

June2016

/\ *FINAL* (!)

Incubator PMC report for June 2016

Timeline

Wed June 01	Podling reports due by end of day
Sun June 05	Shepherd reviews due by end of day
Sun June 05	Summary due by end of day
Tue June 07	Mentor signoff due by end of day
Wed June 08	Report submitted to Board
Wed June 15	Board meeting

Shepherd Assignments

Drew Farris	Gearpump		
Drew Farris	Taverna		
John Ament	Quarks		
John Ament	Tephra		
Justin Mclean	Gossip		
P. Taylor Goetz	SAMOA		
P. Taylor Goetz	Streams		
Timothy Chen	Singa		
Timothy Chen	iota		
Timothy Chen	log4cxx2		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7dc2fe62-7971-4811-aa85-f08b94681850"><ac:plain-text-body><![CDATA[[none]	Airflow]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="34e953ec-68e6-431f-a91b-f30f65651b34"><ac:plain-text-body><![CDATA[[none]	Atlas]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c76a811e-7d26-403a-93ea-f896c61a20c2"><ac:plain-text-body><![CDATA[[none]	Commons RDF]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="01c9de24-0fc3-4527-be63-e457009e4077"><ac:plain-text-body><![CDATA[[none]	Fluo]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d2f01169-b313-40eb-8349-b3b5926cfe30"><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="70c87e79-141a-4b05-8084-8cf3e8e4423b"><ac:plain-text-body><![CDATA[[none]	MRQL]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="1eb3a968-31bd-4e07-b317-6629384d4509"><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="ba5075a6-9cef-4f84-bf06-64d7c0eb6039"><ac:plain-text-body><![CDATA[[none]	Myriad]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7b613ba3-6637-49ab-a13a-3cab4cbbaf8b"><ac:plain-text-body><![CDATA[[none]	Omid]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3f40cc62-66db-4cfd-b3a9-228a9994bf3a"><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="190c0369-cf75-41bf-a3d6-5be7a3d27643"><ac:plain-text-body><![CDATA[[none]	Quickstep]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d0649bb2-4287-4406-ab08-ee8f9e326fbc"><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="b8d30615-ef22-4bf4-a11c-d52b86ecba2d"><ac:plain-text-body><![CDATA[[none]	Trafodion]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for June 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 one podling planning to graduate this month and added two podlings to the roster. May was a generally quiet month, adding a single IPMC member, and completing four releases.

* Community

New IPMC members:

- Joe Witt

* New Podlings

- Pony Mail
- Fluo

* Graduations

The board has motions for the following:

- Twill

* Releases

The following releases entered distribution during the month of May:

- 2016-05-05 Apache Mynewt 0.8.0
- 2016-05-24 Apache Tephra 0.8.0
- 2016-05-24 Apache CommonsRDF 0.2.0
- 2016-05-24 Apache Mnemonic 0.1.2

* Legal / Trademarks

- A comment on a recent Podling Name Search indicated that there was a desire to have PNS's come later on during incubation. Many IPMC members seem to believe that they should happen earlier. An email requesting clarification has been sent.

* Infrastructure

- Multiple former podlings are reporting delays in being converted to TLPs.

* Miscellaneous

- The CMDA podling is now 2 months behind. In response to the lack of activity, a mentor has begun retirement discussions with the podling.

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- Fluo
- Gossip
- Pony Mail

* Not yet ready to graduate

No release:

- Airflow
- Gearpump
- iota
- log4cxx2
- Quarks
- Quickstep
- Streams
- Toree

Community growth:

- Atlas
- CommonsRDF
- HTrace
- Mnemonic
- MRQL
- Myriad
- Omid
- OpenAz
- Ranger
- SAMOA
- Singa
- Taverna
- Tephra
- Trafodion
- Wave

* Ready to graduate

The Board has motions for the following:

- Twill

* Did not report, expected next month

- Climate Model Diagnostic Analyzer (2 months)

Table of Contents

Airflow
Atlas
CommonsRDF
Fluo
Gearpump
Gossip
HTrace
iota
log4cxx2
Mnemonic
MRQL
Myriad
Omid
OpenAz
Pony Mail
Quarks
Quickstep
Ranger
SAMOA
Singa
Streams
Taverna
Tephra
Toree
Trafodion
Wave

Airflow

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

Airflow has been incubating since 2016-03-31.

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

1. Releases
2. Grow up user and contributor communities
3. Improve and extend documentation and samples on the website

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?

- * Since our last podling report 1 month ago, we grew our contributors from 137 to 148
- * Since our last podling report 1 month ago, we accepted/merged 51 PRs
- * We voted on the following matters according to Apache guidelines:
 - * We voted to make all current and future committers part of the PPMC
 - * We voted in a new committer and PPMC member : Steven Yvinec-Kruek and he accepted
 - * We voted in a commit policy of "RTC with a +1 vote from a committer other than the author (assuming no vetos)"

How has the project developed since the last report?

- * All resources have been created and all migrations are complete - code donation included
- * The website is in progress

Date of last release:

N/A - we are still working out releases as we are new to Apache.

When were the last committers or PMC members elected?

On May 20, 1 committer/PPMC member(Steven Yvinec-Kruek) was elected and he accepted.

Signed-off-by:

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

Atlas

Apache Atlas is a scalable and extensible set of core foundational governance services that enables enterprises to effectively and efficiently meet their compliance requirements within Hadoop and allows integration with the complete enterprise data ecosystem

Atlas has been incubating since 2015-05-05.

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

1. Podling name search is pending.
2. Expand the community and add more committers.
3. The upcoming release for Atlas (0.7-incubating) will be nearly 6 months since the last one. We would like to improve on this by making more

frequent releases, and also encourage the community to discuss in release roadmaps etc.

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

No specific issues at this time to report.

How has the community developed since the last report?

1. The user/dev community show increasing interests. In the last two months, the average number of messages has increased to more than 900 messages per month, from an average of 400 per month in the last report.
2. Around 7 contributors were added, number of contributors is at 49 currently. Community is more widespread with more organizations participating in the project.

How has the project developed since the last report?

1. We are working towards a 0.7-incubating release around the next month. This release will include several new features to make Atlas enterprise ready including: Authentication and authorization support, High availability, deeper integrations with other Hadoop components including Hive, HDFS, Falcon etc., Entity versioning, support for rolling upgrade, support for creating business taxonomies, enhanced user interface, better performance etc.
2. There are more strategic contributions coming in to make the system agnostic to the underlying graph database version, and improve horizontal scalability.
3. A total of 216 issues were reported between April 1st 2016 and May 30th 2016. 84 have been resolved in the same time.

Date of last release:

2015-12-31 - 0.6-incubating

When were the last committers or PMC members elected?

1. Suma Shivaprasad was added as a PPMC member on 2016-05-03.
2. Hemanth Yamijala was added as a committer on 2016-04-14 and PPMC member on 2016-04-29.

Signed-off-by:

[](atlas) Arun Murthy
[x](atlas) Chris Douglas
[x](atlas) Jakob Homan
[](atlas) Vinod Kumar Vavilapalli

CommonsRDF

Commons RDF is a set of interfaces and classes for RDF 1.1 concepts and behaviours. The commons-rdf-api module defines interfaces and testing harness. The commons-rdf-simple module provides a basic reference implementation to exercise the test harness and clarify API contracts.

CommonsRDF has been incubating since 2015-03-06.

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

1. Find concrete target goals that could make the project viable as a Commons component.
2. Work with key target communities (Apache Jena and Eclipse RDF4J) to ensure Commons RDF meets potential user requirements and expectations towards next versions of the Commons API.
3. Re-engage and grow the Commons RDF community base.

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

More key people (Andy Seaborne this time) has left the podling due lack of time. Mentor Rob Vesse has also resigned for the same reason.

How has the community developed since the last report?

Community have been compacted with the lost of some key people. We still aim to attract a more broader RDF community, both in and out of ASF.

Mailing list stats:

dev@commonsrdf traffic
Mar 2016: 45
Apr 2016: 75
May 2016: 18

Jira stats for last 90 days:
Created: 7
Resolved: 6

How has the project developed since the last report?

Project's development has resumed the work in the last weeks, being able to cast a 0.2.0-incubating release <https://s.apache.org/rd0E>

Achieving such milestone has reactivated the discussion about the future of the project. All options have been put on the table, even retiring the podling; but currently the most likely option looks to be graduating as a Apache Commons component.

Rather than the original plan of aiming for each RDF frameworks to adapt Commons RDF, the project is now experimenting with its own integration bindings in the style of Commons VFS. Discussions will continue in the next weeks.

Date of last release:

2016-05-24 apache-commonsrdf-0.2.0-incubating

When were the last committers or PMC members elected?

None

Signed-off-by:

[X](commonsrdf) John D Ament
[X](commonsrdf) Gary Gregory

Fluo

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

Fluo has been incubating since 2016-05-17.

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

1. Attract new contributors.
2. Establish an Apache website for Fluo that makes it easy for new users and contributors to get started.
3. Do a release. We should be able to start working on our first release as soon as we are fully established on Apache infrastructure

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?

The community is the same as when it started incubation a few weeks ago.

How has the project developed since the last report?

This is the first Fluo report. We are trying to get established within Apache infrastructure. Mailing lists are setup. Need to import code. We are working with INFRA to explore using Github for issue tracking. See INFRA-11900.

Date of last release:

Never

When were the last committers or PMC members elected?

Never

Signed-off-by:

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

Shepherd/Mentor notes:

John Elser:

Not much to report here. Podling members are working with INFRA to get the existing codebase+issue tracker content ported. Still spinning up.

Gearpump

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

Gearpump has been incubating since 2016-03-08.

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

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

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

None

How has the community developed since the last report?

- Good discussion on dev mailing list related to features, bugs and logo
- Increased participation and interest from developers to provide new sources and examples

How has the project developed since the last report?

- Package renaming completed.
- New logos submitted - reasonable number to finally vote on.
- Continued integration within Apache Beam (Gearpump Runner). PR submitted.
- Continued integration within akka-streams (Gearpump Materializer). PR submitted.
- Release pending this month. Over a dozen JIRA's fixed.

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

Shepherd/Mentor notes:

Drew Farris (shepherd):
One mentor active on the user/dev lists. Healthy activity on mailing lists
commits, moving toward first apache release.

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

Gossip has been incubating since 2016-05.

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

1. Establish an Apache website for Gossip that makes it easy for new users and contributors to get started.
2. Facilitate discussions and build the protocol specification and implementation
3. Produce a usable release

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

No

How has the community developed since the last report?

The community is the same as when it started incubation a few weeks ago.

How has the project developed since the last report?

This is our first report. We have imported the code from github. We are online. We have had good discussions on the dev list that have spawned a couple of tickets. One is in the review stage now. Taylor has been an enormous help.

Date of last release:

Never

When were the last committers or PMC members elected?

Never

Signed-off-by:

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

Shepherd/Mentor notes:

Sean Busbey:
Community def still bootstrapping, but has done a good job of making progress once infra was in place.

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 integrating 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?

There are no issues requiring board attention at this time.

How has the community developed since the last report?

- * GSoC student, Nisala Mendis, mentored by Colin McCabe working on making an Apache Kudu (incubating) SpanReceiver

How has the project developed since the last report?

- * 18 issues resolved since last report
- * YCSB integrated HTrace

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
[](htrace) Lewis John Mcgibbney
[](htrace) Andrew Purtell
[X](htrace) Billie Rinaldi
[x](htrace) Michael Stack

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. Complete code upload
2. Continue discussion on requirements started by Anatole Tresch
3. Need to add additional committer (Barbara Gomes)

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?

Presentation at ApacheCon 2016 - good suggestions from Anatole Tresch in moving iota forward. Have additional potential committer

How has the project developed since the last report?

Code is now being uploaded as well as requirements documents

Date of last release:

None

When were the last committers or PMC members elected?

None

Signed-off-by:

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

Shepherd/Mentor notes:

Justin Mclean (jmclean):
Slow start and still low activity on mailing list.
Web site now has content and there is some code in git.
IMO keep on monthly reports until that changes.

log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

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

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

There was a discussion in the first quarter of the year to move the project to attic unless it is able to release in the next few months. No release happened since then.

The project would like to have some kind of maintenance mode within Apache, where things like community growth and releases are not that important, as long there's a bit of support provided and "anyone" is still alive. The problem with attic is the read-only state of the repo, which makes it impossible to do anything (even if it's rare).

How has the community developed since the last report?

Not grown and only a few mails on the dev list in the last two months.
No support requests the last two months.

Bill Rowe stepped in as a mentor.

How has the project developed since the last report?

The build process using autotools on Linux has been fixed and there was a discussion about supported build tools, because currently the project heavily relies on the unsupported cpptasks framework. While there was no actual result, the common focus was on dropping cpptasks and ANT for the build of the software and use autoconf and/or cmake instead.

Additionally, one project member was willing to make himself more familiar with the release process, but there was no progress reported in the last month.

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:

[X](log4cxx2) Ralph Goers

Shepherd/Mentor notes:

I have explained that there is no "maintenance mode" for a standalone project. However, if the project can demonstrate that it can produce a release and has enough people to maintain the project then it can rejoin the logging project where it can essentially exist in maintenance mode much like log4php is.

The project moved to the incubator because there were people who said they wanted to get involved with the project and were willing to get involved. To date, one person has been actively involved but not enough to produce a release. One other person was involve in a discussion on the release process.

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.

On 05/24/2016, we released our first official build at <https://dist.apache.org/repos/dist/release/incubator/mnemonic/0.1.2-incubating/>

Various metrics for May are as follows: (5/3-5/31/2016)

+-----+		
Metric	counts	
+-----+		
Non Merge Commits	23	
Contributors	3	
Jira New Issues	18	
Resolved Issues	18	
Pull Requests merged	0	
Pull Requests proposed	0	
+-----+		

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

1. Apache releases: Community voted and released Mnemonic-0.1.2-incubating
2. Expand and build strong community, keep active dev list discussions
3. Develop new features: We are working in Native computing layer design

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

There are no issues requiring board attention at this time

How has the community developed since the last report?

Yanping and Gary presented Mnemonic project at ApacheCon Big Data Conference. Yanping presented at In-Memory Computing Summit. Both talks were well received. Several developers approached us to ask additional details on how Mnemonic worked for Spark. They are interested to see if we are able to successfully integrate Mnemonic into other projects.

While we are focusing on develop new features, we will try to contact other projects' communities to explore integration opportunities. As the official release is done this month, for the next 3 months we expect integration activities will be increased.

How has the project developed since the last report?

We have been focusing the first Apache release over the past weeks.

We also explored ideas on design native computing layer which will further enhance durable Java object model on large non-volatile memory.

Date of last release:

2016-05-24

When were the last committers or PMC members elected?

No election, the project is still in early stage of developing.

Signed-off-by:

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

MRQL

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

MRQL has been incubating since 2013-03-13.

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

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

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

None

How has the community developed since the last report?

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

How has the project developed since the last report?

We had a new release on March 2nd. It included various bug fixes and performance improvements. The most important new feature in the new release is support for incremental query processing in MRQL streaming, called Incremental MRQL. There was very little activity on JIRA after the release (only 4 JIRA issues were reported and fixed).

Date of last release:

2016-03-02

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

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

Myriad

Myriad enables co-existence of Apache Hadoop YARN and Apache Mesos together on the same cluster and allows dynamic resource allocations across both Hadoop and other applications running on the same physical data center

infrastructure.

Myriad has been incubating since 2015-03-01.

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

1. Develop project roadmap for longer term community/user engagement.
2. Release frequently - 0.2.0 is underway, but has taken ~6 months since last release.
3. Expand community - users/contributors/committers.

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?

- * dev@ mailing list experienced a low in March, but picked up traffic leading up to 0.2.0 release. 141 messages since the last report.
- * 5 new members on the dev@ mailing list. 2 new contributors.
- * Myriad was presented at ApacheCon Vancouver and at couple of other meetups. Talks submitted at various conferences.
- * Bi-weekly dev syncs happening steadily. Approx. 4-7 members participate. Minutes at <http://s.apache.org/8kF>

How has the project developed since the last report?

- * Myriad 0.2.0 out for PPMC voting. DarinJ is driving the release.
- * 12 commits since 4/1.
- * 13 JIRAs fixed/resolved.

Date of last release:

2015-12-09 myriad-0.1.0-incubating released

When were the last committers or PMC members elected?

2015-10-05 Darin J
2015-10-14 Swapnil Daingade

Signed-off-by:

[](myriad) Benjamin Hindman
[](myriad) Danese Cooper
[X](myriad) Ted Dunning
[x](myriad) Luciano Resende

Shepherd/Mentor notes:

Ted Dunning:
Seems kind of low activity, but with dedicated cadre.

Omid

Omid is a flexible, reliable, high performant and scalable ACID transactional framework that allows client applications to execute transactions on top of MVCC key/value-based NoSQL datastores (currently Apache HBase) providing Snapshot Isolation guarantees on the accessed data.

Omid has been incubating since 2016-03-28.

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

1. Code and documentation successfully moved into the Apache infrastructure (See <https://git-wip-us.apache.org/repos/asf?p=incubator-omid.git> & <http://omid.incubator.apache.org/>)
2. Prepared release guide (<https://cwiki.apache.org/confluence/display/OMID/Omid+Release+process>),

first release (0.8.2) under Apache
(<https://cwiki.apache.org/confluence/display/OMID/Omid+Release+0.8.2.0>)
and upload binaries to Maven Central
3. Active collaboration started with another Apache community: Apache Hive

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?

- * Talk submitted to Hadoop Summit. It was accepted and will be presented on 2016-06-29
- * The Hive community is using Omid to integrate with HBase for their metastore requirements
- * First contact with the Apache Phoenix community for a possible integration of Omid as a transaction manager for the SQL layer Phoenix provides
- * Quarter Stats (from: 2016-05-17 -project moved into Apache- to: 2016-06-03):

Metric	counts
# of msgs in dev list	303
Active Contributors (incl mentors)	6
Jira New Issues	27
Resolved Issues	31
Pull Requests merged	9
Pull Requests proposed	12

How has the project developed since the last report?

N/A, this is the first report since joining the Apache Incubation program

Date of last release:

2016-06-06

When were the last committers or PMC members elected?

At incubation acceptance (2016-03-28)

Signed-off-by:

[x](omid) Daniel Dai (Champion)
[X](omid) Alan Gates
[](omid) Lars Hofhansl
[X](omid) Flavio P. Junqueira
[](omid) Thejas Nair
[](omid) James Taylor

OpenAz

OpenAz

OpenAz is a project to create tools and libraries to enable the development of Attribute-based Access Control (ABAC) Systems in a variety of languages. In general the work is at least consistent with or actually conforming to the OASIS XACML Standard.

OpenAZ has been incubating since 01-2015.

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

1. Increasing community involvement: We are focusing on getting information published on the website to aid the acquisition and orientation of new contributors, as well increasing the number of

committers, in order to reduce individual member workload, reduce reliance on any particular member, and to increase project velocity. David Ash was voted in as a committer last month. There are several other committers that may be added in the near future.

2. Merging the source code from AT&T, JP Morgan and Oracle and creating POMs. AT&T and JP Morgan source code has been merged and POMs created, it needs uploading into the Apache repository. AT&T and JP Morgan ICLAs were received, but Oracle is taking more time. The IP clearance is on it's way.
3. Get our first release done.

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

None. The efforts to revive the project seem to be paying off and discussions are gaining traction.

How has the community developed since the last report?

Pam Dragosh, David Ash and Ajith Nair are the active committers at this point. The project intends to reboot the PPMC which may include removing some of the initial committers who are no longer active. There are, however, several contributors who could be potentially added as committers in the future.

How has the project developed since the last report?

1. Ongoing collective effort to improve project documentation, marketing material for better visibility and adoption.
2. A pending contribution towards PEP implementation is being reviewed and we plan to merge in these changes soon.
3. We are looking into design feedback/comments from Nextlabs - a XACML service provider. We hope to incorporate many of these in the coming weeks.

Date of last release:

N/A

When were the last committers or PMC members elected?

One new committer, David Ash, was elected Tuesday, May 3rd.

Signed-off-by:

[x](openaz) Emmanuel Lecharny
[x](openaz) Colm O Heigeartaigh
[](openaz) Hadrian Zbarcea

Shepherd/Mentor notes:

Emmanuel Lecharny:
This is encouraging ! I'm very pleased with what happened in the past last weeks, with some new committer voted in, and some increasing activity going on.

Pony Mail

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

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

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

1. Grow the community
2. Teach newcomers the Apache Way
3. Build towards the first release

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

None at present.

How has the community developed since the last report?

Just getting started with bootstrapping the podling. Mailing lists have been set up, temporary web site and code repositories ditto. We are working on onboarding the new non-Apache committers in the community and having people familiarize themselves with the project.

How has the project developed since the last report?

N/A

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

No new committers have been elected just yet, as we just got started.

Signed-off-by:

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

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 the product name search and select a suitable name

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

Due to trademark concerns with "Quarks", the Quarks PPMC voted to change the project's name. A renaming discussion and public vote was held on the dev mailing list. The voting results are being discussed privately by the PPMC.

How has the community developed since the last report?

We have added two new committers and PPMC members, Kathy Saunders and Queenie Ma. Kathy has taken the lead on the name search process for Quarks, and her positive attitude and helpfulness are an incredible asset for the Quarks community. Queenie has contributed content to the website and is a frequent participant in community discussions. The Quarks community is still small, as Quarks has just recently incubated, therefore we aim to continue growing our list in the near future.

How has the project developed since the last report?

* According to JIRA, the project has added 27 new issues in May, and 22 issues have been resolved.

* On May 12, several Quarks contributors hosted an "Apache Quarks on Raspberry Pi" webcast where they showed how to use Quarks to work with real sensor data on a Raspberry Pi, and demonstrated a smart sprinkler application using Quarks.

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?

In May, we added two new committers and PPMC members, Kathy Saunders and Queenie Ma.

Signed-off-by:

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

Quickstep

Quickstep is a high-performance data processing platform that aims to fully exploit the vast and increasing amount of parallelism that is packed in individual nodes. The project's initial target is interactive in-memory data warehousing workloads, and early results indicate that Quickstep is over an order-of-magnitude faster than existing platforms including Spark 2.0 and PostgreSQL 9.6 Beta1.

Quickstep has been incubating since 2016-03-29.

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

1. Acquire early adopters
2. Build a community.
3. Create an ASF release.

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

No.

How has the community developed since the last report?

We were incubated into Apache recently, and the developers are actively learning about the Apache way from our mentors.

How has the project developed since the last report?

We have fully transitioned to the ASF for all development work.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Julian Hyde (jhyde):

Getting development onto ASF git took a while, but I think we're there.

I'll work with the project to refine the "Acquire early adopters" goal into more concrete outreach steps.

Roman Shaposhnik (rvs):

Really, at least goal #2 should be "make a first ASF release". Hopefully we can work with the community to get a sense how important it is in addition to growing a community (all 3 goals in this report).

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 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. Mailing list activity since last report (April-01-2016 to May-31-2016):
 @dev 1090
 @user 118
 @commit 174
2. Issues (JIRAs) created/resolved since last report (after April-01-2016 before Mar-31-2016):
 Created: 94
 Resolved: 92

How has the project developed since the last report?

1. Released one minor releases :
 0.5.3 released on 05/31/2016
2. Members have been working on a next major release of 0.6.0
 - a. Working on fixing bugs in the current 0.6.0 features

Date of last release:

May-31-2016

When were the last committers or PMC members elected?

1. Colm O hEigeartaigh has been added as committer on April 19, 2016

Signed-off-by:

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

SAMOA

SAMOA provides a collection of distributed streaming algorithms for the most common data mining and machine learning tasks such as classification, clustering, and regression, as well as programming abstractions to develop new algorithms that run on top of distributed stream processing engines (DSPEs). It features a pluggable architecture that allows it to run on

several DSPEs such as Apache Storm, Apache Flink, and Apache Samza.

SAMOA has been incubating since 2014-12-15.

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

1. Grow the community
2. Elect new PMC members
3. Release more often

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?

Mailing list activity (March 2016 - May 2016):

* @dev: 202 messages

Jira issues backlog (March 2016 - May 2016):

* Created: 3

* Resolved: 9

We have a number of new users showing up on the mailing list, which is encouraging.

How has the project developed since the last report?

After the decision taken by the community to drop support for S4, we removed the support for S4. We fixed automatic build issues, we simplified input from HDFS, and we worked on adding support for Apache Gearpump. Two talks were given to advertise Apache SAMOA, one at ApacheCon NA Big Data by Nicolas Kourtellis, and another one at the "J on the beach" conference by Albert Bifet.

Date of last release:

2015-07-21

When were the last committers or PMC members elected?

None

Signed-off-by:

[X](samoa) Alan Gates
[](samoa) Ashutosh Chauhan
[](samoa) Enis Soztutar
[X](samoa) Ted Dunning

Shepherd/Mentor notes:

Alan Gates:

I sent a mail to the community poking on the long wait since the last release and whether they think there will be a release soon. I also expressed my concern on not having added any committers and asked if there were any contributors who looked like they might be a good candidate for committership.

Singa

Singa is a distributed deep learning platform.

Singa has been incubating since 2015-03-17.

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

1. Grow the community. We anticipate stronger growth of the community when the upcoming, major release (v1.0) is available. SINGA v1.0 will have

new programming abstractions and be able to run on heterogeneous hardware (e.g., using Cuda and OpenCL). It will also have improved usability in terms of installation and training new models, etc (see details below).

2. Clean and improve existing documentations. We plan to replace the website template with a better one.
3. Focus on Rafiki --- SINGA-based deep learning as a service.

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?

The Github repository has scored 602 stars and 205 forks.

There were 107, 73 and 37 emails from dev@ list in March 2016, April 2016 and May 2016 respectively.

There were 9, 35, and 30 commits in March 2016, April 2016 and May 2016.

How has the project developed since the last report?

In April 2016, we released our third version, i.e. v0.3.0, with the following major features:

- * Training on GPU cluster
- * Python wrapper improvements
- * New SGD updaters
- * Remove several dependencies for single node training.
- * Heterogeneous training with CPU and GPU.
- * Support cuDNNv4.

In May, we started working on the first major release (v1.0) with the following notable features(<http://singa.apache.org/develop/schedule.html>):

- * Cross-platform
- * Support heterogeneous hardware: CUDA, OpenCL
- * Enhance programming abstraction, focusing on Tensor as core abstraction.
- * Python binding, cloud integration (Spark, Hadoop, etc.)
- * Runtime optimization.

We are working on a system called Rafiki which uses SINGA to provide deep learning as a service and has the following features:

- * Sharing of pre-training models
- * Visualizing training progress
- * Configuring model via GUI (web browsers)

Date of last release:

2016-04-20

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[](singa) Daniel Dai
[X](singa) Alan Gates
[X](singa) Ted Dunning
[](singa) Thejas Nair

Shepherd/Mentor notes:

Alan Gates:

I sent an email to the community sharing my concerns about the slow progress of adding new committers and asking if they had any contributors

that looked like they might be ready for committership.

Ted Dunning:

The trend in dev mailing list traffic is worrisome as well. The opposite trend in JIRA volume is curious.

Streams

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

Streams has been incubating since 2012-11-20.

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

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

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

of contributors submitting and reviewing pull requests and participating in discussions is still too low. Getting the attention of developers who work with social datasets and highlighting the ways that adopting and contributing to streams could provide value for their projects remains our biggest challenge and opportunity in community growth.

How has the community developed since the last report?

Several recently in-active project contributors participated again this quarter.

How has the project developed since the last report?

<http://streams.incubator.apache.org> has been completely overhauled, and following several talks given at ApacheCon NA 2016, a new blog post, and a renewed focus on marketing we're beginning to see a payoff in the form of a consistent level of unique visitors, according to google analytics.

25 Issues resolved + 4 Issues resolved by open pull requests

Date of last release:

2015-04-26

When were the last committers or PMC members elected?

2015-02-23: Robert Douglas elected as committer / PMC member

Signed-off-by:

[](streams) Matt Franklin
[X](streams) Ate Douma
[](streams) Craig McClanahan

Taverna

Taverna is a domain-independent suite of tools used to design and execute data-driven workflows.

Taverna has been incubating since 2014-10-20.

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

1. Project maturity evaluation

2. Releasing main code base (RC under vote)
3. Engaging user/devops community

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

ASF guidance on US export ECCN cryptography registration
<https://www.apache.org/dev/crypto.html>
has not yet been updated for the 2010 rule changes, confusion around this classification caused Taverna delays.

Lack of active mentors - would appreciate another volunteer.

How has the community developed since the last report?

Two new GSOC students now active.

2-3 other new potential contributors discussing new work.

A total of 4 ICLAs registered - however these folks need further encouragement and guidance to submit regular pull requests.

dev@taverna mailing lists stats:

Mar 2016: 312
Apr 2016: 102
May 2015: 184

users@taverna:

Mar 2016: 12
Apr 2016: 6
May 2015: 1

JIRA issues over last 90 days:

49 created
25 resolved

How has the project developed since the last report?

Preparations for release of the main workflow engine has had focus, including build stability across platforms, documentation update and US Export classifications of encryption items (which caused some delays).

Taverna Engine/Command Line Release candidate now under vote.
<https://lists.apache.org/thread.html/Zpt1lsoyc7qf2so>

Traffic on users@ show interest mainly on Taverna Server, so this component should probably be the development focus after GSOC.

Date of last release:

2016-03-11 taverna-osgi-0.2.1-incubating
2016-03-11 taverna-language-0.15.1-incubating

Currently in [VOTE]:
taverna-engine, taverna-common-activities, taverna-commandline
<https://lists.apache.org/thread.html/Zpt1lsoyc7qf2so>

When were the last committers or PMC members elected?

2015-12-04

Signed-off-by:

[x](taverna) Andy Seaborne
[](taverna) Daniel J Debrunner

[](taverna) Suresh Srinivas
[](taverna) Suresh Maru
[x](taverna) Marlon Pierce

Shepherd/Mentor notes:

Andy Seaborne:

There is still a lack of active mentors. I'm dipping in when I can.

Drew Farris (shepherd):

Little mentor activity on the mailing lists, could use another active mentor. Healthy activity in the community.

Tephra

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

Tephra has been incubating since 2016-03-07.

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

1. Complete project setup
2. Regular releases
3. Improve community engagement

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

There are no issues requiring board attention at this time

How has the community developed since the last report?

- 1) Apache Tephra now powers ACID transactions in Apache Phoenix (since v4.7.0).
- 2) We presented Apache Tephra as a part of the Apache Phoenix talk at HBaseCon 2016.
- 3) We also presented a more elaborate version of the above talk at PhoenixCon 2016.
- 4) We are making progress towards setting up of Project page.

How has the project developed since the last report?

- 1) We had our first release (0.8.0-incubating) since incubation.
- 2) We have added support for Apache HBase v1.2 in the latest release (0.8.0-incubating).

Date of last release:

2016-05-31

When were the last committers or PMC members elected?

None since the May 2016 report

Signed-off-by:

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

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. More external contributions being made

How has the project developed since the last report?

1. Contrinue working with JeroMQ community to further their transition into MPL v2 and away from LGPL. 3 committers to go.
2. Substantial effort in our build system to make Apache compliant releases.
3. Addressing issues opened by community

Date of last release:

None since incubation.

When were the last committers or PMC members elected?

No new additions since incubation

Signed-off-by:

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

Trafodion

Trafodion is a webscale SQL-on-Hadoop solution enabling transactional or operational workloads on Hadoop.

Trafodion has been incubating since 2015-05-24.

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

1. Develop community and bring in more diverse contributors.
2. Finish our second release, including convenience binaries and support for Apache HBase 1.x without a commercial distribution.
3. Continue the momentum and become more integrated with the rest of the Apache community.

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?

- * Participation in public lists has been steady:
722 messages in the codereview forum, 1045@commits,
689@dev, 1294@issues and 240@user.
- * Upcoming release 2.0 was verified by several new users.
- * Release 2.0 is being actively discussed in the dev forum.

How has the project developed since the last report?

- * For our second release in incubation, Trafodion 2.0 release candidate 3 has passed an internal vote. It was submitted for a vote in the incubator PMC on May 31. Many thanks go to Steve Varnau, the release manager!
- * Added documentation in form of HTML and PDF manuals.
- * 314 commits from 28 contributors.
- * 171 JIRAs filed and 165 resolved in the last three months.

Date of last release:

2016-01-11 1.3.0

When were the last committers or PMC members elected?

RuoYu Zuo was announced as a new committer in February.
The last PMC members were announced in November 2015.
We expect to add an additional committer soon.

Signed-off-by:

[](trafodion) Devaraj Das
[](trafodion) Enis Söztutar
[X](trafodion) Michael Stack

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 had a few discussions regarding technical developments.

How has the project developed since the last report?

- Made some code improvements
- Introduced GitHub based pull requests as accepted way for change submissions.

Date of last release:

March 2016

When were the last committers or PMC members elected?

February 2016

Signed-off-by:

```
[ ](wave) Christian Grobmeier  
[x](wave) Upayavira
```

Shepherd/Mentor notes:

Upayavira:

Wave is a slow moving, but functioning, project. It needs to attract additional developers to focus on attracting new developers in order to be able to exit the incubator.