

# September 2015

/\ \*FINAL\* (!)

## Incubator PMC report for September 2015

### Timeline

Wed September 02	Podling reports due by end of day
Sun September 06	Shepherd reviews due by end of day
Sun September 06	Summary due by end of day
Tue September 08	Mentor signoff due by end of day
Wed September 09	Report submitted to Board
Wed September 16	Board meeting

### Shepherd Assignments

Andrei Savu	Zeppelin		
Drew Farris	Climate Model Diagnostic Analyzer		
John Ament	Kalumet		
John Ament	Streams		
John Ament	ODF Toolkit		
Justin Mclean	Geode		
Matthew Franklin	Droids		
Matthew Franklin	MRQL		
Matthew Franklin	Myriad		
P. Taylor Goetz	Cotton		
P. Taylor Goetz	Ranger		
P. Taylor Goetz	REEF		
Raphael Bircher	Corinthia		
Raphael Bircher	Singa		
Ross Gardler	Wave		
Suresh Marru	Calcite		
Suresh Marru	SAMOA		
Suresh Marru	log4cxx2		
Timothy Chen	Brooklyn		
Timothy Chen	Taverna		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="47080b04-913b-43a7-bfd6-296b8b1fbadd"><ac:plain-text-body><![CDATA[	[none]	Apex	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7c30e487-1abb-422d-a764-076af12cf8b5"><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="01f675a9-c09a-4542-9219-888abc46d398"><ac:plain-text-body><![CDATA[	[none]	CommonsRDF	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="dc318910-7360-47b6-becb-8552cc7251f6"><ac:plain-text-body><![CDATA[	[none]	FreeMarker	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7d41def2-b8d5-45f3-a8f7-9d4b4c7a4f90"><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="9765e964-aa1f-4013-81cb-12ab3e42db78"><ac:plain-text-body><![CDATA[	[none]	Johnzon	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="896768c1-4d15-46db-88e2-0e82a1f4f9e5"><ac:plain-text-body><![CDATA[	[none]	Ripple	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c20c1a94-dd26-480f-bf32-d6c369968dfa"><ac:plain-text-body><![CDATA[	[none]	Sentry	]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3429c1e4-428e-485d-be41-6cc7a5a54ea9"><ac:plain-text-body><![CDATA[	[none]	Slider	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e62816e6-d806-48d3-b56f-244255dbb29c"><ac:plain-text-body><![CDATA[	[none]	Trafodion	]]></ac:plain-text-body></ac:structured-macro>

## Report content

Incubator PMC report for September 2015

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 41 podlings currently undergoing incubation.

### \* Community

New IPMC members:

- Sharad Agarwal
- Venkatesh Seeratham

People who left the IPMC:

- Jan Iversen

### \* New Podlings

- Apex
- HAWQ
- HORN

### \* Graduations

- Usergrid graduated last month, but wasn't on the report due to timing of their graduation vote.

### \* Releases

The following releases were made during the month of August:

- 2015-08-10 taverna-language-0.15.0-incubating
- 2015-08-10 taverna-parent-1-incubating
- 2015-08-15 apache-reef-0.12.0-incubating
- 2015-08-17 apache-johnzon-0.9-incubating
- 2015-08-22 tamaya-distribution-0.1-incubating
- 2015-08-28 incubator-corinthia\_release\_0.1
- 2015-08-30 ripple-emulator-0.9.32-incubating
- 2015-08-31 apache-johnzon-0.9.1-incubating
- 2015-08-31 apache-calcite-1.4.0-incubating

### \* IP Clearance

- Adobe contributed the Sling HApi Tools, which provide a set of Hypermedia API tools for Sling.

### \* Legal / Trademarks

- The copyright section of the IP Clearance template was modified.

### \* Infrastructure

- The automated Marvin report reminders fired eight days late this month. The manual Marvin report reminders went instead.

### \* Miscellaneous

- The Corinthia podling experienced a crisis. Many PPMC members, including the most active core contributors, announced their resignations and intent to carry on the project outside Apache.

- A long conversation took place on general@incubator and other ASF lists regarding trademark enforcement for binary distributions which are not directly compiled from unmodified Apache release source code, but instead from modified releases or unreleased code.

- The Kalumet podling is holding a vote to retire.

- The Droids podling is considering retirement.

\* Credits

- Report Manager: Marvin Humphrey

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

\* Still getting started at the Incubator

- Apex

\* Not yet ready to graduate

No release:

- FreeMarker
- Geode
- Singa
- Trafodion

Community growth:

- Atlas
- CommonsRDF
- HTrace
- Johnzon
- log4cxx2
- MRQL
- ODF Toolkit
- Ranger
- SAMOA
- Sentry
- Slider
- Streams
- Taverna
- Zeppelin

\* Ready to graduate

- Brooklyn
- Calcite
- REEF

\* Considering retirement

- Kalumet

\* In crisis

- Corinthia

\* Did not report, expected next month

- CMDA
- Cotton
- Droids
- Myriad
- Ripple

\* Missing Mentor sign-off, expected next month

- Wave

-----  
Table of Contents

Apex  
Atlas  
Brooklyn  
Calcite  
CommonsRDF  
Corinthia  
FreeMarker  
Geode  
HTrace  
Johnzon  
Kalumet  
log4cxx2  
MRQL  
ODF Toolkit  
Ranger  
REEF  
SAMOA  
Sentry  
Singa  
Slider  
Streams  
Taverna  
Trafodion  
Zeppelin

-----  
Apex

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

Apex has been incubating since 2015-08-17.

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

1. Complete software donation
2. Create an incubating web site (<<http://apex.incubator.apache.org/>>)
3. Release Apex-Core, and Apex-Malhar at least twice

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

None. Thanks for continual support.

How has the community developed since the last report?

This is the first report. Community incubated few weeks ago.

How has the project developed since the last report?

This is the first report. The community is now using Apache forums, and other tools to collaborate.

Date of last release:

No release under ASF-Incubation code base as yet. Apex-Core, Apex-Malhar was last launched on July end in pre Apache days.

When were the last committers or PMC members elected?

Committers and Mentors came about via incubation proposal

Signed-off-by:

[X](apex) Chris Nauroth  
[X](apex) Alan Gates

[X](apex) Hitesh Shah  
[X](apex) Justin Mclean  
[X](apex) P. Taylor Goetz  
[ ](apex) Ted Dunning

Shepherd/Mentor notes:

Chris Nauroth (cnauroth):

We are already seeing substantial public discussion on the dev mailing list, both in terms of number of messages and the content of those messages. Collaboration is going well.

-----  
Atlas

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

Atlas has been incubating since 2015-05-05.

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

1. Podling name search is pending
2. Grow the community beyond Jira and ML participation

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 added a few contributors (about 12) to the project
- \* A vote is in progress in the private ML for adding a committer
- \* Atlas is being presented at the security meet up at DataGuise on 9/10/2015 <http://s.apache.org/dA7>

How has the project developed since the last report?

- \* We made the first release of Atlas 0.5.0-incubating on 7/10/2015
- \* A total of 130 issues were created and about 50 of them are resolved

Date of last release:

2015-07-15 - 0.5.0-incubating

When were the last committers or PMC members elected?

- \* A vote is in progress scheduled to close 9/3/2015.

Signed-off-by:

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

-----  
Brooklyn

Brooklyn is a framework for modelling, monitoring, and managing applications through autonomic blueprints.

Brooklyn has been incubating since 2014-05-01.

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

1. One more release (under way)

2. Podling Name Search (under way)
- 3.

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

The Brooklyn community is active, growing, increasingly diverse and doing very well. Chip (champion) suggested a couple of weeks back that the community should start discussing graduation.

How has the community developed since the last report?

The 0.7.0 release got out with only minor comments about a month ago. A 0.8.0 version is being prepared that addresses the comments and should be out shortly.

There are 2 new committers and PPMC members that joined, plus a 3rd one currently under vote.

How has the project developed since the last report?

Very active, growing community.  
A couple of presentations at LinuxCon in Seattle in Aug.  
Talks scheduled for ApacheCon, Budapest in Oct.

Date of last release:

2015-07-27 (0.7.0)

When were the last committers or PMC members elected?

2015-08-06 (svet)  
(one more under vote)

Signed-off-by:

[ ](brooklyn) Matt Hogstrom  
[ ](brooklyn) Alex Karasulu  
[X](brooklyn) David Nalley  
[ ](brooklyn) Marcel Offermans  
[X](brooklyn) Jean-Baptiste Onofré  
[X](brooklyn) Olivier Lamy  
[X](brooklyn) Chip Childers  
[ ](brooklyn) Andrei Savu  
[ ](brooklyn) Joe Brockmeier  
[ ](brooklyn) Jim Jagielski

-----  
Calcite

Calcite is a highly customizable engine for parsing and planning queries on data in a wide variety of formats. It allows database-like access, and in particular a SQL interface and advanced query optimization, for data not residing in a traditional database. (Renamed from Optiq on 2014-09-30.)

Calcite has been incubating since 2014-05-19.

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

1. Hold a graduation vote
- 2.
- 3.

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?

In August, there were 122 messages on dev@, 28 commits, 46 issues

created, 22 issues resolved by 8 contributors.

We released 1.4, the first release for which Julian Hyde was not release manager.

Collaborations with Apache projects (Phoenix, Samza, Kylin, Hive, Drill) continue.

We are discussing graduation, and will move to a vote shortly.

How has the project developed since the last report?

Released 1.4.

Date of last release:

2015-09-02 apache-calcite-1.4.0-incubating

When were the last committers or PMC members elected?

2015-04-24 5 new committers  
2014-12-31 1 new PPMC member

Signed-off-by:

[ ](calcite) Ted Dunning  
[X](calcite) Alan Gates  
[ ](calcite) Steven Noels

-----  
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. Grow the Commons RDF community base
2. Continue development towards 0.2.0-incubating release
3. Work with key target communities (initially within Apache) to ensure Commons RDF meets potential user requirements and expectations

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?

A new community member has mentioned development of a branch which may have the potential to drive the project in a different direction. Commons RDF IPMC members are continuing conversation on the topic. Additionally there is recent discussion of how Commons RDF can be used within Apache Jena. Jena is one of the key consumer communities Commons RDF aims to target within the scope of addressing #3 above.

How has the project developed since the last report?

Development has been reasonably slow post 0.1.0-incubating release.

Date of last release:

2015-05-07

When were the last committers or PMC members elected?

N/A no PPMC or Committers have been elected yet.

Signed-off-by:

[X](commonsrdf) Rob Vesse  
[X](commonsrdf) John D Ament  
[ ](commonsrdf) Gary Gregory  
[X](commonsrdf) Lewis J. McGibbney

Shepherd/Mentor notes:

Rob Vesse (rvesse):

Activity dropped to almost zero after first release but has recently started to pick up again.

-----  
Corinthia

Corinthia is a toolkit/application for converting between and editing common office file formats, with an initial focus on word processing. It is designed to cater for multiple classes of platforms - desktop, web, and mobile - and relies heavily on web technologies such as HTML, CSS, and JavaScript for representing and manipulating documents. The toolkit is small, portable, and flexible, with minimal dependencies. The target audience is developers wishing to include office viewing, conversion, and editing functionality into their applications.

Corinthia has been incubating since 2014-12-08.

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

1. Solve the community crisis and avoid a break
- 2.
- 3.

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

The community is experiencing a major breakdown. The mentors have accessed the situation including several attempts to stop the breakdown without success. IPMC have been informed with a suggestion which are being considered.

It is to be expected that several core developers will leave the project in a very short term unless a solution is found.

How has the community developed since the last report?

We are proud that we have managed to make our first release.

How has the project developed since the last report?

Added support for 32/64bit windows.

Made the setup of a build environment a lot easier, a PPMC provide all third party libraries prebuilt.

A long discussion about licenses lead to major difficulties in deciding whether and how to proceed with one of the project's two main components, an end-user editing app. As a result this development is being redirected.

Our buildbot setup is slowly taking shape, this is important since it will allow a nightly build as well as a nightly test.

Date of last release:

2015-08-28 Release 0.1

When were the last committers or PMC members elected?

2015-06-12 Ian Cunningham  
2015-08-31 Gavin McDonald

Signed-off-by:

[ ](corinthia) Daniel Gruno  
[x](corinthia) Jan Iversen  
[ ](corinthia) Dave Fischer

-----  
FreeMarker

FreeMarker is a template engine, i.e. a generic tool to generate text output based on templates. FreeMarker is implemented in Java as a class library for programmers.

FreeMarker has been incubating since 2015-07-01.

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

1. IP clearance
2. Move to ASF infrastructure, especially the code repository
3. Finding contributors

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

None

How has the community developed since the last report?

We are still in early stage of incubation (like source code was only imported recently), so not much is happening inside ASF yet. One new contributor has found us through ASF.

How has the project developed since the last report?

GitHub import was done, we are starting the IP clearance process. (On GitHub the project was highly active.)

Date of last release:

There was no release from the Incubator yet

When were the last committers or PMC members elected?

We only have the initial committers and initial PMC members yet.

Signed-off-by:

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

-----  
Geode

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

Geode has been incubating since 2015-04-27.

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

1. Have our first Apache (incubating) release (currently blocked on us solving JGroups licensing issues) GEODE-77.
2. Expanding the community to include contributors and committers outside of Pivotal.

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

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

None.

How has the community developed since the last report?

- \* Geode now has a Zeppelin interpreter.
- \* In August 120 issues were created and 59 resolved, bringing the total to 304 created and 136 resolved.
- \* There were 6 pull requests to github in the last 30 days, for a total of 18 pull requests.
- \* There were 322 messages on the dev list and 63 messages on the user list.
- \* There are now 117 subscribers on the user and 127 on the dev list.
- \* Events and Conferences in August:
  - \* QCon Rio: (30 attended)
    - \* In-Memory Analytics and Machine Learning in practice with Spark, Geode, Spring XD and Docker <http://s.apache.org/9ln>
  - \* Geode Clubhouse: Roundtable (<https://www.youtube.com/watch?v=qaOB3d-pvTI>) where 10 issues were discussed in the community. (about 20 participated)
- \* September:
  - \* Geode clubhouse: Building Effective Apache Geode Applications with Spring Data GemFire (Sept 8th)
  - \* Three talks at "Apache: Big Data Europe"
    - An introduction to Apache Geode (incubating)
    - Building a highly scalable open-source Real-time Streaming Analytics system using Spark SQL, Apache Geode (incubating), SpringXD and Apache Zeppelin (incubating)
    - Implementing a Highly Scalable In-Memory Stock Prediction System with Apache Geode (incubating), R and Spring XD

How has the project developed since the last report?

- \* Great progress has been made on JGroups issue, where many unit and integration tests now pass on the feature/GEODE-77 branch.
- \* Good progress has also been made on stabilizing CI, by fixing various intermittent failures in integration tests.

Date of last release:

NA - Just nightly builds.

When were the last committers or PMC members elected?

7 members missed in the initial PMC list were added by vote on June 15, 2015.

Signed-off-by:

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

Shepherd/Mentor notes:

Konstantin Boudnik (cos):

Jan Iversen has resigned from the PPMC today, citing personal reasons.

-----  
HTrace

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

HTrace has been incubating since 2014-11.

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

1. Continue to grow the HTrace community
2. Continue to develop and release stable HTrace incubating artifacts
3. Continue to explore the integration of the HTrace framework into other Apache products

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?

38 subscribers to dev@ (up from 32) and about 55 messages since last report (down from 120 on previous period)

At least three new contributors have showed up in this period.

How has the project developed since the last report?

About 47 issues resolved since last report contributed by 6 different contributors.

New GUI

Working on new htrace-4.0 API with fixes to make htrace easier to host with a release inside the next few weeks.

Date of last release:

None for this period.  
Last release was htrace-3.2.0 Jun 2, 2015

When were the last committers or PMC members elected?

Committer: Abraham Elmahrek - 2.11.15

Signed-off-by:

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

-----  
Johnzon

Implementation of JSR-353 JavaTM API for JSON Processing

Johnzon has been incubating since 2014-06-09.

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

1. Expanding the community, increase dev list activity and adding new committers/pmc members
2. Complete the project incubation status page
3. Add and improve documentation

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

aware of?

\* Nothing

How has the community developed since the last report?

There are new people on the dev list and the project got contributions from non-committers. On the dev list there are approx. 30 msgs (on average) messages per month. We will present Johnzon at ApacheCon Europe, maybe we can attract some people there to join the project.

How has the project developed since the last report?

The project did its 7th incubator release.  
During the last three month 10 jira issues were reported and fixed.  
Project has also started to implement the upcoming JSR-374 specification.

Date of last release:

2015-AUG-17 (0.9-incubating)

When were the last committers or PMC members elected?

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

Signed-off-by:

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

-----  
Kalumet

Kalumet a complete environment manager and deployer including J2EE environments (application servers, applications, etc), softwares, and resources.

Kalumet has been incubating since 2011-09-20.

Unfortunately, due to low activity and lack of time, the Kalumet community didn't success to move forward on the releases and implementation of the expected feature.

We are thinking about retiring Kalumet from the incubator. A formal vote has been sent on the mailing list.

Signed-off-by:

[X](kalumet) Jim Jagielski  
[ ](kalumet) Henri Gomez  
[X](kalumet) Jean-Baptiste Onofre  
[ ](kalumet) Olivier Lamy

Shepherd/Mentor notes:

John D. Ament (johndament):

Activity has been extremely low on the podling, pinged them to see how they were doing, received a quick response. It's always sad to see a podling choose to retire, but it's a valid resolution in this case.

-----  
log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

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

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

No.

How has the community developed since the last report?

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

How has the project developed since the last report?

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

Date of last release:

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

When were the last committers or PMC members elected?

Signed-off-by:

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

-----  
MRQL

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

MRQL has been incubating since 2013-03-13.

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

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

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

None

How has the community developed since the last report?

There were no new developers or new committers since our last report, but the project got contributions from a non-committer (from the Flink community).

How has the project developed since the last report?

During the last three months, 5 Jira issues were reported, from which 4 were fixed. Most of these issues were related to MRQL query evaluation in Flink and Hama modes running on Yarn.

Date of last release:

2015-02-25

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

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

Shepherd/Mentor notes:

Alan Cabrera (adc):

Things seem to be chugging along nicely. The community is not the size of Apache Hadoop, but it's moving in the right direction.

-----  
ODF Toolkit

Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents

ODF Toolkit has been incubating since 2011-08-01.

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

1. Grow the community
2. Generate additional releases
3. Encourage community engagement beyond Jira

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

Lack of mentor sign-off on August 2015 report.

How has the community developed since the last report?

Received+applied patches from new contributor.

How has the project developed since the last report?

Most activity via Jira (12 active issues since May 2015)

Date of last release:

2014-06-02

When were the last committers or PMC members elected?

2012-10-29

Signed-off-by:

[ ](odftoolkit) Sam Ruby  
[ ](odftoolkit) Nick Burch  
[ ](odftoolkit) Yegor Kozlov  
[x](odftoolkit) Rob Weir

Shepherd/Mentor notes:

John D. Ament (johndament):

The podling is not very active but is thriving. There appears to be a tight community behind it. It may help the podling to bring in some new mentors to keep things moving. Potential new contributors have begun discussing on the mailing lists, so we'll see how they grow.

-----  
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. Provide easier user guide/documentation to support new users and contributors.
3. Discuss and decide on roadmap for next release

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

None

How has the community developed since the last report?

1. Added three new contributors: (Hanish Bansal, Mack Hendricks, Sailaja Polavarapu)
2. Added one new committer (Gautam Borad)
3. Mailing list activity since last report (June-01-2015 to Aug-31-2015):  
@dev 856  
@user 259  
@commit 355
4. Issues (JIRAs) created/resolved since last report (after June-01-2015 before Sep-01-2015):  
Created: 113  
Resolved: 166

How has the project developed since the last report?

1. Project has successfully completed - second major release of Apache Ranger - 0.5.0 on 06/10/2015
2. Members have been working on roadmap for the next release - 0.6.0.

Date of last release:

JUNE-10-2015

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  
[ ](ranger) Jakob Homan  
[X](ranger) Owen O'Malley

-----  
REEF

REEF (Retainable Evaluator Execution Framework) is a scale-out computing fabric that eases the development of Big Data applications on top of resource managers such as Apache YARN and Mesos.

REEF has been incubating since 2014-08-12.

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

1. Discussion on graduation is on-going at dev@. Consensus needs to be reached.

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

No

How has the community developed since the last report?

- \* We have seen significant activities from the community since last report
  - 346 JIRA issues created
  - 242 JIRA issues resolved
  - 2500+ emails on dev (including auto-generated messages)
- \* New committers have been added to the team
  - Mariia Mykhailova from Microsoft
  - Andrew Chung from Microsoft
- \* Byung-Gon Chun and Brian Cho will present a tutorial on REEF at the "Big Data Open Source Systems" workshop at VLDB 2015.

How has the project developed since the last report?

- \* 0.12.0-incubating released
  - release date: August 15, 2015
  - Vote: 3 +1 binding votes and 0 -1 votes
  - 177 issues, tasks, features were resolved in this release.
- \* Apache release versioning and packaging scripts added
- \* Group Communications Gather and Scatter added
- \* Driver restartability under active development
- \* IMRU (Iterative Map-Reduce Update) under active development
- \* Vortex (runtime for utilizing volatile resources) under active development, in collaboration with Microsoft Research, Asia

Date of last release:

2015-08-15  
(0.12.0-incubating)

When were the last committers or PMC members elected?

- \* Andrew Chung: 2015-08-28
- \* Mariia Mykhailova: 2015-08-11

Signed-off-by:

[X](reef) Chris Douglas  
[ ](reef) Chris Mattmann  
[ ](reef) Ross Gardler  
[X](reef) Owen O'Malley

-----  
SAMOA

SAMOA provides a collection of distributed streaming algorithms for the most common data mining and machine learning tasks such as classification, clustering, and regression, as well as programming abstractions to develop new algorithms that run on top of distributed stream processing engines (DSPEs). It features a pluggable architecture that allows it to run on several DSPEs such as Apache Flink, 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. Grow the community
2. Elect new PMC members
3. Following up our first release with further 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?

Mailing list activity (June-July-August 2015):

- \* @dev 225 messages

Jira issues backlog (June-July-August 2015):

- \* Created: 10
- \* Resolved: 5

Our main goal is still to grow the community, which is still pretty small. We are doing a large amount of dissemination work in conferences and events to promote SAMOA.

We have had contributions from outside the PMC. In particular, we have been collaborating with the Apache Flink community.

How has the project developed since the last report?

Main developments:

- \* First release 0.3.0 in July.
- \* Cleanup of codebase to ease adoption by new contributors.

We should increase the rate of technical contribution to the project and move to less incremental ones.

Date of last release:

2015-07-21

When were the last committers or PMC members elected?

None

Signed-off-by:

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

-----  
Sentry

Sentry is a highly modular system for providing fine grained role based authorization to both data and metadata stored on an Apache Hadoop cluster.

Sentry has been incubating since 2013-08-08.

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

1. Encourage more feature and direction discussions on the dev list rather than jira.
2. Continue reporting on time
3. Continue making periodic releases following the Apache guidelines.

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

Community has made been making good progress on all of the above mentioned items.

1. There have been multiple discussions on the dev list about features/ design decisions/ roadmap/ dev practices.
2. Community has been doing timely monthly reporting.
3. Community is actively working on bug fixes and working to release 1.6.0 next month

How has the community developed since the last report?

We had 160 messages on dev list last month. (Got number from <http://markmail.org/list/org.apache.sentry.dev>)

How has the project developed since the last report?

About 42 issues were created and about 42 resolved(Numbers from jira).

Date of last release:

2015-07-14

When were the last committers or PMC members elected?

Colin Ma, Dapeng Sun, Guoquan Shen and Xiaomeng Huang were added as committers on 12/24/2014. No new PPMC members have been added since the project has entered the incubator.

Signed-off-by:

[X](sentry) Arvind Prabhakar  
[ ](sentry) Joe Brockmeier  
[X](sentry) David Nalley  
[ ](sentry) Olivier Lamy  
[X](sentry) Patrick Hunt  
[ ](sentry) Thomas White

-----  
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. Complete unit testing
2. Follow the release guide to complete the first release
3. Plan the features for the second release

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?

1. Researchers from a Stanford group contacted Singa committers regarding adding new features to Singa.
2. Some Singa team members from NUS attended workshops associated with SIGMOD15 (Melbourne), VLDB15 (Hawaii), and a local workshop on deep learning in Singapore, to extend the reach of the project to developers from other institutions and countries.
3. There are new developers contributing to the code base.

How has the project developed since the last report?

1. Implementation of the features included in the first release was completed. Particularly, we added contrastive divergence algorithm for training RBM models. Checkpointing and restoring are also implemented.
2. Code review is completed for 3/4 of the code. We also tuned and optimised the code to improve efficiency and scalability.
3. User documents are added at <http://www.comp.nus.edu.sg/~dbsystem/singa//docs.html>. We are moving them to the Apache website.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Thejas Nair (thejas):

The documentation url mentioned above is a staging url. Content will be moved/copied to apache website after initial version is ready.

-----  
Slider

Slider is a collection of tools and technologies to package, deploy, and manage long running applications on Apache Hadoop YARN clusters.

Slider has been incubating since 2014-04-29.

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

1. Getting more external users
2. Getting more diverse set of developers
3. Getting more diverse set of committers/PMC

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

No

How has the community developed since the last report?

We are seeing a lot of traction in the community, over the improvement in deployment and management of applications using Slider. Good number of emails are pouring in the dev list, primarily on creation of Slider packages for newer applications and willingness to contribute code. Few presentations and meetups were organized by the community, socializing their Slider-ized applications. Occasional patches were submitted and a number of features added to the recent releases were driven by end-user needs. We still have to get those people into long term coding, and then bring them in to the committer group.

How has the project developed since the last report?

We released slider-0.80.0-incubating in May with 64 JIRAs resolved. It was one of Slider's most important release, with several critical business application features like rolling upgrade, co-processor support for dynamic addition of libraries, and support for Docker containers.

We've been debating when to switch from Hadoop 2.6+ support to 2.7.1+ only. Given how Slider has always tended to push much of the YARN services work, it's time to pick up on the changes they've put in.

Date of last release:

2015-05-18 slider-0.80.0-incubating

When were the last committers or PMC members elected?

2015-07-07: Yu (Thomas) Liu

Signed-off-by:

[ ](slider) Arun C Murthy  
[X](slider) Devaraj Das

(slider) Jean-Baptiste Onofré  
 (slider) Mahadev Konar

-----  
Streams

Apache Streams (incubating) is a lightweight (yet scalable) framework for ActivityStreams.

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. Collaborate with other Apache projects to improve integrations with their software.

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

# of contributors submitting and reviewing pull requests has been uncharacteristically low in 2015.

How has the community developed since the last report?

# of contributors submitting and reviewing pull requests has been uncharacteristically low in 2015.

How has the project developed since the last report?

26 issues created / 21 merged / 4 in pull request this quarter

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  
 (streams) Ate Douma  
 (streams) Craig McClanahan

Shepherd/Mentor notes:

John D. Ament (johndament):

Most mailing list activity is JIRA/Github. Little live communication occurring.

-----  
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. Motivate community engagement
2. Community growth
3. Continued release planning

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

Thank you to GSOC and three very successful students!

How has the community developed since the last report?

New contributor of documentation.

3 GSOC projects completed, all contributions merged.

Mailing list activity for Jun/Jul/Aug 2015:

- dev@taverna: 392/336/194
- users@taverna: 7/14/11
- commits@taverna: 338/367/211

How has the project developed since the last report?

First Apache release: taverna-language-0.15.0-incubating.  
Documented release procedure and ironed out LICENSE/NOTICE.

Focus on documentation.

Date of last release:

2015-08-11 taverna-language-0.15.0-incubating

When were the last committers or PMC members elected?

n/a

Signed-off-by:

- [x](taverna) Andy Seaborne
- [ ](taverna) Chris Mattmann
- [ ](taverna) Suresh Srinivas
- [ ](taverna) Suresh Marru
- [x](taverna) Marlon Pierce

Shepherd/Mentor notes:

Marlon Pierce (mpierce):

Nice quarter with first Apache release and 3 successful GSOC projects.

-----  
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 a deeper understanding of the Apache Way.
2. Resolve remaining license conflicts.
3. Make our first Apache (incubating) release.

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

None.

How has the community developed since the last report?

- \* JIRAs being created and commented on by end users from at least 2 organizations.
- \* Contribution in terms of code from an end user.
- \* First committer elected from community that was introduced to Trafodion since incubation.
- \* User participation in public lists remains steady.
- \* Distinct culture shift -- discussions have moved out from private

distribution lists into public Apache distribution lists

How has the project developed since the last report?

- \* Wiki migration almost complete. Several new pages added.  
<https://cwiki.apache.org/confluence/display/TRAFODION/>
- \* 14 authors have made 51 commits to master in August. 5,574 files changed, 111,847 line additions, 93,047 deletions.
- \* 41 JIRAs created and 1000 resolved in August 2015
- \* Copyright header changed in all source files.

Date of last release:

We have not done a release of Trafodion since incubation.

When were the last committers or PMC members elected?

Steve Varnau and Pierre Smits elected as committers in August, 2015.

Signed-off-by:

[X](trafodion) Andrew Purtell  
[ ](trafodion) Devaraj Das  
[ ](trafodion) Enis Söztutar  
[ ](trafodion) Lars Hofhansl  
[X](trafodion) Michael Stack  
[X](trafodion) Roman Shaposhnik

-----  
Zeppelin

A collaborative data analytics and visualization tool for distributed, general-purpose data processing systems such as Apache Spark, Apache Flink, etc.

Zeppelin has been incubating since 2014-12-23.

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

1. Do an Apache release
2. Community growth: more committers from different organisations
3. Increase adoption

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?

Traffic in users@ and dev@ grew and +46 subscribers in @dev, +132 subscribers in users@ since last report.

How has the project developed since the last report?

139 patches has been merged and 181 patches has been submitted from 57 different contributors since last report.

Made first release under Apache incubation. Integration to Apache Flink, Apache Ignite, Apache Lens, Apache Geode, Cassandra, Apache Phoenix, PostgreSQL, Apache Kylin has been added. Lots of improvements in front-end side for better user experiences.

Date of last release:

2015-07-31

When were the last committers or PMC members elected?

2015-03-27

Signed-off-by:

(zeppelin) Konstantin Boudnik  
 (zeppelin) Henry Saputra  
 (zeppelin) Roman Shaposhnik  
 (zeppelin) Ted Dunning  
 (zeppelin) Hyunsik Choi