

September 2016

/\ *FINAL* (!)

Incubator PMC report for September 2016

Timeline

Wed September 07	Podling reports due by end of day
Sun September 11	Shepherd reviews due by end of day
Sun September 11	Summary due by end of day
Tue September 13	Mentor signoff due by end of day
Wed September 14	Report submitted to Board
Wed September 21	Board meeting

Shepherd Assignments

Drew Farris		Ranger	
Drew Farris		log4cxx 2	
John Ament		Pirk	
John Ament		Trafodion	
Justin Mclean		Gearpump	
P. Taylor Goetz		Carbon Data	
P. Taylor Goetz		Juneau	
P. Taylor Goetz		Tamaya	
Timothy Chen		Myriad	
Timothy Chen		SAMOA	
Timothy Chen		Streams	
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="38f859bb-e562-4abb-872b-fba0c906db2b"><ac:plain-text-body><![CDATA[[none]	AriaTosca]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="9c368531-ec58-4305-97e3-7df732381de9"><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="487a5abd-b781-414d-bd07-bb81175d4b11"><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="b98a11e7-2783-4a7a-9258-56814228f680"><ac:plain-text-body><![CDATA[[none]	Distributed Log]]></ac:plain-text-body></ac:structured-macro>
Kathy Saunders		Edgent	
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="4cc1051d-843f-4ce5-b4c0-44e32b3e9b09"><ac:plain-text-body><![CDATA[[none]	Fineract]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6b4d1a72-8383-490d-8332-34ed6fb31a33"><ac:plain-text-body><![CDATA[[none]	HTTrace]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="8ffeab3f-fda6-4951-b7e5-928f050326d1"><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="80d2a04e-5a7b-4ec6-97d3-33a3a2a81535"><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="3bc2a13d-d766-4f11-86de-a989b8c583ac"><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="e7fd269e-7986-4461-a997-c8bca5fa1f3d"><ac:plain-text-body><![CDATA[[none]	Pony Mail]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="2a8efa17-72fb-4da2-a458-d80156b11657"><ac:plain-text-body><![CDATA[[none]	PredictionIO]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="8e3e2e88-91cc-4af4-b34a-2e5363d4a2ca"><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="e1709719-5b3a-42b7-a67e-e05d96a8f9c4"><ac:plain-text-body><![CDATA[[none]	SensSoft]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3f6b71c5-1c7f-46c3-871f-d7d9b8a14f13"><ac:plain-text-body><![CDATA[[none]	Singa]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e083b9e8-3d77-4f2b-8bc5-1f8f47fab7ca"><ac:plain-text-body><![CDATA[[none]	Taverna]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7ebf4d63-aed7-45fd-a11d-64851a6f49f7"><ac:plain-text-body><![CDATA[[none]	Tephra]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="04030ef7-83d9-43de-80c5-a2822b2cc2b4"><ac:plain-text-body><![CDATA[[none]	Traffic Control]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5c11fe58-c1ce-4362-9577-977a4e041838"><ac:plain-text-body><![CDATA[[none]	Toree]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="0c73ce91-c347-446d-8175-4db02e682478"><ac:plain-text-body><![CDATA[[none]	Wave]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for September 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 61 podlings currently undergoing incubation -- an all-time high.

* Community

No IPMC roster changes this month.

* New Podlings

- Annotator
- AriaTosca

* Graduations

None this month.

* Releases

The following releases entered distribution during the month of August:

- 2016-08-01 Apache Pony Mail 0.9.incubating
- 2016-08-02 Apache Mnemonic 0.2.0-incubating
- 2016-08-03 Apache Eagle 0.3.0-incubating
- 2016-08-09 Apache Datafu (incubating) 1.3.1
- 2016-08-10 Apache Gearpump 0.8.1-incubating
- 2016-08-20 Apache Ranger (incubating) 0.6.1
- 2016-08-22 Apache Metron 0.2.0BETA-incubating
- 2016-08-22 Apache Beam 0.2.0-incubating
- 2016-08-24 Apache CarbonData 0.1.0-incubating
- 2016-08-29 Apache Pirk 0.1.0-incubating

* IP Clearance

None this month.

* Legal / Trademarks

- The Fluo podling grappled with how to handle the fluo.io domain, io.fluo Maven group ID, and the fluo-io GitHub organization. Existing high profile domains which contain our trademarks (accumulosummit.com, stratahadoopworld.com, search.maven.org, etc.) had contributed to misunderstandings about to what extent Fluo could have a two-pole presence split across fluo.apache.org and a fluo.io entity outside Apache control. Members of the IPMC made it clear that Fluo's graduation would be contingent on resolving this issue.
- The Guacamole podling received guidance on how to integrate with Docker: either participate in the Infra-supported offering (<chub.docker.com/r/apache>), or treat Docker as a downstream distribution channel and create a customized offering in compliance with

release/branding/incubation/etc policies.

* Miscellaneous

- The OpenAZ podling has retired.
- The recurring discussion as to whether the Incubator is too full came back for another installment. Mentoring for healthy podlings seems to scale well, but other aspects such as release checking and cleaning up after troubled podlings present more difficulty.

* Credits

- Report Manager: Marvin Humphrey / John D. Ament

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

* Still getting started at the Incubator

- AriaTosca
- Traffic Control

* Not yet ready to graduate

No release:

- DistributedLog
- Edgent
- Fineract
- HTrace
- Juneau
- log4cxx2
- PredictionIO
- SensSoft
- Toree

Community growth:

- CarbonData
- CommonsRDF
- Gearpump
- Mnemonic
- Myriad
- Omid
- Pirk
- Pony Mail
- Singa
- SAMOA
- Tephra
- Trafodion
- Wave

* Ready to graduate

- Atlas
- Ranger
- Taverna

* Considering retirement

- Streams

* Did not report, expected next month

- Quickstep
- Tamaya

* No Mentor signoff, report removed

- MRQL

Table of Contents

AriaTosca
Atlas
CarbonData
CommonsRDF
DistributedLog
Edgent
Fineract
Gearpump
HTrace
Juneau
log4cxx2
Mnemonic
Myriad
Omid
Pirk
Pony Mail
PredictionIO
Quickstep
Ranger
SAMOA
SensSoft
Singa
Streams
Tamaya
Taverna
Tephra
Traffic Control
Trafodion
Toree
Wave

AriaTosca

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

AriaTosca has been incubating since 2016-08-27.

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

1. Establish a formal release process and schedule, allowing for dependable release cycles in a manner consistent with the Apache development process.
2. Grow the community around the project.
3. Migrate existing users to apache mailing lists, redirect the present external project links to Apache.

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

None

How has the community developed since the last report?

The project is still being setup in Apache Incubator and is presently.

Received Software Grant Agreement (SGA) from Gigaspaces.

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

How has the project developed since the last report?

This would be the first report for AriaTosca.

The project is presently being setup in Incubator.

Date of last release:

None

When were the last committers or PMC members elected?

Project is being established in incubator with the proposed initial set of committers.

Signed-off-by:

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

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. We have added 4 new committers since the last report from different organizations. Around 15 new contributors were added since the last report. The current number of contributors is around 60.
2. The activity on the atlas-dev mailing list stands at around 900 messages per month between June and August. [1]

How has the project developed since the last report?

1. A major release, 0.7-incubating, was made in the month of July. This release added 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, Kafka message versioning, support for creating business taxonomies, enhanced user interface, better performance etc.
2. Work is proceeding in 0.8-incubating to stabilize many of these features and also add several strategic improvements including abstractions to the underlying graph database to make it easier to move to later versions, rolling upgrades, etc.
3. A total of 302 issues were reported between June 1st 2016 and August 31st 2016. 266 issues were resolved in the same time frame [2],[3]

Date of last release:

2016-07-09: 07-incubating [4]

When were the last committers or PMC members elected?

1. The PPMC voted and elected 4 new members as committers to the project on 2016-07-05: Tom Beerbower, Keval Bhatt, David Kantor and Darshan Kumar.

[1] http://mail-archives.apache.org/mod_mbox/incubator-atlas-dev/

[2] <https://issues.apache.org/jira/browse/ATLAS-1149?jql=project%20%3D%20ATLAS%20AND%20created%20%3E%3D%202016->

06-01%20AND%20created%20%3C%3D%202016-08-31

[3] <https://issues.apache.org/jira/browse/ATLAS-775?jql=project%20%3D%20ATLAS%20AND%20resolved%20%3E%3D%202016-06-01%20AND%20resolved%20%3C%3D%202016-08-31>

[4] https://mail-archives.apache.org/mod_mbox/www-announce/201607.mbox/%3CCAAbv78fuzCv51E4Pn8fg7nwkJyOUQo9z2WNhsiz2EstbTvjlGw@mail.gmail.com%3E

Signed-off-by:

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

CarbonData

Apache CarbonData is a new Apache Hadoop native file format for faster interactive query using advanced columnar storage, index, compression and encoding techniques to improve computing efficiency, in turn it will help speedup queries an order of magnitude faster over PetaBytes of data.

CarbonData has been incubating since 2016-06-02.

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

1. Prepare new releases
2. Increase both dev and user communities
3. Publish and update website to promote CarbonData

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

None

How has the community developed since the last report?

The community activity increased: many new users started to use and test CarbonData, we had more than 100 new issues created during Aug; One user formally deployed CarbonData to their business system, the query efficiency speeded up 50 times (the user is the biggest internet finance company in China)

How has the project developed since the last report?

Code donation has been done and all resources have been created by INFRA (git, github mirror, mailing list, Jira, ...).

We also created the Jenkins CI jobs.

We did the first release (0.1.0-incubating) and we are preparing a new one.

We're also working on the website (with improved look'n feel).

We're also preparing some talks and presentations about CarbonData.

Date of last release:

2016-08-27

When were the last committers or PMC members elected?

2016-07-15

Signed-off-by:

[X](carbodata) Henry Saputra
[X](carbodata) Jean-Baptiste Onofre
[X](carbodata) Uma Maheswara Rao G

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. 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.
2. Engage wider RDF community
3. Graduate to Apache Commons

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

None

How has the community developed since the last report?

The community has been fairly quiet during the summer.

Mailing list stats:

dev@commonsrdf traffic

Jun 2016: 10

Jul 2016: 21

Aug 2016: 3

Jira stats for last 90 days:

Created: 5

Resolved: 1

How has the project developed since the last report?

Building on the idea of Commons VFS-style integration we have added branches that add bindings for Eclipse RDF4j (formerly Open RDF Sesame), Apache Jena and JSON-LD-Java.
<https://github.com/apache/incubator-commonsrdf/branches>

This has raised interest within Jena
<https://issues.apache.org/jira/browse/JENA-1015>

We should now be moving towards a 0.3.0 release, and working towards graduation to Apache Commons PMC, we are adapting the build to be more compatible with the Commons release process <https://commons.apache.org/releases/>

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

[](commonsrdf) Gary Gregory

[X](commonsrdf) Lewis J McGibbney

DistributedLog

DistributedLog is a high-performance replicated log service. It offers durability, replication and strong consistency, which provides a fundamental

building block for building reliable distributed systems.

DistributedLog has been incubating since 2016-06-24.

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

1. Getting an Apache release out
2. Improve general documentation
3. Grow user base

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

There is currently no issue.

How has the community developed since the last report?

1. A few people began actively to involve in the project discussion
2. Some project ideas have been initiated and discussed on the mailing list

How has the project developed since the last report?

1. The wiki has been filled up with development and contribution information.
2. The existing website content has been pushed to asf-site branch and live under <http://distributedlog.incubator.apache.org>. The new website has been discussed.
3. The pom file was updated to follow Apache's principle and the status page is updated.
4. Various project ideas are under developing.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Edgent

Apache Edgent 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 Edgent provides efficient and timely analytics across the whole IoT ecosystem: from the center to the edge.

Edgent 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 Edgent project
2. Set up new build system, allowing us to remove code dependencies from our repository
3. Create the first Apache release of Edgent

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?

Significant progress has been made on a new build system, using Gradle. One of the goals is to remove third party code, so we can work toward a release.

The community has been active and participating in voting. Heavy commit activity, but light Jira activity, recently as the name change was implemented in the code, documentation and web site.

How has the project developed since the last report?

According to JIRA, the project has added the following:

- * June: 29 new issues; 22 issues resolved.
- * July: 20 new issues; 15 issues resolved.
- * August: 7 new issues; 7 issues resolved.

August activity was largely build and renaming activities, tracked under a few Jira issues.

The name for the project was changed from Apache Quarks to Apache Edgent. The new name went through a suitable name search and was approved. The infrastructure, code and documentation changes have been completed.

A logo was voted on and chosen, and a wiki was set up for the Edgent community.

Date of last release:

We have not created an Apache release as of yet. We need to complete build changes and licensing work to create a release.

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:

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

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. Getting a first release out - we are still inhibited by Hibernate issues which are more deeply rooted than we realized.
2. Moving the community from the Mifos infrastructure to the Apache infrastructure - most discussions are taking place out of public view.
3. Resolving the questions around introduction of a rework of the original code - in particular repository/infrastructure/release questions which result from the microservice architecture.

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

The person who creates and submits the board reports will cease to perform this task until the code rework is ready to be introduced to the community. It is possible that this will endanger the board reports.

How has the community developed since the last report?

Several prospective new committers have recently joined the Mifos community through Google Summer of Code and as new partners that we are working to add as committers to the Apache Fineract.

How has the project developed since the last report?

We have tried and failed to produce an apache release again. A release was produced by the mifos community outside of apache and announced on the apache fineract user list.

The release included work intended to move us closer to apache policies. In particular, the removal of Hibernate annotations from the code uncovered a circular dependency in the data structure which was not caught in testing. Two production systems from early adopters are out-of-service today.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

Jim Jagielski was added as a mentor.

Signed-off-by:

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

Shepherd/Mentor notes:

Roman Shaposhnik:

I'll be following up with the community on the report generation point which clearly needs to be clarified.

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 with Apache 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?

Increased community contribution on integration with Apache Cassandra.

How has the project developed since the last report?

We released Apache Gearpump 0.8.1-incubating, which is Gearpump's first release as an Apache incubator project.

Integration within Apache Beam (Gearpump Runner) has been merged into a feature branch.

Continued integration within akka-streams (Gearpump Materializer). Three

PRs committed.

28 issues created and 18 resolved.

Date of last release:

2016-08-16

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:

Justin Mclean:

PPMC are looking for additional mentors.

HTrace

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

Apache HTrace has been incubating since 2014-11-11.

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

1. Continue to grow the Apache HTrace community
2. Continue to explore projects integrating Apache HTrace
3. Continue to develop and release stable Apache 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?

- * The dev@htrace mailing list continues to see new comers which is great.
- * Although a GSoC-related project was not successfully completed, the Apache Kudu and Apache HTrace communities have made contact. This is part of the drive to have more projects integrate Apache HTrace.

How has the project developed since the last report?

- * A discussion thread is currently ongoing to release HTrace 4.2. This will include over 40 issues addressed and may even be our last incubating release before we attempt to graduate.

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
 (htrace) Billie Rinaldi
 (htrace) Michael Stack

Juneau

Apache Juneau is a toolkit for marshalling POJOs to a wide variety of content types using a common framework, and for creating sophisticated self-documenting REST interfaces and microservices using VERY little code.

Juneau has been incubating since 2016-06-24.

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

1. Build up community with non-IBM and non-Salesforce contributors.
2. Release a version of the product to spark external interest in the project.
3. Need to determine where built libraries will be hosted for download.

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

None.

How has the community developed since the last report?

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

How has the project developed since the last report?

Website is now running.
Jira is now running and accepting bug reports.
Jenkins is now running and producing builds.
Confluence space has been created with articles for developers to get started contributing.
Passes apache-rat check.

Date of last release:

None.

When were the last committers or PMC members elected?

Signed-off-by:

(juneau) Craig Russell
 (juneau) Jochen Wiedmann
 (juneau) John D. Ament

log4cxx2

Logging for C++. N.B. This is a reboot of the Log4cxx podling which previously graduated.

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.

There was a discussion about some kind of "maintenance mode" where releases and community growth doesn't matter much. As that mode doesn't exist, it has been suggested to release, graduate from incubator and just live along other logging projects and in a comparable state like log4cxx.

Currently it seems that there's only one official project member with commit access left answering mails. Even though that member and some contributors worked on being able to create a release in the past, contrary to expectations that project member didn't find the time to deal with a release as necessary.

It seems most likely that if a release doesn't happen until the next report is due, it won't at all anymore. At least not by that one left member.

How has the community developed since the last report?

One new contributor took quite some time to look and fix the release process. Two support mails from users, no new project members.

How has the project developed since the last report?

Robert Middleton took quite some time to look at the build and release process on Windows and Linux in the last two months and came to the conclusion that the current Maven based process for both seems broken. He provided a shell based alternative for the release and CMAKE build starting points, which hasn't been merged into or tested by the project yet.

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:

Drew Farris:

One mentor active within in the past month. There is a single active committer and minor progress towards a release. The current build system is proving problematic for the contributor involved. There is some discussion about moving to the attic and this seems appropriate

Mnemonic

Apache Mnemonic is an advanced hybrid memory storage oriented library, we've proposed a non-volatile/durable Java object model and durable computing service that bring several advantages to significantly improve the performance of massive real-time data processing/analytics. developers are able to use this library to design their cache-less and SerDe-less high performance applications.

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

1. Apache releases: we have release version 0.2.0-incubating
2. Encouraging community participation and contribution
3. Invented Durable Computing Model and Service on top of Durable Object Model

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?

Our last report was in April. Since then

- * Twelve new contributors have joined and seven contributors have submitted their patches.
- * We need to encourage more discussion on the dev list in the coming months

How has the project developed since the last report?

- * Added memory service and testcases
- * Added computing service and testcases
- * Re-organized sub-modules
- * Added Apache Rat and Checkstyle to ensure compliance
- * Added release.sh, runall.sh script tools
- * Supported Docker dev. construction
- * Removed the dependency of customized 3rd party library
- * Removed the dependency of special version of nvml
- * Constructed project website (WIP)
- * Supported discrete addressing for memory service
- * and other bugfixes, features, improvements

Date of last release:

2016-08-02

When were the last committers or PMC members elected?

None

Signed-off-by:

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

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. Expand user/contributor/committer community
2. Gather more support from Hadoop and Mesos vendors
3. Release new versions (0.3 in planning), with more enterprise features

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 picked up traffic (153 messages in June) around the 0.2.0 release. 236 messages since the last report.
- * 6 new members on the dev@ mailing list. 3 new patch contributors.
- * Bi-weekly dev syncs with 2-5 participants. Some cancelled due to lack of attendance. Minutes at <http://s.apache.org/8kF>
- * Myriad to be presented at HadoopWorld/Strata NYC and a talk submitted for ApacheCon BigData EU.

How has the project developed since the last report?

- * Myriad 0.2.0 released. Myriad 0.3.0 in development/planning.
- * 10 commits since 6/1. 8 JIRAs resolved.

Date of last release:

2016-06-20 myriad-0.2.0-incubating

When were the last committers or PMC members elected?

2016-06-16 Yuliya Feldman

Signed-off-by:

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

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. Boost project dissemination in the Apache community.
2. Collaborate/integrate Omid with other Apache projects.
3. Get positive feedback from other projects currently integrating/using Omid.

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?

- * Omid presented at Hadoop Summit 2016 in San Jose. Some attendants manifested interest and asked for future plans of integrating Omid with Apache Phoenix.
- * First contact with Apache DistributedLog community for a possible integration with Omid.
- * Blog entry on Apache Omid ready to be published in Yahoo's hadoop blog site.
- * Quarter Stats (from: 2016-06-01 to: 2016-08-31):

Metric	counts
# of msgs in dev list	204
Active Contributors (incl mentors)	6
Jira New Issues	10
Resolved Issues	12
Pull Requests merged	7
Pull Requests proposed	10

How has the project developed since the last report?

First version of Omid released under the Apache incubator program. It has also been integrated with the Apache Hive project.F

Date of last release:

2016-06-24

When were the last committers or PMC members elected?

At incubation acceptance (2016-03-28)

Signed-off-by:

(omid) Alan Gates
 (omid) Lars Hofhansl
 (omid) Flavio P. Junqueira
 (omid) Thejas Nair
 (omid) James Taylor

Pirk

Pirk is a framework for scalable Private Information Retrieval (PIR).

Pirk has been incubating since 2016-06-17. The initial code for Pirk was granted on 2016-07-11.

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

1. Grow the community to establish diversity of background and expertise.
2. Establish a formal release process and schedule, allowing for dependable release cycles in a manner consistent with the Apache development process.
3. Add to and improve the Pirk user documentation and examples both on the website and in the codebase

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 officially welcomed two new committers (Tim Ellison and Suneel Marthi, both also mentors). The mailing list activity, pull requests, and page views to the Apache Pirk website (via Google Analytics) have also greatly increased since the last report.

How has the project developed since the last report?

The project completed it's first release, 0.1.0-incubating, since the last report and provided a release guide on its website documenting the process. Additionally, the project has ensured that all dependencies are compliant with the Apache License version 2.0 and that all code and documentation artifacts have the correct Apache licensing markings and notice. The project has also undertaken a refactor of some of the initial packages and provided initial streaming implementations.

Date of last release:

August 29, 2016

When were the last committers or PMC members elected?

Pirk elected two new committers (Tim Ellison and Suneel Marthi, both mentors) on August 18, 2016.

Signed-off-by:

(pirk) Billie Rinaldi
 (pirk) Joe Witt
 (pirk) Josh Elser
 (pirk) Suneel Marthi
 (pirk) Tim Ellison

Shepherd/Mentor notes:

Josh Elser:

Extremely happy to see the first release completed so soon after entering incubation. Very pleased with the interaction and responsiveness at PPMC and IPMC levels during the vote process. They are well on their way.

Suneel Marthi:

I must add that the first release candidate was put out for PPMC/IPMC Vote in less than a month since the initial code was pushed to the apache github repo.

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. Make contributions easier (lower the bar for hacking on code/docs/UI)
3. Increase awareness of the 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?

Slight uptick in active members of the community, we now count roughly 10 active people participating either on the mailing list or IRC.

How has the project developed since the last report?

Work has begun on a redesign of the UI, allowing for contributions more easily. This is codenamed 'coffee and cake', and is expected to reach a usable phase in a month or two.

Date of last release:

2016-08-02

When were the last committers or PMC members elected?

None since inception. See graduation info point 1. The codebase needs to be more accessible before contributions can occur.

Signed-off-by:

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

PredictionIO

PredictionIO is an open source Machine Learning Server built on top of state-of-the-art open source stack, that enables developers to manage and deploy production-ready predictive services for various kinds of machine learning tasks.

PredictionIO has been incubating since 2016-05-26. The initial code for PredictionIO was granted on 2016-06-16. A second grant of PredictionIO Templates is in the works and anticipated to be complete in September 2016.

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

1. Establish a formal release schedule and process, allowing for dependable release cycles in a manner consistent with the Apache way.
2. Grow the community to establish diversity of background.
3. Transition remaining former PredictionIO users from google-groups to ASF mailing lists.

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. Users have begun to migrate from the old Google group to the Apache user list
2. Three new contributors have become active and have been making consistently acceptable contributions and are on path to becoming committers.

How has the project developed since the last report?

1. The old Prediction.io docs site has been migrated to <http://predictionio.incubator.apache.org>
2. About seven templates from the original PredictionIO have been ported to Apache PredictionIO, pending a decision on a 2nd grant. The pending 2nd grant will not block a PredictionIO release.

Date of last release:

No releases have been made yet. Present activities are focussed towards a planned first Apache Release in September 2016.

When were the last committers or PMC members elected?

Paul Li was elected as committer and PMC member on Aug 30, 2016.

Signed-off-by:

(predictionio) Andrew Purtell
 (predictionio) James Taylor
 (predictionio) Lars Hofhansl
 (predictionio) Luciano Resende
 (predictionio) Xiangrui Meng
 (predictionio) Suneel Marthi

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 (June-01-2016 to August-31-2016):
 @dev 1423
 @user 124

@commit 294

2. Issues (JIRAs) created/resolved since last report (after June-01-2016 before August-31-2016):
Created: 155
Resolved: 168

3. Adding more community members to involve in Apache Ranger
 - * Added new contributors: Suneel Marthi
 - * Apache Hawq community is working with Ranger to build a authorization engine for Apache Hawq (RANGER-782)

How has the project developed since the last report?

1. Released one major release :
0.6.0 released on 07/18/2016
2. Released one minor release :
0.6.1 released on 08/20/2016
3. Members have been working on a next major release of 0.7.0
 - a. Working on scoping of the release and features

Date of last release:

August-20-2016

When were the last committers or PMC members elected?

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

Signed-off-by:

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

Shepherd/Mentor notes:

Drew Farris:

Three mentors active. New release available in August. Community growth. Discussion of Graduation Steps. A healthy podling.

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 S4, and Apache Samza.

SAMOA has been incubating since 2014-12-15.

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

1. Perform another release
2. Grow the developer base
3. Grow the user 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?

Mailing list activity (June 2016 - August 2016):

* @dev: 77 messages

Jira issues backlog (June 2016 - August 2016):

* Created: 1

* Resolved: 0

We have had interest from both the Apex and Gearpump communities.

How has the project developed since the last report?

We have been preparing a second release of the project, which was delayed due to some minor issues in the RC and then the summer break for most developers.

Date of last release:

2015-07-21

When were the last committers or PMC members elected?

None

Signed-off-by:

(samoa) Alan Gates

(samoa) Ashutosh Chauhan

(samoa) Enis Soztutar

(samoa) Ted Dunning

SensSoft

SensSoft is a software tool usability testing platform

SensSoft has been incubating since 2016-07-13.

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

1. Establishing our Users mailing list and getting incubating releases to the user community
2. Completing the initial build on our website(s) (in progress)
3. Consolidating our Quick Start deployment guides on Apache infrastructure (Confluence, website) that allow for an end-to-end deployment of data collection services with containers for assets (Docker Images) for speedy deployment (in progress)

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

No

How has the community developed since the last report?

1. Community members have been working with DoD and Industry partners to bring them on as contributors to the project (ex. Other industry leaders in HCI/UX offerings for DoD have begun offering SensSoft products as part of their technical offerings).
2. The community has established a Twitter feed (@ApacheSensSoft) and are growing our community by following others (e.g., Kitware, etc.)
3. Community members from Draper Labs have made arrangements to present at uie21 and the emetrics summit as exhibitors to capture community followers in the UI/UX and business analytics communities.

How has the project developed since the last report?

1. The transition to Apache Infrastructure is complete!!!
2. We have moved over all our development tickets into Apache Jira. All development is now taking place @Apache which is great.
3. We are working on new patches and version releases for TAP, UserAle.js, Distill products (sub projects under SENSSOFT), which will now be

visible through commits.

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

N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Chris Mattmann:

Project has primarily being Championed by Lewis John McGibbney. Lots of lots of JIRA activity and mailing list activity from the Draper contributors, including Joshua Poore. Doing a good job starting up.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

XXXX-XX-XX

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. One committer was added in July, more would be nominated soon after releasing V1.0. We are also expanding the user base by participating in meetups of local organizations, including PyDataSG.
2. Optimize V1 release in terms of memory and efficiency for both stand-alone training and distributed training.
3. Add more examples and tutorials to attract users, including deployment on Android phones.

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

No

How has the community developed since the last report?

There were 187, 100 and 69 emails from dev@ list in June 2016, July 2016 and August 2016 respectively. The number of commits has increased 50% (about 200 commits) since last report. One new committer was added.

How has the project developed since the last report?

We have been working on V1.0 since the last release. The vote for releasing V1.0 has been sent to the general mailing list.

The new version has major improvements in term of programming abstraction and usability. Some features are listed here,

- * Support heterogeneous hardware: CPP for CPU, CUDA for Nvidia GPU, OpenCL for other devices (including AMD GPUs)
- * New programming abstractions supporting complex deep learning models
- * Re-write the Python binding and provide running examples
- * Cross-platform (Linux and Mac OSX)

- * Upgrade the Cudnn library version to V5 for running RNN models
- * New website for documentation

Date of last release:

2016-04-20

When were the last committers or PMC members elected?

2016-07-20

Signed-off-by:

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

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?

No commits or mailing list activity since the last report.

How has the community developed since the last report?

No commits or mailing list activity since the last report.

How has the project developed since the last report?

No commits or mailing list activity since the last report.

No issues resolved

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

Shepherd/Mentor notes:

Ate Douma:

As indicated in the report, the activity and community of the project has dwindled down to practically zero. As a consequence I've now raised the question on the dev list if better to retire the project. Unless

this triggers some serious activity and interest from the community, I expect an actual vote to retire before the next board report.

Tamaya

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

Tamaya has been incubating since 2014-11-14.

Shepherd/Mentor notes:

John D. Ament:

The project continues to struggle. I'm not sure its recoverable as an ASF project at this time. May make sense to reboot in an external resource, e.g. GitHub, build a true community, and then come back.

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 and graduation
2. Complete IP/license review of unreleased repositories
3. Further engaging user/devops community

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

Preparing to graduate.

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

The community is moving towards graduation, by preparing a Graduation Maturity Assessment
<https://cwiki.apache.org/confluence/display/TAVERNADEV/2016-03+Taverna+Graduation+Maturity+Assessment>

While we wished we could have released all of Taverna's git repositories through the Incubator before graduating, we realize that with limited developer time, getting all of them to a cleanly releasable state could remove focus from development of the core repositories.

As we an important point of Incubator releases is to ensure the project understands the ASF release process, including release review and voting, IP management and licensing compatibility; as well as accepting contributions. We believe we have gained that after 3 release votes
<https://archive.apache.org/dist/incubator/taverna/source/>
and 7 GSOC/student contributors.

We have however identified the need for a License Review of the remaining git repositories that have not yet been formally released through the Apache Incubator (e.g. to ensure ASF file headers), which we are now completing:
<https://cwiki.apache.org/confluence/display/TAVERNADEV/2016-09+License+review>

We might consider not including all of these repositories as we move towards graduation - we already have an alternative home
<https://github.com/taverna-extras> for dormant extensions.

dev@taverna mailing lists stats:

Jun 2016: 347
Jul 2016: 221
Aug 2016: 170

users@taverna:

Jun 2016: 4
Jul 2016: 2
Aug 2016: 4

JIRA issues over last 90 days:

34 created
12 resolved

How has the project developed since the last report?

We have seen an upsurge in activity with two successful GSOC projects as well as several new plugins from other students. Some of them have shown interest in continued contributions, which we must take into account for proposing new committers.

This has admittedly taken some focus from the committers from the incubator graduation path - so we're now putting effort back into finishing our License Review.

Date of last release:

2016-07-01 taverna-engine-3.1.0-incubating
2016-07-01 taverna-common-activities-2.1.0-incubating
2016-07-01 taverna-commandline-3.1.0-incubating

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 Marru
[](taverna) Marlon Pierce

Shepherd/Mentor notes:

Andy Seaborne:

The Taverna podling is hindered by a lack of active mentors (myself included) - the project has appealed to general@ on a couple of occasions

3 of the 4 active contributors are from one institution and contribution is related to EU project commitments. The codebase is large and has a long history. One option is to identify a "core" and move out the rest until such time as they are sustainable. More active PMC/committers would be a big boost.

The comment from the last report applies:

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. Regular releases
2. Improve community engagement
3. Increase adoption

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

- None at this time.

How has the community developed since the last report?

- 12 subscribers in dev mailing list
- 5 new JIRAs filed since last report
- 2 new contributors submitted patches

How has the project developed since the last report?

- Working on 0.9.0 release

Date of last release:

2016-05-31

When were the last committers or PMC members elected?

- None since coming to incubation

Signed-off-by:

[X](tephra) Alan Gates
[X](tephra) Andrew Purtell
[X](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: Great news! The team has just merged a PR and removed the LGPL license in favor of MPL (<https://github.com/zeromq/zeromq/pull/358>)
2. Make a release: 0.1.x branch should be ready for release. Master has moved to start support for Spark 2.0
3. Grow a diverse community: We should put some emphasis on growing the community and making it diverse (the rule is at least three independent contributors) We have 1 major independent contributor (mariusvniekerk) and at least 2 minor ones (<https://github.com/apache/incubator-toree/graphs/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 - Awaiting MPL release of JeroMQ. Will update branches to remove caveats re: LGPL once the JeroMQ release is available.

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. Making significant progress here.
3. More external contributions being made. Community member 2mariusvniekerk key in porting Toree to work on Spark 2.0. Sitting on PR 56. Working with committer to finish up some details and will merge when ready. This is a significant contribution.

How has the project developed since the last report?

1. JeroMQ community has finalized work to transition to MPL. Waiting for maven release.
2. Addressing issues opened by community

Date of last release:

None since incubation.

When were the last committers or PMC members elected?

No new additions since incubation

Signed-off-by:

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

Traffic Control

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

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

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

1. Software grant from Comcast and others active in the current github.com/Comcast project to ASF. Discussion between Comcast and ASF legal on-going, this is blocking all other progress.

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

See above issue 1.

How has the community developed since the last report?

Established apache.org accounts for all of the initial committers.

How has the project developed since the last report?

- Setup git repo and mailing lists.
- Started Software Grant process.

Date of last release:

We released 1.6.0 on 2016-07-29 under the old process, because the release process (candidates and votes) had already started before our official entry in the incubator. We have started release 1.7.0 under the old process as well, since we don't have the necessary ASF infrastructure in place, and this is pending the SGA.

When were the last committers or PMC members elected?

2016-07-12 - initial committer list.

Signed-off-by:

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

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. Continue the momentum and become more integrated with the rest of the Apache community.
3. Continue to create software releases and to find a release manager for the next one.

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 relatively steady:
789 messages in the codereview forum, 752@commits,
483@dev, 1240@issues and 101@user.
- * We released Trafodion 2.0.0 (source and some binaries) and
Trafodion 2.0.1 (source and convenience binaries).
Steve Varnau was the release manager for both releases.
- * We announced two new committers and three new PPMC members:
 - * Gunnar Tapper (committer and PPMC member), July 5
 - * Steve Varnau (existing committer, now also a PPMC member),
Aug 18
 - * Ming Liu (committer and PPMC member), Aug 19
- * We have seen increased activity in China.
- * We are working on a roadmap for the next release, Atanu Mishra
has volunteered to coordinate it.
- * Community members speaking at public events:
 - * Rohit Jain talked at Data Day Seattle on 7/23
 - * Rao Kakarlamudi talked at the Milpitas Big Data meetup on 8/5

How has the project developed since the last report?

- * Some highlights of Trafodion 2.0.0:
 - * Support for Apache HBase 1.0 and Hadoop distros CDH 5.4, HDP 2.3
 - * Elasticity (adding/removing nodes from a Trafodion instance)
 - * DDL operations are now transaction-protected
 - * BLOB/CLOB support (preview)
 - * ALTER COLUMN for all column attributes
 - * Support for GBK character set in Hive tables
- * 249 commits from 24 contributors.
- * 180 JIRAs filed and 123 resolved June 1 - Aug 30.

Date of last release:

2016-07-07 2.0.1

When were the last committers or PMC members elected?

in August (committers and PPMC members)

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 have ongoing community discussion about how to bring in new developers to ensure sustainability of the project.

A votation session has been scheduled (Sept 28th) to decide if Wave's fork Swellrt is brought/merged into Wave. This would bring SwellRT active developers into Wave and it could facilitate use of Wave by new developers thanks to a new API, removing GWT legacy code.

How has the project developed since the last report?

- Made some code improvements

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:

Once more Wave is on the brink of retirement. However, this time, an offer has been made of code from SwellRT, which is a fork of Wave itself, and a concall has been scheduled for interested parties to discuss whether it is a go-er. It is my (limited) understanding that many of the complexity issues in the Wave code that have prevented community development have been resolved in SwellRT. I will watch to see where this goes, if anywhere.