

December2016

/\ FINAL (!)

Incubator PMC report for December 2016

Timeline

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

Shepherd Assignments

| | | |
|--|--------|---|
| Drew Farris | | Annotator |
| Drew Farris | | Tephra |
| John Ament | | Range r |
| John Ament | | SAMOA |
| Justin Mclean | | MRQL |
| P. Taylor Goetz | | Myriad |
| P. Taylor Goetz | | Quickstep |
| P. Taylor Goetz | | Taverna |
| Timothy Chen | | Singa |
| Timothy Chen | | Spot |
| Timothy Chen | | log4cxx2 |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="37eaa00b-387f-44d7-b50c-b0e2f82784ef"><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="1dab81ce-62ae-4002-9499-2f742392d813"><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="aec8d7ee-2aaf-4cda-9e8f-6225a10fecae"><ac:plain-text-body><![CDATA[| [none] | Gearpump]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="226822a4-ca5c-4361-9cef-3c53cf890179"><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="e65493bf-4048-4f1b-ae7d-ed1fc8ad6c4c"><ac:plain-text-body><![CDATA[| [none] | Hivemall]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="4bf5ecee-dedf-4ac1-9a7b-2f9ae09f8157"><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="cc2bb8a2-2795-4b87-a768-aded167d9163"><ac:plain-text-body><![CDATA[| [none] | NetBeans]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3e84e4ad-8680-4cb9-ad42-1b82d362ffb4"><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="95327a1f-250e-4970-8bc8-338ce32209de"><ac:plain-text-body><![CDATA[| [none] | OpenWhisk]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3bd8cfb2-8e15-4810-ba52-226a65f2f072"><ac:plain-text-body><![CDATA[| [none] | Pirk]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a3ac77d2-359a-47ba-beca-ebbd4366fd51"><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="37329407-9a13-49ab-b279-56a0b9a3b889"><ac:plain-text-body><![CDATA[| [none] | Streams]></ac:plain-text-body></ac:structured-macro> |
| <ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="69e37426-ad6c-4451-9240-6d603cda4d4d"><ac:plain-text-body><![CDATA[| [none] | Trafodion]></ac:plain-text-body></ac:structured-macro> |

Report content

Incubator PMC report for December 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 currently 65 podlings incubating, another record high. There were seven releases published in the month of November. There were seven new IPMC members elected and one member who left.

* Community

New IPMC members:

- Kasper Sørensen
- Willem Ning Jiang
- Luke Han
- Par Niclas Hedhman
- Jukka Zitting
- Stian Soiland-Reyes
- Kasper Sørensen

People who left the IPMC:

- Sean Busbey

* New Podlings

- RocketMQ
- OpenWhisk
- Weex

* Graduations

The board has motions for the following:

- Beam
- Eagle

The following have already graduated into sub-projects:

- CommonsRDF -> Apache Commons RDF

* Releases

The following releases entered distribution during the month of November:

- Apache Ranger 0.6.2-incubating 2016-11-08
- Apache Streams 0.4-incubating 2016-11-10
- Apache SystemML 0.11.0-incubating 2016-11-12
- Apache CommonsRDF 0.3.0-incubating 2016-11-15
- Apache CarbonData 0.2.0-incubating 2016-11-19
- Apache Metron 0.3.0-incubating 2016-11-28
- Apache Gearpump 0.8.2-incubating 2016-11-28

* IP Clearance

- None

* Legal / Trademarks

- None

* Infrastructure

- None

* Miscellaneous

- Work has begun on updating Apache Incubator documentation to better align with current infrastructure requirements and provide better guidelines to new podlings on how to get setup.

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- Annotator
- AriaTosca
- NetBeans
- OpenWhisk
- Spot

* Not yet ready to graduate

No release:

- MRQL
- Myriad
- Pirk
- SAMOA
- Trafodion
- Wave

Community growth:

- Atlas
- Gearpump
- Hivemall
- HTrace
- Mnemonic
- Pony Mail
- Quickstep
- Singa
- Tephra

* Ready to graduate

- Ranger
- Streams
- Taverna

The Board has motions for the following:

- Beam
- Eagle

* In trouble, further actions required

- log4cxx2
- Omid

Table of Contents

Annotator
AriaTosca
Atlas
Gearpump
Hivemall
HTrace
log4cxx2
Mnemonic

MRQL
Myriad
NetBeans
Omid
OpenWhisk
Pirk
Pony Mail
Quickstep
Ranger
SAMOA
Singa
Spot
Streams
Taverna
Tephra
Trafodion
Wave

Annotator

Annotator provides annotation enabling code for browsers, servers, and humans.

Annotator has been incubating since 2016-08-30.

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

1. Achieving active development cadence
2. Attracting new contributors
3. Establishing presence (web, repo, activity)

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

Beginning work within the ASF has been slow (and rather confusing). Mailing list activity has improved, but still low volume. Development priorities are still being discussed. We do now have repositories for code and the website, but are blocked by needing access to them--the request as been made to the general@incubator list. However, we have successfully gotten Jira and Confluence setup for the project and have a promotional web page ready and waiting once access has been granted. Making the incubation process clearer (and ideally more automated) would have saved a lot of time and confusion. That said, our Champion (Daniel Gruno), has been most helpful in helping us navigate.

How has the community developed since the last report?

Discussions have begun on development priorities. Infrastructure helped us resolve some outstanding project setup needs (Jira, Confluence, and git for the web site and code). Once the site is online, we intend to promote the community more and make concrete plans out of the proposals made on the mailing list so far.

How has the project developed since the last report?

Same as above. Mostly discussion and infrastructure/access setup--some still in progress.

Date of last release:

None.

When were the last committers or PMC members elected?

None since start of incubation.

Signed-off-by:

[*](annotator) Nick Kew
[](annotator) Brian McCallister
[](annotator) Daniel Gruno
[](annotator) Jim Jagielski

Shepherd/Mentor notes:

Drew Farris (shepherd): One mentor active on mailing lists and another's assistance is referenced in the board report. Annotator is still getting off the ground. Observed recent mailing list activity seeking to set a path forward for the project.

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. Grow the community and enroll new committers.
2. Have frequent Release cycles to be compliant with the Apache Way.
3. Publish "How To Contribute" guide at ARIA-TOSCA Website.
4. Update ARIA website according to ASF guideline(Include release process for ARIA-TOSCA Project)

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

There's an INFRA jira we created over a month ago which has yet to be addressed - please see here [INFRA-12733](#) - Ability to create a Sprint Board for AriaTosca WAITING FOR USER

It is not entirely critical for the project's progress, but it could be helpful if we indeed get a sprint board and the other things requested in that issue.

Update: the ticket is now in progress.

How has the community developed since the last report?

- * Open-O is in process for consuming the ARIA ASF package
- * Cloudfiy is in process for consuming the ARIA ASF package as well.
- * Need to work on building the community and accepting external contributions.

How has the project developed since the last report?

- * TOSCA parser migrated into ASF
- * CI on Appveyor, Jenkins
- * Sphinx documentation added to the project
- * CLIs of workflow engine and parser coalesced
- * TOSCA CSAR packager added
- * APIs for workflows, operations
- * Workflow engine task retry support
- * Workflow engine execution cancel support

Date of last release:

No Releases yet

When were the last committers or PMC members elected?

Project is still functioning with the initial set of committers.

Signed-off-by:

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

Shepherd/Mentor notes:

John D. Ament:

The podling is still getting its feet wet. Discussions are starting to happen on the mailing lists rather than in physical offices which is good and will help to bring about a more diverse community. I would also strongly encourage the podling to work with their mentors to get infra changes in to avoid confusion.

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 3 new committers since the last report from different organizations. About 10 new contributors have been added since the last report. The current number of contributors is around 70.
2. The activity on the atlas-dev mailing list stands at around 610 messages per month between September and November [1]
3. People across various organizations are interested in exploring Atlas. We have received comments and suggestions from people who are evaluating Atlas for their governance requirements [5]

How has the project developed since the last report?

1. We are working towards a 0.8-incubating release in the next couple of months. This upcoming release would include major features like structured high level REST APIs, entity creation/updation from Atlas UI, DSL Search improvements, Atlas Hook support for Hive2, Knox SSO support for Atlas, addition of user defined and system attributes to entities
2. An abstraction layer has been added to make Atlas agnostic to the underlying graph database version
3. A total of 186 issues were reported between Sep 1st 2016 and Nov 30th 2016. 122 issues have been resolved in the same time [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 3 new members as committers to the project on 2016-10-20: Madhan Neethiraj, Jeffrey Hagelberg and Vimal Sharma.

Links

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

[2] <https://issues.apache.org/jira/browse/ATLAS-1341?jql=project%20%3D%20ATLAS%20AND%20created%20%3E%3D%20%272016-09-01%27%20and%20created%20%3C%3D%20%272016-11-30%27>

[3] <https://issues.apache.org/jira/browse/ATLAS-1335?jql=project%20%3D%20ATLAS%20AND%20resolved%20%3E%3D%20%272016-09-01%27%20and%20resolved%20%3C%3D%20%272016-11-30%27>

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

[5] <http://www.ibmbigdatahub.com/blog/case-open-metadata>

Signed-off-by:

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

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 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?

Two Apache source releases and one binary release

- Added Window DSL and Triggers.

Continued integration within Apache Beam (Gearpump Runner).

Codes have been merged to runner-gearpump branch on Beam and Beam capabilities are being built into Gearpump.

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

4 PRs have been committed to akka and an akka-stream feature branch has been created on Gearpump.

- 75 issues created and 55 resolved

- Manu presented Gearpump at Strata Beijing 2016

- Kam presented Gearpump Materializer at Reactive Summit 2016

- Huafeng and Karol co-presented Gearpump at ApacheCon Big Data Europe 2016

Date of last release:

2016-11-30

When were the last committers or PMC members elected?

No new committers or PMC members elected yet.

Signed-off-by:

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

Hivemall

Hivemall is a library for machine learning implemented as Hive UDFs/UDAFs/UDTFs.

Hivemall has been incubating since 2016-09-13.

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

1. Create the first Apache release
2. Community growth (committers and users)
3. Migrate the development community from Github to ASF infra

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

- * All SGAs and I-CLA has been submitted and recorded
- * Daniel Dai is elected as a new Mentor on Nov 22
<https://lists.apache.org/thread.html/48a61eb44188d3f204e7c41ff7c102128e9325eed029ab41a23ff3a9@%3Cdev.hivemall.apache.org%3E>

How has the community developed since the last report?

- * Presented a talk in Cloudera World Tokyo on Nov 8.
- * Created Twitter Account: @ApacheHivemall

How has the project developed since the last report?

- * Collected and submitted all SGAs and I-CLA required for IP clearance
<https://issues.apache.org/jira/browse/HIVEMALL-9>
- * Finished updating License headers and NOTICE
<https://issues.apache.org/jira/browse/HIVEMALL-12>
- * Finished the initial code dump on Dec 1 and the development moved to ASF infra
<https://issues.apache.org/jira/browse/HIVEMALL-6>
<https://github.com/apache/incubator-hivemall>
- * Created 5 JIRA Issues and resolved 4 issues in the last 30 days
<https://issues.apache.org/jira/browse/HIVEMALL/>

Date of last release:

No release yet (planning the first Apache release in Q1, 2017)

When were the last committers or PMC members elected?

- * Daniel Dai (Mentor) on Nov 22

Signed-off-by:

[](hivemall) Reynold Xin
[](hivemall) Markus Weimer
[](hivemall) Xiangrui Meng
[X](hivemall) Daniel Dai

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 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?

Mailing lists have been quiet. Little activity during this period.

How has the project developed since the last report?

Released htrace-4.2.0 and added a new committer.

Date of last release:

htrace-4.2.0 October 6th, 2016

When were the last committers or PMC members elected?

+Mike Drob (mdrob) October 3rd, 2016 - Committer+PMC Member

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

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.

Currently it seems that there's only one official project member with commit access left and even though some contributor worked on fixing the release process, that project member was unable to invest the necessary time to create a release yet.

For that reason, the left member decided to not even care anymore about a release and published this decision on the dev mailing list of the project. While the member is still willing to provide some support and would regret a move of the project to attic, creating a release has no priority nor is seen as strictly necessary and will

therefore most likely always prioritize lower than other tasks.

How has the community developed since the last report?

3 support questions, no new project members. There was a lengthy discussion about switching to C++ standard smart pointers with possibly breaking the API, the influence of such an approach on a release etc. Some bugs have been created to coordinate that work.

Robert Middleton started to move the project to smart pointers using boost or std::shared_ptr in a private GitHub repo. He's actually changing more code than the one left official project member.

https://github.com/rm5248/log4cxx-testing/tree/smart_pointers

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:

Ralph Goers:

The Log4j PMC will need to have a discussion about the state of this project. Since Log4cxx2 was always intended to be a subproject of the Logging project moving it to the attic may or may not be the right move. Instead, it might just become a "dormant" subproject.

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, SerDe-less and native-direct computing high performance applications.

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

1. Apache releases: we have release version 0.3.0-incubating
2. Improved dev. docker to facilitate development.
3. provided information to potential customers e.g. bank user, other project's Apache committers

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

No.

How has the community developed since the last report?

Our last report was in April. Since then

- * Two more developers are going to join our community.
- * We need to encourage more discussion on the dev list in the coming months

How has the project developed since the last report?

- * Added durable native sorting service (n-dim linked data) with testcases
- * Removed the 3rd Maven repository and use Maven central repository

only

- * DOM and DNCS could apply to volatile memory service except fixed k-v store
- * Added a bench workload to verify the benefits from DNCS with scripts, generator
- * Added a sys-vmem service that makes it possible to utilize system memory for DNCS in an unified way
- * Added a script to bundle signed portable artifacts for maven repository submitting
- * and other bugfixes, features, improvements

Date of last release:

2016-10-21

When were the last committers or PMC members elected?

None

Signed-off-by:

[X](mnemonic) Patrick Hunt
[](mnemonic) Andrew Purtell
[](mnemonic) James Taylor
[X](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 was one new developer who developed the new evaluation mode for MRQL that runs on top of Apache Storm. There were no new committers since our last report.

How has the project developed since the last report?

A new evaluation mode for MRQL queries was developed based on Apache Storm (using the Trident API). MRQL can now process stream queries on a single stream using Storm, but will be soon be extended to support multiple streams.

Date of last release:

2016-03-02

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

[X](mrql) Alan Cabrera

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

Shepherd/Mentor notes:

Alan Cabrera: Not a huge amount of activity but at least there's a small consistent trickle. I think that the fact that someone was interested in adding an Apache Storm adapter speaks somewhat to the viability of the project.

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 progress), 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 had 73 messages since the last report. 2 new correspondents.
- * Bi-weekly dev syncs with 2-3 participants. Some cancelled due to lack of attendance. Minutes at <http://s.apache.org/8kF>
- * Myriad presented at HadoopWorld/Strata NYC.

How has the project developed since the last report?

- * Myriad 0.3.0 in progress.
- * 1 commit/PR merged and 2 JIRAs resolved.
- * Myriad submitted to DC/OS Universe:
<https://github.com/mesosphere/universe/pull/841>

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
[x](myriad) Danese Cooper - <https://lists.apache.org/thread.html>
/f7635ace681984cacfca6692a35221ede0000bcc70f4fcfc990433e9@%3Cdev.myriad.apache.org%3E
[x](myriad) Ted Dunning
[](myriad) Luciano Resende

Shepherd/Mentor notes:

johnndament: I am concerned about lack of engagement in this podling. Mentors need to be a bit more engaged and repsonsive.

NetBeans

NetBeans is a development environment, tooling platform and application framework.

NetBeans has been incubating since 2016-10-01.

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

1. Licensing, i.e., identifying and solving GPL-related code.
2. Coming up with a process of contributing code that makes sense to everyone.
3. Working on roadmaps, features, and plans together as a 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?

Discussed CCLA/SGA on the dev mailing list.

Discussed existing mailing lists migration plan to Apache:
<https://cwiki.apache.org/confluence/display/NETBEANS/Mailing+lists+migration>

How has the project developed since the last report?

CCLA and SGA still in the process of being approved and signed by Oracle, which will also apply the copyright headers as instructed on the dev mailing list.

Started initial code cleanup analysis
<https://cwiki.apache.org/confluence/display/NETBEANS/Overview%3A+NetBeans+Structure>

Oracle has already agreed to relicense the Graal.js parser under the Apache license. Resolved in LEGAL-279 the nb-javac dependency license.

Almost finished new static website for Apache NetBeans.

Experimented Mercurial to Git repository migration with a focus on repository size. Solution found, size not an issue anymore.

Date of last release:

No releases yet.

When were the last committers or PMC members elected?

No one has been elected so far.

Signed-off-by:

[X](netbeans) Ate Douma
[X](netbeans) Bertrand Delacretaz
[X](netbeans) Emmanuel Lecharny
[X](netbeans) Daniel Gruno
[](netbeans) Jim Jagielski
[](netbeans) Mark Struberg

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?

Initiated contact with Apache Phoenix community to integrate Omid.
 Blog entries published in Yahoo Engineering and Hadoop blogs.
 Article about Apache Omid incubation published in Datanami.
 Talk proposal on Omid sent to DataWorks Summit EU 2017
 Quarter Stats (from: 2016-09-01 to: 2016-11-30):

| Metric | counts |
|------------------------------------|--------|
| # of msgs in dev list | 20 |
| Active Contributors (incl mentors) | 6 |
| Jira New Issues | 2 |
| Resolved Issues | 1 |
| Pull Requests merged | 0 |
| Pull Requests proposed | 0 |

How has the project developed since the last report?

In the process of completing its second release under in the Apache incubator.

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
- [x](omid) Flavio P. Junqueira
- [](omid) Thejas Nair
- [](omid) James Taylor

Shepherd/Mentor notes:

Alan Gates:

I'm concerned about this podling. The last commit was on August 15. Discounting broadcast and reminders, there have been no emails on the dev list since October, and there were only 3 that month. I've sent email on the dev list asking if people are still actively involved.

Flavio Junqueira:

I second Alan's concern, there has been very little activity in the last period for this podling.

 OpenWhisk

OpenWhisk is an open source, distributed Serverless computing platform able to execute application logic (Actions) in response to events

(Triggers) from external sources (Feeds) or HTTP requests governed by conditional logic (Rules). It provides a programming environment supported by a REST API-based Command Line Interface (CLI) along with tooling to support packaging and catalog services.

OpenWhisk has been incubating since 2016-11-23.

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

1. The project just entered the Incubator and infrastructure is in progress to be setup.
- 2.
- 3.

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

-

How has the community developed since the last report?

-

How has the project developed since the last report?

-

Date of last release:

None yet.

When were the last committers or PMC members elected?

-

Signed-off-by:

[X](openwhisk) Felix Meschberger
[X](openwhisk) Isabel Drost-Fromm
[X](openwhisk) Sergio Fernández

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 attracted a couple of new contributors since the last report. The mailing list activity, pull requests, and page views to the Apache Pirk website (via Google Analytics) have also increased since the last report. An 'Introducing Pirk' talk was presented at ApacheCon Europe on November 15, 2016.

How has the project developed since the last report?

The project is in the process of completing its second release, 0.2.0-incubating. The project has also undertaken a refactor of some of the packages and drivers and added additional user/developer documentation on the website.

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:

[x](pirk) Billie Rinaldi
[x](pirk) Joe Witt
[x](pirk) Josh Elser
[X](pirk) Suneel Marthi
[](pirk) Tim Ellison

Shepherd/Mentor notes:

Josh Elser:

Dev traffic appears to have slowed down a little in the past month, but nothing really to be concerned about yet, IMO.

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. Grow awareness of the project
3. Master releases and project management

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

None that we know of.

How has the community developed since the last report?

One new committer (Sebastian Bazley) was invited to the project.

How has the project developed since the last report?

Email traffic has been roughly the same, 75-ish mails per month to the lists. Quite a few bugs have been worked on.

Date of last release:

2016-08-02

When were the last committers or PMC members elected?

- Sebastian Bazley was elected committer on October 9th, 2016.

Signed-off-by:

[](ponymail) Andrew Bayer

[X](ponymail) John D. Ament

Quickstep

Quickstep is a high-performance database engine.

Quickstep has been incubating since 2016-03-29.

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

1. Make the first release of the system.
2. Write blogs/papers on the release.
3. Help foster a user community to try out 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?

The developer community is actively working on having more discussions on line both on code and larger project related work. In other words, embracing the Apache way both mechanistically and philosophically.

How has the project developed since the last report?

The code continues to make progress and is close to the point where a release can be created.

Date of last release:

None so far.

When were the last committers or PMC members elected?

Tarun Bansal was elected as a committer on Oct 27, 2016.

Signed-off-by:

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

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. Working with Mentors on all action/preparation for Graduation
2. Re-organize documentation for easier access to end-users, new users and contributors.
3. Expanding user/dev 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?

1. Mailing list activity since last report (September-01-2016 to November-30-2016):
@dev

@user 96
@commit 140

2. Issues (JIRAs) created/resolved since last report (after September-01-2016 before November-30-2016):
Created: 71
Resolved: 48

3. Adding more community members to involve in Apache Ranger

How has the project developed since the last report?

1. Released one minor release:
0.6.2 released on 11/08/2016
2. Members have been working on a next major release of 0.7.0
 - a. Working on the roadmap features to be delivered soon.
3. Members are actively involved in Graduation process in the dev community
 - * "Graduation Maturity Assessment" is being discussed in the community
 - * Elected VP/PMC Chair for Ranger: Selvamohan Neethiraj

Date of last release:

November-08-2016

When were the last committers or PMC members elected?

1. Velmurugan Periasamy has been added as PMC member on September-05-2016

Signed-off-by:

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

Shepherd/Mentor notes:

Alan Gates:

The community is actively working on graduating. They have been discussing a draft proposal on the dev list and are working out a few final details before taking it to the Incubator PMC. Based on their self assessment against the maturity model, community growth, and several Apache releases I believe they are ready to graduate.

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 Flink, Apache Storm, Apache Apex 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 developer base
2. Grow the user base
3. Add some more ML techniques

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 (September 2016 - November 2016):

* @dev: 87 messages

Jira issues backlog (September 2016 - November 2016):

* Created: 0

* Resolved: 1

- We manage to resolve the pending issue regarding the integration of Apache Apex with Samoa.
- Also, we made some testing progress with Apache Gearpump.
- We also had a presentation of Apache Samoa and its integration with Apache Apex in Apache BigData Europe 2016
- We also published a scientific paper using Apache Samoa and one of its machine learning techniques (Vertical Hoeffding Tree) (to be presented in IEEE BigData Conference in December 2016)

How has the project developed since the last report?

- We have performed a new release with several bug fixes and new features (0.4.0)
- We are looking into new ML techniques for development.

Date of last release:

2016-09-30

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:

Very low activity on this project over the last few months, with last commit showing as Sept 22 and very few commits since last March. A big need here beyond activity is to grow the community. There is a one individual who made several contributions last spring; I sent an email to the mailing list asking if it made sense to invite him to become a committer.

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. Add more documentations and tutorials, to make the installation and use of SINGA easier. We have done some work as described below in the answer of 'how has the project developed since the last report'.
2. Grow the community. We added one more committer after the last report and are getting more user feedbacks. We want to further grow the community and to graduate in early 2017.

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 68, 125 and 61 emails from dev@ list in September 2016, October 2016 and November 2016 respectively.

There are 100 more commits since the last report.

One new committer was added into SINGA committer list.

How has the project developed since the last report?

We released the v1.0 shortly after the last report, which include major changes in terms of the programming model and other aspects:

- + Tensor abstraction for supporting more machine learning models.
- + Device abstraction for running on different hardware devices, including CPU, (Nvidia/AMD) GPU and FPGA (to be tested in later versions).
- + Replace GNU autotool with cmake for compilation.
- + Support Mac OS.
- + Improve Python binding, including installation and programming
- + More deep learning models, including VGG and ResNet
- + Replace website building tool from maven to Sphinx and updated the documentations.
- + Integrate Travis-CI

We also added python notebooks and windows support, and are working on the model zoo, python whl files and Docker images to make it easier for users to get started with SINGA.

We are focusing on our niches, including distributed training at the system level and health-care models at the application level.

Date of last release:

2016-09-09

When were the last committers or PMC members elected?

2016-10-21

Signed-off-by:

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

Spot

Apache Spot is a solution stack that provides the capability to ingest network related telemetry (network flows, domain name service information and proxy server logs) and provide unsupervised machine learning capabilities to identify suspicious activity. The information is organized and presented using operational analytics so that a security analyst can investigate the most suspicious connections. Apache Spot is built on an open data model using Apache Spark and Apache Hadoop.

Spot has been incubating since 2016-09-23.

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

1. Move infrastructure and development to ASF (code, issues, mailing list, ...)
2. Build diverse community
3. Demonstrate ability to create releases

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

Please elaborate on the process for code scanning and IP reviews.

How has the community developed since the last report?

Community interaction has ramped velocity, as new adopters of the project have been posting questions through Twitter and Slack. We have added approximately 20% to our Slack channel and added several contributors through the ICLA.

How has the project developed since the last report?

We have moved our repo to ASF as of December 8, 2016. We have removed our LDA-C dependencies by switching to Spark LDA, eliminating licensing issues on this particular branch.

Date of last release:
N/A

When were the last committers or PMC members elected?
N/A

Signed-off-by:

[](spot) Jarek Jarcec Cecho
[](spot) Brock Noland
[X](spot) Andrei Savu
[X](spot) Uma Maheswara Rao G

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?

The community has begun an internal assessment using the Apache Maturity Model to identify gaps that should be addressed before graduation.

How has the community developed since the last report?

Dev List

117 emails sent by 16 people, divided into 32 topics.

Web Page

498 Sessions (up ~150%), 341 Users (up ~35%), 1178 Pageviews (up 10%).

How has the project developed since the last report?

Source Control

<https://github.com/apache/incubator-streams>

Excluding merges, 2 authors have pushed 17 commits to master and 17 commits to all branches. On master, 816 files have changed and there have been 36,280 additions and 38,567 deletions.

<https://github.com/apache/incubator-streams-examples>

Excluding merges, 2 authors have pushed 2 commits to master and 2 commits to all branches. On master, 8 files have changed and there have been 323 additions and 301 deletions.

Steve Blackmon presented Apache Streams at Shark Tank, ApacheCon Europe, Seville, Spain on Nov 18, 2016.

Team is presently working on 0.4.1-incubating release scheduled for the week of Dec 5, 2016.

Work has been scoped out for an upcoming 0.5-incubating release scheduled for later this year.

JIRA

52 issues closed with 0.4-incubating release
12 issues closed after 0.4-incubating release
16 new issues opened in November.

Date of last release:

2016-11-11 : 0.4-incubating release

When were the last committers or PMC members elected?

2016-10-27: Joey Frazee elected as committer / PPMC member
2016-09-28: Suneel Marthi elected as mentor / PPMC member

Signed-off-by:

[X](streams) Matt Franklin
[X](streams) Ate Douma
[X](streams) Suneel Marthi

Shepherd/Mentor notes:

Suneel Marthi:

There's been lot of activity on the project since Sep 2016 and the project has had 2 releases since then. The project's on track for TLP graduation.

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. IP/licence review of unreleased repositories
3. Graduate!

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

Preparing to graduate.

Stian added as new mentor.

How has the community developed since the last report?

Added 3 new committers who were introduced via the Google Summer of Code. 2 of them did not directly work on Apache Taverna through GSoC but this was where they heard about the project.

How has the project developed since the last report?

The Taverna Mobile Android application is nearing release standard with all licence checks passed.

Most of the codebases under the Taverna banner are also nearing release standard with an extensive licence review currently underway.

Preparing maturity report, almost ready for graduation vote.

Email stats dev@taverna:

2016-09: 190 emails
2016-10: 69 emails
2016-11: 78 emails

Email stats users@taverna:

2016-09: 0 emails
2016-10: 3 emails
2016-11: 16 emails

Jira stats:

2016-09 (6 created, 0 resolved)
2016-10 (3 created, 7 resolved)
2016-11 (0 created, 15 resolved)

Git stats:

2016-09: 158 commits
2016-10: 165 commits
2016-11: 34 commits

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?

2016-10-20

Signed-off-by:

[x](taverna) Andy Seaborne
[](taverna) Daniel J Debrunner
[x](taverna) Stian Soiland-Reyes
[](taverna) Suresh Srinivas
[](taverna) Suresh Marru
[](taverna) Marlon Pierce

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?

- 18 subscribers in dev mailing list
- 21 new JIRAs filed since the last report
- 4 external contributors submitted patches since the last report

How has the project developed since the last report?

- Working on 0.10.0 release
- Released 0.9.0

Date of the last release:

2016-09-18

When were the last committers or PMC members elected?

- None since coming to incubation

Signed-off-by:

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

Shepherd/Mentor notes:

- Drew Farris (Shepherd): Two mentors active on mailing lists. Project is building towards a healthy community.

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.

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: 715 messages in the codereview forum, 728@commits, 257@dev, 1155@issues and 61@user.
- * We are working towards our next release, 2.1.
- * Gunnar Tapper has volunteered to be the Release Manager.
- * Identified items for the roadmap of our next release:
 - TRAFODION-2067 Support CDH 5.7
 - TRAFODION-1988 Better java exception handling in the java layer of Trafodion
 - TRAFODION-2117 Add support of SQL extension 'EXCEPT'
 - TRAFODION-2060 Add support for native tinyint datatype
 - TRAFODION-1931 ODB tool support on Windows platform
 - TRAFODION-1861 Support Trafodion running on CentOS 7
 - TRAFODION-1673 Implement the WITH clause in Trafodion SQL for simple use cases
 - TRAFODION-2047 add SQL syntax support to INTERSECT
 - TRAFODION-2127 enhance Trafodion implementation of WITH clause
- * The community continues to be active in China, with some communications happening outside the Apache dev lists, we had a discussion on the Incubator general list about that.
- * Community members speaking at public events:
 - Liu Ming had a Trafodion open-class for undergraduates at FuDan University, Shanghai on Nov 10.
 - Jin Jian had a Trafodion open-class for post-graduate students at FuDan University, Shanghai on Nov 24

How has the project developed since the last report?

- * 253 commits from 27 contributors.
- * 182 JIRAs filed and 113 resolved Sep 1 - Nov 30.
- * Many of the roadmap items listed above have been implemented, plus many more enhancements not listed above.

Date of last release:

2016-07-07 2.0.1

When were the last committers or PMC members elected?

in August

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?

The community has several new/re-activated contributors and participants across different disciplines since last report. The most

significant developments for our community have come about since September.

After an open community conference call on Sept 28th, a joint effort between Wave and SwellRT development was initiated. SwellRT API has been tried out by some members of the community and a branch containing the SwellRT code is now in the repository. Discussions are still taking place to determine the full extent of the integration as the projects share a common goal but have side plans as well.

SwellRT is a fork of Wave resolving some complexity issues of the Wave code that have prevented community growth and adoption issues.

How has the project developed since the last report?

Some members of Wave community have tried SwellRT API and some common tasks have been performed, such as:

- Cleaned up some obsolete dependencies.
- New branch created on wave repo containing a mirror of SwellRT code.
- Updated major dependencies and fixed minor bugs, some taken from SwellRT.

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:

SwellRT appears to be a useful addition to Wave - it was originally a fork of Wave that has been improved/simplified. The Wave community still needs to handle IP Clearance for the SwellRT codebase.