

March2016

/\ *FINAL* (!)

Incubator PMC report for March 2016

Timeline

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

Shepherd Assignments

Drew Farris	Siro na		
John Ament	Stre ams		
John Ament	Wave		
Justin Mclean	Metr on		
P. Taylor Goetz	Josh ua		
P. Taylor Goetz	Zep elin		
P. Taylor Goetz	log4 cxx2		
Timothy Chen	Imp ala		
Timothy Chen	Joh nzon		
Timothy Chen	MR QL		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="987ddc2c-d6a9-4eb3-8437-109d1e3041bf"><ac:plain-text-body><![CDATA[[non e]	Atlas]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="aae18588-ac35-4b78-bd47-c7cb09568e0e"><ac:plain-text-body><![CDATA[[non e]	Climate Model Diagnostic Analyzer]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c2c51709-9879-4df1-b15e-68fa11a36603"><ac:plain-text-body><![CDATA[[non e]	CommonsRDF]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7e993915-f1ac-44e4-ab96-9e88b06f9ad2"><ac:plain-text-body><![CDATA[[non e]	Fineract]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="360a4044-888b-4527-a62f-e59d3d3708c8"><ac:plain-text-body><![CDATA[[non e]	HTrace]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="62d45c84-c725-43cb-994d-4b89c6c4fc2f"><ac:plain-text-body><![CDATA[[non e]	Kudu]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="48d6d384-f4b0-4738-883d-37a5d27e73ce"><ac:plain-text-body><![CDATA[[non e]	Milagro]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6bfb0a3-ba1c-494f-80ba-c2c8035d61ae"><ac:plain-text-body><![CDATA[[non e]	Myriad]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="742e60a2-0f9a-4576-8e73-bb69a2418beb"><ac:plain-text-body><![CDATA[[non e]	OpenAz]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="bf0df4ff-b26e-4f91-8932-4c5182758066"><ac:plain-text-body><![CDATA[[non e]	Ranger]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b339266a-605e-4ab4-9699-a5ff5af8470a"><ac:plain-text-body><![CDATA[[non e]	SAMOA]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6b85053b-795d-4c93-95c4-c8c3c68df05c"><ac:plain-text-body><![CDATA[[non e]	Sentry]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="c89fd2f6-9e9d-4497-abd9-fd7f61a8371d"><ac:plain-text-body><![CDATA[[non e]	Singa]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ac2dda15-f2a2-47ae-914f-c3fee4296089"><ac:plain-text-body><![CDATA[[non e]	Taverna]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="928d4204-4b6f-4cd9-a288-82595d6edc8e"><ac:plain-text-body><![CDATA[[non e]	Toree]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="965049d4-43d0-404d-87f3-82dc95f8c0f6"><ac:plain-text-body><![CDATA[[non e]	Trafodion]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e56dbb1d-5142-4055-bfb3-c13c8b8df4e7"><ac:plain-text-body><![CDATA[[non e]	iota]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for March 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 52 podlings currently undergoing incubation.

* Community

New IPMC members:

- Kathey Marsden
- Daniel John Debrunner

* New Podlings

- Guacamole
- Joshua
- Quarks

* Graduations

The board has motions for the following:

- Sentry

* Releases

The following releases entered distribution during the month of February:

2016-02-06 Apache Geode 1.0.0-incubating.M1
2016-02-07 Apache Apex v3.3.0-incubating
2016-02-11 Apache Wave 0.4.0-incubating
2016-02-11 Apache SystemML 0.9.0-incubating
2016-02-11 Apache TinkerPop 3.1.1-incubating
2016-02-20 Apache Johnzon 0.9.3-incubating
2016-02-25 Apache Kudu (incubating) 0.7.0
2016-02-26 Apache AsterixDB 0.8.8-incubating

* IP Clearance

- No podling related IP Clearances were processed during this time.

* Legal / Trademarks

- There was substantial discussion on general@incubator about how challenging it is to produce compliant LICENSE and NOTICE files.
- A couple of podlings with runtime LGPL dependencies pledged to remove those dependencies before graduation but were permitted to make

incubating releases in the meantime. There was uncertainty as to whether VP Legal's approval was required; ultimately VP Legal appeared on general@incubator and approved the plans.

* Infrastructure

- The Freemarker podling had a long standing request to use their existing domain as the primary domain. This was approved by Shane Curcuru but closed by infra without comment, see INFRA-10787.
- The Kudu podling is currently running its website as <http://getkudu.io>. The standard webpage of <http://kudu.incubator.apache.org> does not work. An infra ticket has been created to do the conversion INFRA-11407.

* Miscellaneous

- The Incubator's Retirement Guide has been substantially refreshed, and it should now be easier for a Mentor or other volunteer to perform the administrative steps to retire a podling.
- This month saw a large number of mentors resign from podlings.
- Brooklyn submitted an IP Clearance on the incubator mailing lists. Last that was seen about the incubators role in IP Clearance was that it shouldn't be the gatekeeper of top level projects.
- The current count of 52 podlings is likely the largest the incubator has ever seen.

* Credits

- Report Manager: Marvin Humphrey
- Interim Report Manager: John D. Ament

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

* Still getting started at the Incubator

- Impala
- Joshua
- Milagro

* Not yet ready to graduate

No release:

- CMDA
- Fineract
- log4cxx2
- Metron
- Toree

Community growth:

- CommonsRDF
- Kudu
- Myriad
- SAMOA
- Singa
- Taverna
- Trafodion
- Zeppelin

* Ready to graduate

The Board has motions for the following:

- Sentry

* Did not report, expected next month

- HTrace
- Iota
- Johnzon
- OpenAz
- Ranger
- Sirona
- Wave

Table of Contents

Atlas
 Climate Model Diagnostic Analyzer
 CommonsRDF
 Fineract
 Impala
 Joshua
 Kudu
 log4cxx2
 Metron
 Milagro
 MRQL
 Myriad
 SAMOA
 Sentry
 Singa
 Streams
 Taverna
 Tere
 Trafodion
 Zeppelin

Atlas

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

Atlas has been incubating since 2015-05-05.

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

1. Podling name search is pending
2. Expand the community and add more committers

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

- * Everything seems to be smooth, nothing specific at this time.

How has the community developed since the last report?

1. The user/dev community show increasing interests.
The average number of messages in a month is around 400.
3. Around 10 contributors were added, number of contributors is at 42 currently.

How has the project developed since the last report?

1. Integrations with storm, falcon and sqoop were added. Tagging feature is now available in the dashboard. Community is working on entity deletes, business catalog, HA and improvements to existing integrations.
2. 0.6-incubating version was released in Dec, 2015.
3. A total of 143 issues were created and 48 of them were resolved since Jan, 2016

Date of last release:

- * 2015-12-31 - 0.6-incubating

When were the last committers or PMC members elected?

- * Suma Shivaprasad was elected as a committer on 10/02/2015.

Signed-off-by:

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

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. Moving project communication to the mailing list;
2. Moving development activities to be owned by Apache;
3. Working on community building.

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?

- * Started using the mailing list dev@cmda.incubator.apache.org to communicate project development issues;
- * Created a Docker container that will be used for our future development as well as deployment vehicle (in progress);
- * Working on a unified git repository to host both the CMU and the JPL code bases;

How has the project developed since the last report?

- * The JPL team developed the following new services:
 - * Time series for multiple datasets
 - * Conditional probability density function calculation
 - * Empirical orthogonal function calculation
 - * Random forest feature importance calculation
 - * Anomaly calculation
- * The JPL team improved these features:
 - * A new color scheme of data visualization of some services
 - * The URL for a RestFul service call has been constructed
- * The JPL team is working on
 - * supporting a workflow with CMDA services. identified first examples to work on.
 - * modifying the CMDA services to be "chainable" (callable in a workflow).
- * The CMU team created a unified github repository that has both the CMU and JPL code bases. Needs cleanup;
- * The CMU team created a new Docker container that has the latest Ubuntu operating system and the run-time environment for the CMU deployment;
- * The JPL team installed the run-time environment in the Docker container for the JPL deployment. An integrated test is pending;
- * The JPL team acquired a GPU server that will be used as the testbed for speeding up their web services;

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

Signed-off-by:

[X](climatemodeldiagnosticanalyzer) James W. Carman
[](climatemodeldiagnosticanalyzer) Chris Mattmann
[X](climatemodeldiagnosticanalyzer) Michael James Joyce
[](climatemodeldiagnosticanalyzer) Kim Whitehall
[](climatemodeldiagnosticanalyzer) Gregory D. Reddin

Shepherd/Mentor notes:

Great news! The community wrote the report this month! Thank you to Lei Pan for pushing forward and I have high hopes over the next few months we can light a fire under this project.

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. Re-engage and grow the Commons RDF community base
2. Work with key target communities (initially within Apache) to ensure Commons RDF meets potential user requirements and expectations towards next versions of the Commons API.
3. Find lateral target goals that could make the project viable.

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?

Sleeper community in the last months; very few mails and no commits.

How has the project developed since the last report?

Project's development has completely stop in the last months. We expected a different engagement from major RDF toolkits. But since that does not happen, we have started to discuss options to find a viability for the project <https://s.apache.org/oxqb>

By the next reporting period we should have something more concrete towards the future, or we could think about a possible exit strategy.

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:

[](commonsrdf) Rob Vesse
[x](commonsrdf) John D Ament
[](commonsrdf) Gary Gregory

[X](commonsrdf) Lewis John McGibbney

Shepherd/Mentor notes:

johndament: The report reflects current community engagement.

Fineract

Fineract is an open source system for core banking as a platform.

Fineract has been incubating since 2015-12-15.

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

1. Inform the non-technical part of the existing Mifos.org community of the change in project structure
2. Getting licensing issues figured out for our first release
3. Replacing dependencies that are not compliant with the license policies

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

None, currently.

How has the community developed since the last report?

Since the code was initially pushed over, the community is working on the first release.

In addition to the initial committer new contributors started to work on the project.

There is a technical conference in Amsterdam in the second week of March during which some contributors and some community members will be introduced to developing under the Apache model.

How has the project developed since the last report?

The initial code is now available and the work to meet the license policies has started.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Impala

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

Impala has been incubating since 2015-12-03.

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

1. Movement of existing JIRA / Git / wiki resources to Apache equivalents
2. Initial release as incubating project.
3. Community growth

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 have been no additions to the committer or PMC lists since incubation began. However, we have seen an uptick in external contributions, both through code, and in discussion on the mailing list. One contributor has been attempting to port Impala to PPC, and has reported some success after asking many questions.

How has the project developed since the last report?

We have made some slow progress with our initial infrastructure tasks. Code review traffic is now copied to `dev@impala.incubator.apache.org`, which means that developer discussions are now happening on the mailing lists.

We have a number of infrastructure tasks ahead of us which are blocked on the current Impala team at Cloudera being very busy with an internal release (this is one reason we look forward to a more diverse community!). For example:

1. We would like to move our Git repository to `.apache.org` in short order, but as it stands the existing repo is 10GB large and historically contains many binary artifacts that, while acceptably licensed, have no useful place in Impala's repository. We need to strip these artifacts from the Git history, and then adjust Gerrit to commit to the new branch in the new repo. This is not hard but takes some time.
2. We would also like to move our JIRA tickets from `issues.cloudera.org` to `issues.apache.org`. Experience in a sister podling has shown that this isn't straightforward if we wish to preserve existing release labels, user assignments and so on, so requires some time.

We anticipate having much more time to work on these basic issues after the end of February. We look forward to getting Impala into a position where it is easier for the larger community to collaborate on these kinds of project management issues.

Date of last release:

None since entering incubation.

When were the last committers or PMC members elected?

None since entering incubation.

Signed-off-by:

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

Joshua

Joshua is a statistical machine translation toolkit

Joshua has been incubating since 2016-02-13.

The INFRA bootstrap ticket has been opened[1], and the PPMC and INFRA are active in working through the setup tasks.

No concerns for the board at this time.

[1] (<https://issues.apache.org/jira/browse/INFRA-11264>)

Signed-off-by:

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

Shepherd/Mentor notes:

johndament: Report contents are very sparse. All sections should be retained.

Kudu

Kudu is a distributed columnar storage engine built for the Apache Hadoop ecosystem.

Kudu has been incubating since 2015-12-03.

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

1. Build a more diverse developer and user community
2. Migrate our web site under apache.org
3. Continue to make regular Apache 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?

- In February, we only received one patch from a new contributor (3 total from non-committers). We're continuing to focus on building a more diverse developer community.

Although recruiting new contributors to Kudu itself have been slow, we had community members report beginnings of a Presto integration, an Apache Flume sink, and Rust language bindings. We are considering whether it makes sense to develop these things within the Kudu project or external locations.

- Overall development velocity continues to be pretty good: 94 patches committed since the last report, representing about 3.5kloc of new code and 4kloc of changes. This is slightly down from last month, as we spent a bit more time on managing our first release and fixing bugs rather than adding new features.

Dev list traffic is similar to January (1032 messages vs 1144 in Jan).

- Community members continue to give public talks about Kudu at various conferences and meetups. Talks were delivered at conferences or meetups in Tokyo (x2), Sunnyvale, Krakow, San Francisco (x2), and Seattle. More are planned for March.

We are also encouraged that at least two Kudu meetup talks were delivered by new members of the community outside the core development team. We hope to see the "grassroots community" trend continue.

- The public Kudu Slack channel continues to be popular. It has grown from 116 members last month to 150 this month. Discussion between users and devs is quite active, with daily participation from a global community (USA, China, Poland, Japan, et al). Many users seem to prefer this method over the traditional user mailing list.

How has the project developed since the last report?

Since the last report, we've accomplished the following milestones:

- We successfully migrated our issue tracking to the Apache JIRA. Special thanks to Gavin on the infra team and Mike Percy on the Kudu project for hard work on this transition.
- We cleaned up our license documentation and successfully published our first official release as an incubating project (0.7.0)!

Date of last release:

Feb 25, 2016: Apache Kudu 0.7.0-incubating

When were the last committers or PMC members elected?

None since initial incubation.

Signed-off-by:

```
[ ](kudu) Jake Farrell
[ ](kudu) Brock Noland
[X](kudu) Michael Stack
[ ](kudu) Jarek Jarcec Cecho
[X](kudu) Chris Mattmann
[X](kudu) Julien Le Dem
[ ](kudu) Carl Steinbach
```

Shepherd/Mentor notes:

KUDU is very active. We need to get our website done and branded as an Incubator and eventually TLP site. Good progress.

log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

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

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

There's an ongoing discussion to move the project to attic unless it is able to release in the next few months.

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?

After the discussion about the activity of the project and the success of the incubation process, there's currently a member willing to step through the release process. The initial plan on incubation anyway was to primary apply available patches and release those bug fixes.

Additionally, there's once more a discussion about the supported build platforms, because currently the project heavily focusses on cpptasks, which is not maintained and simply doesn't support current versions of all formerly supported IDEs anymore. The current focus is on dropping cpptasks and ANT for the build of the software and use autoconf and/or cmake instead.

Date of last release:

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

When were the last committers or PMC members elected?

Signed-off-by:

[X](log4cxx2) Ralph Goers

Metron

Metron integrates a variety of open source big data technologies in order to offer a centralized tool for security monitoring and analysis. Metron provides capabilities for log aggregation, full packet capture indexing, storage, advanced behavioral analytics and data enrichment, while applying the most current threat-intelligence information to security telemetry within a single platform.

Metron has been incubating since 12-08-2015

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

- Building a diverse community of developers for Metron
- Getting security practitioners to provide feedback on requirements
- Make an Apache release

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

- We are currently trying to produce the first Apache release of Metron
- We would like to request Apache infrastructure (a Jenkins server) for CI

How has the community developed since the last report?

- We voted to add one additional committer, Nick Allen

How has the project developed since the last report?

- We closed on 15 more Jiras and made our code base more compliant with Apache Licenses. We are at the latest stages of cleaning up the code base to come up with the Apache build.

Signed-off-by:

[X](metron) Billie Rinaldi
[](metron) Chris Mattmann
[](metron) Owen O'Malley
[X](metron) P. Taylor Goetz
[](metron) Vinod Kumar Vavilapalli

Milagro

Distributed Cryptography; M-Pin protocol for Identity and Trust

Milagro has been incubating since 2015-12-21.

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

1. Import main project repos
2. Work towards our first release
3. Broaden the community

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

aware of?

The Initial Committer list is being revised (see below).

How has the community developed since the last report?

We have identified shortcomings in the Initial Committer list, and adopted a fast-track process to add new members, valid for the project's first three months (until March 20th). We have welcomed our first genuinely new committer and are inviting another. Initial committers are familiarising with Apache.

How has the project developed since the last report?

Both companies MIRACL and NTT are now readying their initial contributions into the Milagro repositories.

Date of last release:

n/a

When were the last committers or PMC members elected?

We have no distinction between committers and PMC.

The following have accepted invitations during the period:

Peter Scanlon
Tomasz Matecki
Anthony Shaw
Vladislav Mitov
Nikolai Stoilov

All except Anthony Shaw are Miracl employees, and should have been on the Initial Committer list. We anticipate improving diversity soon.

Signed-off-by:

[](milagro) Sterling Hughes
[](milagro) Jan Willem Janssen
[x](milagro) Nick Kew

MRQL

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

MRQL has been incubating since 2013-03-13.

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

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

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

None

How has the community developed since the last report?

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

How has the project developed since the last report?

We have released our fourth release under Apache incubation. There were various bug fixes and performance improvements before the release. The

most important new feature in the new release is support for incremental query processing in MRQL streaming, called Incremental MRQL.

Date of last release:

2016-03-02

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

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

Myriad

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

Myriad has been incubating since 2015-03-01.

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

1. Grow the Myriad user base.
2. Expand the contributor community.
3. Release Myriad 0.2.0 in April/May, with focus on multitenancy, security, and oversubscription.

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 traffic went into a holiday lull after the 0.1.0 release. Should pick up again as we plan the 0.2.0 release.
http://mail-archives.apache.org/mod_mbox/incubator-myriad-dev/
- Our biweekly community syncs still get 6-12 participants, including new participants from Accenture. See minutes at: <http://s.apache.org/8kF>
- Freenode #myriad IRC channel has been mostly quiet since the 0.1.0 release, except for an occasional user question or discussion around a community sync. Chatter should pick up again leading up to the 0.2.0 release, and moreso after we get a few more users. Archived by ASFBot at: <http://wilderness.apache.org/channels/#logs-#myriad>
- Website and wiki were updated with the 0.1.0 release.
<http://myriad.incubator.apache.org>
- Presented at DockerConEU. Submitted talks for ApacheCon, MesosCon, Hadoop Summit, Container World.

How has the project developed since the last report?

- Resolved 6 JIRAs and merged 6 PRs since Dec 2, 2015.

Date of last release:

2015-12-10 myriad-0.1.0-incubating

When were the last committers or PMC members elected?

2015-10-05 Darin J
2015-10-14 Swapnil Daingade
In the process of nominating the next committer.

Signed-off-by:

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

SAMOA

SAMOA provides a collection of distributed streaming algorithms for the most common data mining and machine learning tasks such as classification, clustering, and regression, as well as programming abstractions to develop new algorithms that run on top of distributed stream processing engines (DSPEs). It features a pluggable architecture that allows it to run on several DSPEs such as Apache Storm, Apache S4, and Apache Samza.

SAMOA has been incubating since 2014-12-15.

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

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

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

None

How has the community developed since the last report?

Mailing list activity (December 2015 - February 2016):

* @dev: 114 messages

Jira issues backlog (December 2015 - February 2016):

* Created: 7

* Resolved: 4

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

How has the project developed since the last report?

The community has decided to drop support for S4, which is now inactive. There is also some continued interest by Apache Apex about a possible integration. Internally, we have been working on strengthening our main classifier (VHT). There has also been work on improving the connection with Avro and Kafka. The adapted for Flink has been updated to the latest stable version (0.10).

Date of last release:

2015-07-21

When were the last committers or PMC members elected?

None

Signed-off-by:

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

Shepherd/Mentor notes:

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. Graduation vote has been passed, see below for details.
- 2.
- 3.

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

1. Graduation vote has been passed and the resolution has been proposed to the board for consideration in the board meeting.
References:
https://s.apache.org/dev_discuss
https://s.apache.org/dev_vote_result
https://s.apache.org/general_discuss
https://s.apache.org/Incubator_vote
Other:
https://s.apache.org/general_notify
<https://s.apache.org/8rty> (Maturity Model assessment)
2. Olivier Lamy resigned as a mentor.

How has the community developed since the last report?

1. We had ~500 messages on dev list in the past three months. (Got number from <http://markmail.org/list/org.apache.sentry.dev>)
2. Community is continuing to do monthly hangouts.

How has the project developed since the last report?

About 131 issues were created and about 87 resolved in the past 3 months (Numbers from jira).

Date of last release:

2015-09-22, preparing for 1.7.0 release

When were the last committers or PMC members elected?

Hao Hao was added as committer on 01/29/2016. On 2015-12-21, community voted to make Committer == PPMC.

Signed-off-by:

[X](sentry) Arvind Prabhakar
[](sentry) Joe Brockmeier
[](sentry) David Nalley
[X](sentry) Patrick Hunt
[X](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. Improve the system, including performance and features like distributed GPU training over GPU clusters, R binding, etc.
2. Add more examples and documentation to attract users
3. Grow the community. We have been trying to reach out to a more diverse community of developers and users by presenting SINGA for various

audiences. Completing items 1 and 2 above would attract more contributors and users.

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

No

How has the community developed since the last report?

The Github repository has scored 540 stars and 190 forks.

There were 60, 180 and 7 commits in Dec 2015, Jan 2016 and Feb 2016 respectively.

There were 77, 105, 13 emails from dev@ list in Dec 2015, Jan 2016 and Feb 2016 respectively.

How has the project developed since the last report?

In Jan 2016, we released our second version, i.e., incubating-v0.2.0, with the following major features:

- * GPU support, which enables training of complex models on a single node with multiple GPU cards.
- * Hybrid neural net partitioning, which supports data and model parallelism at the same time.
- * Python wrapper, which makes it easy to configure the job, including neural net and SGD algorithm.
- * RNN model and BPTT algorithms, to support applications based on RNN models, e.g., GRU.
- * Cloud software integration includes Mesos, Docker and HDFS.

Date of last release:

2016-01-14

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Streams

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

Streams has been incubating since 2012-11-20.

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

1. Community growth and PMC maturity.
2. Demonstrate a consistent release schedule.
3. 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 remained low this quarter.

How has the community developed since the last report?

no demonstrable progress in community growth.

How has the project developed since the last report?

5 issues created / 1 resolved this quarter
at mentors suggestion, this quarter the project focused on
documenting project purpose, philosophies, place in ecosystem,
how to get started, etc... as part of a web site overhaul.
the new version is hosted at streams.staging.apache.org and will go
live at streams.incubator.apache.org soon.

Date of last release:

2015-04-26

When were the last committers or PMC members elected?

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

Signed-off-by:

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

Taverna

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

Taverna has been incubating since 2014-10-20.

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

1. Community growth
2. Release more of the imported code to create engagement
3. Move to a sustainable development pattern

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

Our mentors have reported they have little or no time
available as Taverna is picking up pace.

Dan Debrunner volunteered from the incubator list and
became a new mentor.

How has the community developed since the last report?

Welcomed new PPMC member.

GSOC 2016 project ideas are already getting interest.

dev@taverna mailing lists stats:

Dec 2015: 42
Jan 2016: 128
Feb 2016: 325

users@taverna:

Dec 2015: 2
Jan 2016: 0
Feb 2016: 2

JIRA issues over last 90 days:

50 created
32 resolved

How has the project developed since the last report?

Good buzz around release preparation, release candidate currently under vote.

Date of last release:

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

(new release currently under vote)

When were the last committers or PMC members elected?

2015-12-04

Signed-off-by:

[x](taverna) Andy Seaborne
[X](taverna) Chris Mattmann
[](taverna) Suresh Srinivas
[X](taverna) Suresh Marru
[X](taverna) Marlon Pierce
[](taverna) Dan Debrunner

Shepherd/Mentor notes:

Andy:: It would be good to have a second active mentor in addition to Dan.
Chris: I have to retire as a Taverna mentor. That said - what's stopping Taverna from being ready to graduate? If it's just mentor activity then consider graduating. The community looks active, looks like it's releasing, looks like involved in GSoC, etc.

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. Removal of LGPL dependency
2. Ease of installation
3. Documentation of APIs

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. Work with Jupyter community to update their docker stacks to include the latest Toree
2. Blog post at <http://www.spark.tc/announcing-apache-toree/> to announce the incubation
3. Active communication in mailing list and gitter with early adopters
4. Work with Brunel project (<https://github.com/Brunel-Visualization/Brunel>) to make technology available on Toree
5. Still working on transitioning users from Spark Kernel project into Toree.

How has the project developed since the last report?

1. Source package have been renamed to org.apache.toree
2. All original contributors are not actively developing and fixing issues on the apache project
3. Project can be published and installed through Pypi (only snapshot/dev packaged made available for testing)
4. Added a plug-in framework to enable integration of 3rd party code to add capabilities
5. Create Spark Streaming example notebook to feature many of the capabilities of Toree in collaboration with Jupyter ecosystem (dashboards and widgets)
6. Misc bug fixes and features
7. Actively working with JeroMQ community to further their transition into MPL v2 and away from LGPL.

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
[X](toree) Julien Le Dem

Trafodion

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

Trafodion has been incubating since 2015-05-24.

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

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

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

* Lars Hofhansl and Andrew Purtell resigned as mentors.

How has the community developed since the last report?

- * Participation in public lists has been increasing: 809 messages in the codereview forum, 1051@commits, 561@dev, 1377@issues and 320@user.
- * Steve Varnau has agreed to be the release manager for our next release.
- * A logo for Apache Trafodion was adopted by vote in February.

How has the project developed since the last report?

- * Completed our first release in incubation, Trafodion 1.3.0, on Jan 11, 2016. Many thanks go to Roberta Marton, the release manager. See <http://trafodion.apache.org/release-notes-1-3-0.html> for details.
- * Completely redesigned the web page, <http://trafodion.apache.org> and moved content from the wiki to the site, following the example of other Apache projects.
- * Published documentation in form of HTML and PDF manuals.
- * 287 commits from 23 contributors.

* 200 JIRAs filed and 139 resolved in the last three months.

Date of last release:

2016-01-11

When were the last committers or PMC members elected?

RuoYu Zuo was announced as a new committer in February.
The last PMC members were announced in November 2015.
In both cases the election happened a few weeks earlier
but it took some time to get the necessary karma.

Signed-off-by:

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

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. Prove PPMC's maturity by handling dispute in community gracefully
2. Community growth: more committers from different organizations

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

Updating long term roadmap [1] to embrace community needs is in
discussion.

[1] <https://s.apache.org/Ksgm>

How has the community developed since the last report?

Traffic in users@ and dev@ grew: +38 subscribers in @dev, +83 subscribers
in users@ since last report.

Talks at several conferences are scheduled [2].

[2] <https://cwiki.apache.org/confluence/display/ZEPPELIN/Community>

How has the project developed since the last report?

209 patches have been merged from 42 different contributors since last
report.

Integration to JDBC, Apache HBase, Elasticsearch, Scalding, Tachyon has
been added.

Authentication and notebook authorization have implemented.

Date of last release:

2016-01-22

When were the last committers or PMC members elected?

2016-01-16

Signed-off-by:

[](zeppelin) Konstantin Boudnik

[X](zeppelin) Henry Saputra
[X](zeppelin) Roman Shaposhnik
[X](zeppelin) Ted Dunning
[](zeppelin) Hyunsik Choi