

August2016

/\ *FINAL* (!)

Incubator PMC report for August 2016

Timeline

Wed August 03	Podling reports due by end of day
Sun August 07	Shepherd reviews due by end of day
Sun August 07	Summary due by end of day
Tue August 09	Mentor signoff due by end of day
Wed August 10	Report submitted to Board
Wed August 17	Board meeting

Shepherd Assignments

Drew Farris	Climate Model Diagnostic Analyzer		
Drew Farris	Guacamole		
John Ament	Beam		
John Ament	Quickstep		
Justin Mclean	Pony Mail		
P. Taylor Goetz	Blur		
P. Taylor Goetz	Fineract		
P. Taylor Goetz	Unomi		
Timothy Chen	Impala		
Timothy Chen	S2Graph		
Timothy Chen	SystemML		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f17f61cd-7e2f-4102-bde9-5006d22769b9"><ac:plain-text-body><![CDATA[[none]	BatchEE]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="29eadaf3-85da-4f1a-ab46-5c3db325a967"><ac:plain-text-body><![CDATA[[none]	CarbonD ata]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ae299ddf-3ea9-4f7c-9ac3-804acc89b656"><ac:plain-text-body><![CDATA[[none]	Distribut edLog]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="082a47e5-9b48-43eb-8045-97470ca39c78"><ac:plain-text-body><![CDATA[[none]	Eagle]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="3c098014-d745-481d-ad92-56c4ed156458"><ac:plain-text-body><![CDATA[[none]	Fluo]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b3f60d2f-5ab0-44f6-977f-aac463fd5324"><ac:plain-text-body><![CDATA[[none]	Joshua]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="09c5d47f-951d-4eb3-a93c-b2f0080a22d0"><ac:plain-text-body><![CDATA[[none]	Juneau]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="0f945449-df32-4af9-94b5-27d7eca69ca8"><ac:plain-text-body><![CDATA[[none]	OpenAz]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ec778d8e-70e6-4970-92db-bd008b1b240e"><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="7f04ff3b-b6a1-4ef9-9705-89697b4e609c"><ac:plain-text-body><![CDATA[[none]	Predictio nIO]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5e8b6288-3400-4fa9-b504-2af739ea2e02"><ac:plain-text-body><![CDATA[[none]	Quarks]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="9989de5c-b377-4d0c-a583-72ea4a08758e"><ac:plain-text-body><![CDATA[[none]	SensSoft]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="52f09ba8-2f74-4a8b-9830-e209e5c50ca0"><ac:plain-text-body><![CDATA[[none]	Sirona]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="523ae49b-85d9-49b2-aa32-d62cd19a8b14"><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="0c3bef7c-22f2-4595-9bf8-ec92087480c7"><ac:plain-text-body><![CDATA[[none]	Tamaya]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="8d0e0628-263d-4bb0-a519-8a8d9cbce82c"><ac:plain-text-body><![CDATA[[none]	Toree]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="dc813bc4-f571-4b7d-a97b-9c4e8a04312c"><ac:plain-text-body><![CDATA[[none]	Traffic Control]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="1d336976-ab20-4d54-8ffc-cc25b66eb9fc"><ac:plain-text-body><![CDATA[[none]	iota]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for August 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 59 podlings currently undergoing incubation. We added two podlings to the roster and have no graduations this month. July was a quiet month, completing four releases and no other major activities.

* Community

New IPMC members:

- None

People who left the IPMC:

- None

* New Podlings

- SensSoft
- Traffic Control

* Graduations

The board has motions for the following:

- None

* Releases

The following releases entered distribution during the month of July:

- Apache Trafodion 2.0.1-incubating
- Apache Atlas 0.7-incubating
- Apache Ranger 0.6.0-incubating
- Apache Eagle 0.4.0-incubating

* Legal / Trademarks

- See also: Fluo's report

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- CarbonData
- DistributedLog
- Juneau
- Pirk
- Pony Mail

- SensSoft
- Traffic Control

* Not yet ready to graduate

No release:

- Blur
- CMDA
- Guacamole
- Impala
- iota
- PredictionIO
- Quickstep
- S2Graph

Community growth:

- Beam
- Eagle
- Fluo
- Joshua
- Slider
- SystemML

* Did not report, expected next month

- Fineract
- Quarks

Table of Contents

BatchEE
 Beam
 Blur
 CarbonData
 Climate Model Diagnostic Analyzer
 DistributedLog
 Eagle
 Fluo
 Guacamole
 Impala
 iota
 Joshua
 Juneau
 OpenAz
 Pirk
 Pony Mail
 PredictionIO
 Quickstep
 S2Graph
 SensSoft
 Sirona
 Slider
 SystemML
 Tamaya
 Toree
 Traffic Control
 Unomi

BatchEE

BatchEE projects aims to provide a JBatch implementation (aka JSR352) and a set of useful extensions for this specification.

BatchEE has been incubating since 2013-10-03.

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

None, as the PPMC feels the project is ready to graduate

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 got several new users and interest in advanced integration (in particular our GUI).

How has the project developed since the last report?

Not much but it is quite stable and no new spec was out since last report so it behaves as expected.

BatchEE is used in TomEE and many industry projects.

So we also get feedback in indirect ways via TomEE, Geronimo, etc.

Date of last release:

2015-12-08 batchee-0.3-incubating

When were the last committers or PMC members elected?

* Reinhard Sandtner on 2015-12-01

Signed-off-by:

[](batchee) Jean-Baptiste Onofré

[](batchee) Olivier Lamy

[X](batchee) Mark Struberg

Shepherd/Mentor notes:

struberg: I think this podling is ready to graduate.

Beam

Apache Beam is an open source, unified model and set of language-specific SDKs for defining and executing data processing workflows, and also data ingestion and integration flows, supporting Enterprise Integration Patterns (EIPs) and Domain Specific Languages (DSLs). Beam pipelines simplify the mechanics of large-scale batch and streaming data processing and can run on a number of runtimes such as Apache Flink, Apache Gearpump, Apache Spark, and Google Cloud Dataflow (a cloud service). Beam also brings SDKs in different languages, allowing users to easily implement their data integration processes.

Beam has been incubating since 2016-02-01.

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

1. Additional and continued Beam releases
2. Grow the community of Beam users and contributors
3. Add to and improve upon documentation, code samples, and project website

Any issues that the Incubator PMC (IPMC) or ASF Board wish/need to be aware of? None.

How has the community developed since the last report?

* 425 closed/merged pull requests

* High engagement on dev and user mailing lists (590 / 455 messages)

* Several public talks, articles, and videos including:

- * Hadoop Summit San Jose ("Apache Beam: A Unified Model for Batch and Streaming Data Processing" & "The Next Generation of Data Processing & OSS")
- * O'Reilly & The New Stack ("Future-proof and scale-proof your code")
- * QCon NY ("Apache Beam: The Case for Unifying Streaming API's")
- * JBCN Barcelona ("Introduction to Apache Beam")

How has the project developed since the last report?

Major developments on the project since last report include the following:

- * First incubating release (0.1.0-incubating)
- * Second incubating release (0.2.0-incubating)
- * Addition of Apache Beam Python SDK
- * Addition of the Apache Gearpump runner
- * Added support for writing to Apache Kafka clusters
- * Added support for reading from and writing to Java Message Services, including Apache ActiveMQ, GeronimoJMS, and RabbitMQ
- * Ratified new Beam model APIs to improve efficiency and failure handling: DoFn setup, teardown, and reuse
- * Optimized key components such as data serialization and shuffle
- * Continued improvements to the Flink, Spark, and Dataflow runners
- * Continued reorganization and refactoring of the project
- * Continued improvements to documentation and testing

Date of last release:

- * 2016/06/15 - 0.1.0-incubating
- * 2016/08/08 - 0.2.0-incubating)

When were the last committers or PMC members elected?

N/A - no changes since last report.

Signed-off-by:

```
[X](beam) Jean-Baptiste Onofre
[ ](beam) Venkatesh Seetharam
[X](beam) Bertrand Delacretaz
[X](beam) Ted Dunning
```

Blur

Blur is a search platform capable of searching massive amounts of data in a cloud computing environment.

Blur has been incubating since 2012-07-24.

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

1. Greater community involvement.
2. Produce releases.
- 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?

Subscriptions: user@ - 61[0]; dev@ - 79[+1]
The community involvement has not really changed over the past few months.

How has the project developed since the last report?

- Not much has changed, a few bug fixes.

Date of last release:

2014-07-29

When were the last committers or PMC members elected?

2014-07-28

Signed-off-by:

[](blur) Doug Cutting
[X](blur) Patrick Hunt
[x](blur) Tim Williams

CarbonData

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

CarbonData has been incubating since 2016-06-02.

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

1. Prepare first CarbonData release
2. Prepare first CarbonData website
3. Promote the project and grow user and 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?

We voted two new committers: Venkata Ramana and Ravindra Pesala.

We are preparing the CarbonData website, blog posts and talks to promote CarbonData and grow the user and dev communities.

The community activity increased: many new users started to use and test CarbonData, we had more than 100 new issues created during July.

On the other hand, presentation material has been created and a first talk has been given:

<http://www.slideshare.net/liangchen18/apache-carbondatanew-high-performance-data-format-for-faster-data-analysis>

How has the project developed since the last report?

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

We also created the Jenkins CI jobs.

We are now in the process of cleanup and polishing the build and legal to prepare the first release.

Date of last release:

Not yet available

When were the last committers or PMC members elected?

2016-07-15

Signed-off-by:

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

Climate Model Diagnostic Analyzer

CMDA provides web services for multi-aspect physics-based and phenomenon-oriented climate model performance evaluation and diagnosis through the comprehensive and synergistic use of multiple observational data, reanalysis data, and model outputs.

Climate Model Diagnostic Analyzer has been incubating since 2015-05-08.

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

1. Enlarge the CMDA community;
2. Develop stable platform to allow broad contributors to turn in code;
3. Develop a streamline procedure to allow the entire community to develop the CMDA project

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

No.

How has the community developed since the last report?

In this month, the CMDA community focuses on developing the summer school program, in order to advocate and distribute CMDA 2.0, to educate more students on our latest product, and to attract more people to join our community. So far, we have received about 25 graduate students who show interest and will come to our summer school to learn CMDA 2.0 product and the platform.

Meanwhile, with the great guidance from our mentors, the CMDA community has learned to run the project from various aspects in an Apache way. Many community meetings were held to streamline and develop our operation regulations. Here we just name a few strategies. Weekly community meetings are arranged with pre-distributed meeting agenda. Meeting invites are sent to all community members with open access to the weekly meetings. After each community meeting, various contributors/groups summarize their to-do items on the community mailing list, to welcome all community members' input. Community communication always adopts the community mailing list as carrier. More committers join the voting group. More committers are recommended based on their active performance in the community.

How has the project developed since the last report?

We have successfully encapsulated and deployed CMDA 2.0 onto Amazon cloud environment, leveraging docker technology to ensure scalability. Provenance engine is also deployed onto Amazon cloud to ensure that we gather all provenance during the summer school for further study and development of our project after the summer school.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

Signed-off-by:

[](climatemodeldiagnosticanalyzer) James W. Carman
[X](climatemodeldiagnosticanalyzer) Chris Mattmann

```
[ ](climatemodeldiagnosticanalyzer) Michael James Joyce
[ ](climatemodeldiagnosticanalyzer) Kim Whitehall
[ ](climatemodeldiagnosticanalyzer) Gregory D. Reddin
```

Shepherd/Mentor notes:

- * Chris Mattmann - The community has shown progress in addressing a near vote to terminate the podling. There were and still are concerns that need to be tended to and watched but as a mentor I am pleased by the efforts of Jia, Seungwon, and Lei whom I have seen contribute to the discussion. I think we should continue to closely evaluate the progress, but I am happy with this month.
- * Drew Farris - Three mentors active on the mailing lists. Kudos to Chris to the amount of engagement here. After choosing to avoid retirement the community is taking concrete steps to increase its engagement with ASF by establishing a web presence at Apache, creating committer accounts, preparing code for contribution and moving processes to the mailing lists. This podling is moving to a monthly reporting cycle for the next three months progress can be monitored closely.

DistributedLog

DistributedLog is a high-performance replicated log service. It offers durability, replication and strong consistency, which provides a fundamental building block for building reliable distributed systems.

DistributedLog has been incubating since 2016-06-24.

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

- * Getting an Apache release out
- * Improve general documentation
- * Grow user base

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

There is currently no issue.

How has the community developed since the last report?

- * A few people have started to report issues
- * Some conversations have been initiated on the mailing list
- * Sijie will be present at Strata in China to present the project

How has the project developed since the last report?

- * The wiki has been set up (cwiki.apache.org/confluence/display/DL/)
- * A few contributions were made to the repository on GitHub and are in the process of being merged
- * Various mailing lists have been set up
- * SGA from Twitter has been completed
- * Status page is up and we have started filling it out
- * Most accounts of initial committers have been created
- * Jira queue has been created (<https://issues.apache.org/jira/browse/DL>)
- * Repository has been created (and mirrored on GitHub)

Signed-off-by:

```
[x](distributedlog) Flavio Junqueira
[x](distributedlog) Chris Nauroth
[ ](distributedlog) Henry Saputra
```

Shepherd/Mentor notes:

Flavio Junqueira: Project has been progressing well. All three mentors have been actively participating in bringing the project up to speed and helping out with issues. Several contributors have expressed interest in the project

and there is some initial discussion around issues, improvements, and new features. I'm personally happy with the progress.

Eagle

Apache Eagle (incubating) is an open source analytics solution for identifying security and performance issues instantly on big data platforms, e.g. Apache Hadoop, Apache Spark etc. It analyzes data activities, Yarn application, Hadoop JMX metrics and daemon logs etc., provides state-of-the-art alert engine to identify security breach, performance issues and shows insights.

Eagle has been incubating since 2015-10-26.

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

1. Eagle community is discussing graduation and use Apache maturity assessment to measure the gap to graduation. That should be completed, see <https://cwiki.apache.org/confluence/display/EAG/Eagle+Podling+Maturity+Assessment>
2. Complete one requirement from self assessment which is to provide a well-documented channel to report security issues, along with a documented way of responding to them.
3. Identify other housekeeping before graduation, for example agreement on bylaws, PMC composition etc.

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

NIL

How has the community developed since the last report?

- Presented in Apache Con, May 2016
- Presented in London Stratus + Hadoop, May 2016
- Presented in San Jose Hadoop Summit, June 2016
- Presented in DTCC2016 Beijing, China, May 2016
- Presented in JavaCon, China, May 2016
- Presented in Hadoop Summit, San Jose, China, June 2016
- Presented in Data Asset Management Summit, Shanghai, China, July 2016

Communities showed continuous interest in Apache Eagle project. Some company tried to integrate Eagle as part of whole solution and some contributed different use case to Eagle platform, for example MapR support, Oozie monitoring support etc.

How has the project developed since the last report?

Technically we added more documents for user to easily use Eagle

Besides Hadoop security monitoring, we have started working on performance monitoring design and development based on the requests from community.

As we may have multiple different use cases running on top of Eagle, a design discussion and prototype is conducted to make sure Eagle is a framework to host multiple use cases and user will use Eagle to manage those use cases.

Alert engine is going to be decoupled from data processing so that alert engine will be multi-tenant and easy to be used by multiple use cases by incorporating well-designed metadata.

Date of last release:

2016-07-19

When were the last committers or PMC members elected?

- Daniel Zhou, 2016-06-15
- Michael Wu, 2016-06-17

Signed-off-by:

[](eagle) Owen O'Malley
[](eagle) Henry Saputra
[x](eagle) Julian Hyde
[x](eagle) P. Taylor Goetz
[x](eagle) Amareshwari Sriramdasu

Fluo

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

Fluo has been incubating since 2016-05-17.

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

1. Attract new contributors
2. Do a release
3. Conduct a podling name search

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

No

How has the community developed since the last report?

We have a new contributor to Fluo since the last report.

How has the project developed since the last report?

We completed refactoring of code to use Apache names. We attempted our first release of a parent pom (for both Fluo projects) and it failed on the incubator list. A parent pom is a way to share build configuration among multiple projects. We plan for Fluo and Fluo Recipes to both inherit the same parent pom and obtain common config for code formatting, findbugs, checkstyle rules, etc. For this to work the parent pom needs to be released before the projects.

The reason it failed is because members of the IPMC identified the following issues :

- * Improper linking to pre-Apache releases of Fluo from website.
- * Branding and linking issues with fluo.io domain, fluo-io GitHub org and projects. The domain name and github org were used by the project before moving to Apache.

This was really good feedback. We want to remedy the issues identified and eliminate possible confusion around the Apache Fluo brand. The following link identifies discussion of a possible way forward.

<https://lists.apache.org/thread.html/b478645815bf7b1fee8342a21924162decc173d91c22eaf5f708a435@%3Cdev.fluo.apache.org%3E>

Date of last release:

Never

When were the last committers or PMC members elected?

Never

Signed-off-by:

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

Shepherd/Mentor notes:

JE: There was a very heavy discussion on branding spurred from the attempted VOTE thread. In the end, I believe this was a very positive discussion and resulted in fixing many branding issues. I believe most of them were due to the transition from an independent Github community to Apache (which included bringing the website). The PPMC has been extremely responsive to incubator feedback and I am very happy with the level of response they spearheaded (without mentor prodding/coaxing).

Guacamole

Guacamole is an enterprise-grade, protocol-agnostic, remote desktop gateway. Combined with cloud hosting, Guacamole provides an excellent alternative to traditional desktops. Guacamole aims to make cloud-hosted desktop access preferable to traditional, local access.

Guacamole has been incubating since 2016-02-10.

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

1. Making the first Guacamole release under the Apache Incubator
2. Encouraging community participation and contribution
3. Accepting additional committers

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

None.

How has the community developed since the last report?

Activity on the mailing lists has increased now that the new project website is up. At the time of the last report, the @user list was unused. Since then, the number of emails has increased each month, with 52 in May, 95 in June, and 123 in July. Users do still occasionally post to the old SourceForge forums, but are then gently redirected to the mailing lists.

Excluding project committers, the community has opened 5 pull requests, 2 of which have been merged.

How has the project developed since the last report?

The project has finished migrating to Apache infrastructure, including the project website. There is one major outstanding task, screen sharing, which has been blocking the release but is finally nearing completion.

Since last report, the project has had roughly 250 commits, about 60% of which were directly related to the outstanding task mentioned above.

Date of last release:

2015-12-18 (0.9.9, prior to Apache Incubator)

When were the last committers or PMC members elected?

The most recent committer, Frode Langelo, was accepted into the project by VOTE on 2016-04-03, with the required ICLA received on 2016-04-05.

Signed-off-by:

[x](guacamole) Jean-Baptiste Onofre
[](guacamole) Daniel Gruno
[](guacamole) Olivier Lamy

[X](guacamole) Jim Jagielski
[X](guacamole) Greg Trasuk

Impala

Impala is a high-performance C++ and Java SQL query engine for data stored in Apache Hadoop-based clusters.

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

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

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

No.

How has the community developed since the last report?

Our last report was in April. Since then

- * Six new contributors have submitted patches for review, and two contributors new to the project since incubation have continued to send patches.
- * Mailing list activity more than doubled in the four months since our last report compared to the four months before that, from 31 threads to 75 threads (excluding patch review comments)

How has the project developed since the last report?

- * The podling name search was completed: <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-96>
- * Trademark handoff was completed
- * Impala's git repository is now hosted on ASF infrastructure
- * Impala's website's source is hosted on ASF's git infrastructure and the website is now available on <https://impala.apache.org>
- * Project bylaws have been ratified: <https://impala.apache.org/bylaws.html>
- * Developer documentation has started to move to the ASF-hosted wiki: <https://cwiki.apache.org/confluence/collector/pages.action?key=IMPALA>
- * Work has begun in migrating to ASF-hosted JIRA
- * A patch changing the copyright headers is in review

Date of last release:

No releases have been made yet.

When were the last committers or PMC members elected?

No committers or PMC members have been added since incubation began.

Signed-off-by:

[X](impala) Tom White
[](impala) Todd Lipcon
[](impala) Carl Steinbach
[](impala) Brock Noland

iota

Open source system that enables the orchestration of IoT devices.

iota has been incubating since 2016-01-20.

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

1. Building a community now that we have completed source code upload of one for the main components of the system.
2. Moving towards an initial release before the end of the year
3. Working on IP clearance for the 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?

A number of individuals have made contributions some of which have been incorporated. We need to encourage more discussion on the dev list in the coming months.

How has the project developed since the last report?

We continue to improve (the Fey engine) we have added tests for the core engine and documentation for the engine and the component set.

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

At the start of the project.

Signed-off-by:

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

Shepherd/Mentor notes:

Justin Mclean: Probably time to come off the extended monthly reporting?

Joshua

Joshua is a statistical machine translation toolkit

Joshua has been incubating since 2016-02-13.

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

1. Ensure first release of Joshua Incubating artifacts (6.1)
2. Continue to build the Joshua PPMC and user community
3. Investigate targeted user communities within Apache

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

None.

How has the community developed since the last report?

We have gained a new contributor, and have continued developing and updating the web page to increase interest. We have not made any real advertising or publicity pushes, but hope to around the time of our first formal release under the Apache banner (targeted for September).

How has the project developed since the last report?

We have been steadily pushing up stability and design improvements, including a move from an ant+ivy to a maven build system. We have

made some changes to our build process, including enabling Travis-CI for continual integration testing. We are in discussion about deeper architectural changes that will facilitate an API.

Date of last release:

N/A

When were the last committers or PMC members elected?

Kellen Sunderland and Felix Hieber (April 11, 2016)
Thamme Gowda (May 26, 2016)

Signed-off-by:

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

Juneau

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

Juneau has been incubating since 2016-06-24.

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

1. Build up community with non-IBM and non-Salesforce contributors.
2. Finish creating website.
3. Finish populating Git repo.

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?

Currently only two of the original proposers have gone through the process of getting commit approval.

Some of the Salesforce proposers have been busy with their release schedule. Others need to be bugged to finish up the paperwork.

How has the project developed since the last report?

Grant from IBM Corporation for Juno (Apache Juneau) has been received by the ASF.

The Juno codebase has been uploaded to the Git repository. We're still in the process of cleaning it up and converting to using Maven for builds.

Work has been done on the website, although build issues appear to be preventing it from being fully deployed.

Date of last release:

N/A

When were the last committers or PMC members elected?

Signed-off-by:

[x](juneau) Craig Russell

[](juneau) Jochen Wiedmann
[X](juneau) John D. Ament

OpenAz

Tools and libraries for developing Attribute-based Access Control (ABAC) Systems in a variety of languages.

OpenAz has been incubating since 2015-01-20.

Shepherd/Mentor notes:

johndament: The community is currently voting to retire.

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. Establish a formal release process and schedule, allowing for dependable release cycles in a manner consistent with the Apache development process.
2. Establish a process which allows different release cycles for the core framework, extensions/adaptors, and additional algorithms.
3. Grow the community to establish diversity of background and expertise.

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 mailing list activity, pull requests, and page views to the Apache Pirk website (via Google Analytics) have greatly increased since the last report.

How has the project developed since the last report?

The project has ensured that all dependencies are compliant with the Apache License version 2.0 and that all code and documentation artifacts have the correct Apache licensing markings and notice. The project has also undertaken a refactor of some of the initial packages and provided more robust support for user interaction.

Date of last release:

No releases have been made yet.

When were the last committers or PMC members elected?

Pirk is operating with its initial committers and PMC and has not yet elected new committers or PMC members.

Signed-off-by:

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

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 base
2. Ensure that releases and processes adhere to ASF standards
3. Spread the word, get more developers onboard.

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?

The community is growing, albeit at a very slow pace, possibly due to summer vacations (and some newborns). We intend to ramp up our recruitment drive in the coming months.

How has the project developed since the last report?

We have managed to vote on and release Pony Mail 0.9.

Trademark issues have been discussed, mostly done but perhaps with a few loose ends to tie up.

Date of last release:

2016-08-02

When were the last committers or PMC members elected?

2016-05-27

Signed-off-by:

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

PredictionIO

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

PredictionIO has been incubating since 2016-05-26.

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

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

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

None

How has the community developed since the last report?

The mailing list activity, pull requests and page views to the PredictionIO website (per google analytics) have greatly increased since the last report. We have announced our new mailing lists on ASF infra on the old Google Groups and have made those groups read only. Since then, we have seen previous joining the new ASF-based community, and some new contributors actively making proposals and patches.

How has the project developed since the last report?

The project has ensured that all dependencies are Apache License 2.0 compliant, and all code. Working to ensure that all code and documentation markings have the correct apache license 2.0 attributions. We have seen contributors who used to submit patches pre-donation are now actively submitting patches to the migrated codebase. There are also new contributors who have submitted proposals about migrating the template gallery to be based on ASF, which we have happily merged. New integration tests have also been proposed and being actively developed by new contributors.

Date of last release:

No releases yet.

When were the last committers or PMC members elected?

PredictionIO is still operating with the initial committers and PMC and has not yet elected new committers or PMC members.

Signed-off-by:

[X](predictionio) Andrew Purtell
[](predictionio) James Taylor
[](predictionio) Lars Hofhansl
[x](predictionio) Luciano Resende
[](predictionio) Xiangrui Meng
[X](predictionio) Suneel Marthi

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. Acquire early adopters.
2. Continue to build the developer community.
3. Create an ASF 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. Added two new members to the (core) Committers group.
2. Added one new member to the Contributor group.
3. Started discussions for a community and governance structure. Initial proposal has been circulated the project developer list.

How has the project developed since the last report?

1. We have closed over 40 pull requests.
2. Created a website.
3. Created a getting started guide.
4. Started work on cleaning up the entire repository to have licensing and

copyrights in line with ASF projects.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

S2Graph

S2Graph is a distributed and scalable OLTP graph database built on Apache HBase to support fast traversal of extremely large graphs.

S2Graph has been incubating since 2015-11-29.

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

1. Make a release
2. Attract users and contributors
3. Foster more and diverse committers

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

None

How has the community developed since the last report?

- * Held a session on S2Graph: Internal at Apache Big Data NA
- * Made various networking efforts at Apache Big Data NA, Vancouver and Druid Meet-up, Seoul
- * Made plans to integrate with Apache TinkerPop in order to reach a wider audience (S2GRAPH-72)
- * Discussion with Apache Zeppelin regarding joint meet-ups at Seoul
- * 3 inquiries from user mailing list
- * Submit paper to INFOCOM 2017

How has the project developed since the last report?

- * We decided project logo.
- * Project website is up.
- * Working on issues on providing package for distribution.
- * 20 issues are created, 18 issues are resolved.
- * Working on first release(S2GRAPH-86)

Date of last release:

No yet

When were the last committers or PMC members elected?

No

Signed-off-by:

[X](s2graph) Andrew Purtell
[](s2graph) Venkatesh Seetharam
[X](s2graph) Sergio Fernández

SensSoft

SensSoft is a software tool usability testing platform

SensSoft has been incubating since 2016-07-13.

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

1. Ensure that Software grant/CCLA are submitted to secretary@
2. Port all infrastructure and community over to the ASF Incubator infrastructure.
3. Grow the Senssoft community and work towards incubating releases.

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?

Senssoft has only literally arrived in the Incubator. Podling bootstrapping is currently underway. CCLA was received by secretary@ and all CLAs have been submitted and received as well. We are ready to begin incubating.

How has the project developed since the last report?

Project repos have been established at the ASF and the community are undertaking QA in order to have them writable.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Sirona

Monitoring Solution.

Sirona has been incubating since 2013-10-15.

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

1. Extend the community
2. Release a bit more often
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?

User community increased a bit but not much the core project community.

How has the project developed since the last report?

We tried to release but have few fixes to do.

Date of last release:

2015-11-03

When were the last committers or PMC members elected?

Signed-off-by:

[](sirona) Olivier Lamy
[](sirona) Henri Gomez
[](sirona) Jean-Baptiste Onofre
[](sirona) Tammo van Lessen
[X](sirona) Mark Struberg

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. Apache Slider community/PPMC has voted to move portions of Slider into Apache Hadoop YARN as modules. It is possible that the remaining pieces will be moved at a later point in time or become obsolete or evolve to work closely with YARN. Slider PPMC will decide what makes the most sense as we progress through this exciting time.
2. Getting more external users
3. Growth of a diverse set of developers/committers/PMC members is also crucial towards the final state of Slider

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

No.

How has the community developed since the last report?

The Apache Hadoop YARN community/PMC and Apache Slider community/PPMC has decided to move portions of Slider into YARN, to make fast and significant progress on: YARN-4692 - [Umbrella] Simplified and first-class support for services in YARN. As a result, portions like Slider Core, Application Master, Client and Tests will roll into Apache Hadoop YARN as module(s). This effort is being captured in: YARN-5079 : Native YARN framework layer for services. At this point, a branch has been created in the Hadoop codebase, and Slider has been selectively migrated. As a result of this, we are already seeing interest in the Hadoop committers/PMC, who have started to contribute and submit patches to Slider.

In order to support existing users of Slider, and to provide seamless migration, there will be sufficient overlap between the time when a stable state of long running services support is available in some future version of YARN and the time till an independent Slider release is available. The community/PPMC will also determine the future state of Slider as we navigate through these changes.

The discussions on the Slider and YARN community DLs can be viewed here -

<https://s.apache.org/0hoh>
<https://s.apache.org/MncV>

How has the project developed since the last report?

We released version 0.91.0-incubating in June, 2016 with significant work

done on token management, unique component support, SLIDER-906 (Support for Docker based application packaging with first class YARN support) and several bug fixes. Work continues on support for complex services (assemblies) and agent-less applications in Slider. The efforts on the Core will continue in the new services branch created in YARN.

Date of last release:

2016-06-28 slider-0.91.0-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

SystemML

SystemML provides declarative large-scale machine learning (ML) that aims at flexible specification of ML algorithms and automatic generation of hybrid runtime plans ranging from single node, in-memory computations, to distributed computations running on Apache Hadoop MapReduce and Apache Spark.

SystemML has been incubating since 2015-11-02.

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

- Grow SystemML community: increase mailing list activity, increase adoption of SystemML for scalable machine learning, encourage data scientists to adopt DML and PyDML algorithm scripts, respond to user feedback to ensure SystemML meets the requirements of real-world situations, write papers, and present talks about SystemML.
- Continue to produce releases.
- Increase the diversity of our project's contributors and committers.

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

NONE.

How has the community developed since the last report?

Our mailing list from May through July had 237 messages including a wide range of topics. We have gained 6 new contributors since May 1st. On GitHub, the project has been starred 365 times and forked 129 times. Fred Reiss spoke at Spark Summit West on June 7 about building custom machine learning algorithms with SystemML. Mike Dusenberry presented about SystemML on May 19 at Datapalooza Denver. Elgohary, Boehm, Haas, Reiss, and Reinwald published Compressed Linear Algebra for Large-Scale Machine Learning.

How has the project developed since the last report?

We produced our second Apache release, version 0.10.0-incubating. The project has had 205 commits since May 1. 187 issues have been reported on our JIRA site and 127 issues have been resolved. 70 pull requests have been created since May 1, and 63 of these have been closed.

Date of last release:

2016-06-15 (version 0.10.0-incubating)

When were the last committers or PMC members elected?

2016-05-07 Glenn Weidner
2016-05-07 Faraz Makari Manshadi

Signed-off-by:

[x](systemml) Luciano Resende
[](systemml) Patrick Wendell
[](systemml) Reynold Xin
[](systemml) Rich Bowen

Tamaya

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

Tamaya has been incubating since 2014-11-14.

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

1. Allow community to build consensus in APIs.
2. Get new contributors.
3. Understand and work through end user use-cases.

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

The community recently had a bit of a falling out between this project and some changes proposed to go into the Geronimo project. Its generally believed that what was being proposed was not malicious. It does however represent some of the core issues around this project's ability create consensus and grow.

The project has already been incubating for two years, with very few releases and very little contributor growth.

How has the community developed since the last report?

- See issues section.

How has the project developed since the last report?

- Releases are still in progress, and work is being done to refactor the Tamaya APIs based on some input from Mark Struberg and others.

Date of last release:

2016-04-06 (0.2-incubating)

When were the last committers or PMC members elected?

Philipp Ottlinger has been elected on 2016-04-22 as new committer.

Signed-off-by:

[X](tamaya) John D. Ament
[X](tamaya) Mark Struberg
[](tamaya) Gerhard Petracek
[](tamaya) David Blevins

Shepherd/Mentor notes:

johndament: Discussions are beginning to talk through the future of the podling. Going to request an extra report next month since this was created by mentors.

Toree

Toree provides applications with a mechanism to interactively and remotely access Apache Spark. It enables interactive workloads between applications and a Spark cluster. As a Jupyter Notebook extension, it provides the user with a preconfigured environment for interacting with Spark using Scala, Python, R or SQL.

Toree has been incubating since 2015-12-02.

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

1. Resolve LGPL dependency: This is in progress, and hopefully we will have the dependency license changed to a license that is compatible with AL2.
2. Make a release: 0.1.x branch should be ready for release. Master has moved to start support for Spark 2.0
3. Grow a diverse community: We should put some emphasis on growing the community and making it diverse (the rule is at least three independent contributors)

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

TOREE-262 - Progress on removal of LGPL dependency

How has the community developed since the last report?

1. Active communication in mailing list and gitter with early adopters
2. Still working on transitioning users from Spark Kernel project into Toree.
3. More external contributions being made. Community member 2mariusvniekerk key in porting Toree to work on Spark 2.0. Sitting on PR 56.

How has the project developed since the last report?

1. Continue working with JeroMQ community to further their transition into MPL v2 and away from LGPL. 3 committers to go.
2. Addressing issues opened by community

Date of last release:

None since incubation.

When were the last committers or PMC members elected?

No new additions since incubation

Signed-off-by:

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

Traffic Control

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

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

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

1. Get set up with all resources (git / mailing lists / JIRA / irc / etc)
2. Name search / name sign off
3. Prepare for next release under apache process

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?

This is the first report, we are just getting started.

How has the project developed since the last report?

This is the first report, we are just getting started. ICLAs are done.

Date of last release:

We released 1.6.0 on 2016-07-29 under the old process, because the release process (candidates and votes) had already started before our official entry in the incubator. Our next release will be the first release where we will follow the Apache incubator process.

When were the last committers or PMC members elected?

2016-07-12 - initial committer list.

Signed-off-by:

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

Unomi

Unomi is a reference implementation of the OASIS Context Server specification currently being worked on by the OASIS Context Server Technical Committee. It provides a high-performance user profile and event tracking server.

Unomi has been incubating since 2015-10-05.

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

1. Continued releases, and updated dependencies
2. Grow up user and contributor communities, seeing more contribution/PR

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 activity of the dev mailing list smoothly increased, discussing key enhancements in the project.

We are now targeting the development of the user community. For that, we discussed about improving and polishing the website. Today, it's obvious that it's not easy to understand what Unomi can do and actually does. The purpose is to give more use cases and introduction on the mailing list.

How has the project developed since the last report?

The first Apache Unomi 1.0.0-incubating has been released.

We updated key dependencies and implemented couple of new releases in addition of the bug fixes. The purpose is to cut off a new release asap.

Date of last release:

2016-03-09

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](unomi) Jean-Baptiste Onofré

[X](unomi) Bertrand Delacretaz

[X](unomi) Chris Mattmann