication. It can be used like email, chat, or a document.

Wave has been incubating since 2010-12-04.

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

1. Growing community
2. Improving code 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?

- We made our first release.
- Elected new committer.
- Previously private fork (wiab.pro) opened the source code that includes many improvements that can be integrated back into main project.
- Also, there's an active development in additional open source fork (swellrt) that potentially allows us to integrate back important features.

How has the project developed since the last report?

- Replaced the old ANT build system with a new Gradle one.
- Fixed a few bugs
- Made a few code improvements

Date of last release:

March 2016

When were the last committers or PMC members elected?

February 2016

Signed-off-by:

[x](wave) Upayavira