

May2016

/\ *FINAL* (!)

Incubator PMC report for May 2016

Timeline

Wed May 04	Podling reports due by end of day
Sun May 08	Shepherd reviews due by end of day
Sun May 08	Summary due by end of day
Tue May 10	Mentor signoff due by end of day
Wed May 11	Report submitted to Board
Wed May 18	Board meeting

Shepherd Assignments

Drew Farris	Slider		
John Ament	Kudu		
Justin Mclean	Open Az		
P. Taylor Goetz	S2Graph		
P. Taylor Goetz	Tamaya		
Timothy Chen	Beam		
Timothy Chen	Eagle		
Timothy Chen	Mnemonic		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5bb019c5-8ba0-49ce-8a2c-a1060765959f"><ac:plain-text-body><![CDATA[[none]	Airflow]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="cd44060a-ac97-4301-a801-a428f5fb0162"><ac:plain-text-body><![CDATA[[none]	Blur]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d141bfb0-6428-4fd5-b417-ce246b831db9"><ac:plain-text-body><![CDATA[[none]	Concrted]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="4a83fb3d-db52-43eb-baf2-0ce38c1b37fd"><ac:plain-text-body><![CDATA[[none]	Gearpump]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="84224efc-b1af-4d21-9650-28215e181c4c"><ac:plain-text-body><![CDATA[[none]	Guacamole]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="60b38a82-4ba8-4476-9485-f9972256d846"><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="f7807d48-f592-47d9-b707-04fa370d3cac"><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="994f4766-3be5-4a68-afea-6921061eb1ba"><ac:plain-text-body><![CDATA[[none]	Quickstep]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="018df290-6664-4900-aa9e-f95a146290aa"><ac:plain-text-body><![CDATA[[none]	SystemML]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="dfd16962-d39b-4e5b-892c-028cc57f62af"><ac:plain-text-body><![CDATA[[none]	Tephra]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="6e3b7e30-13cb-4bde-a1fd-c7980e399a51"><ac:plain-text-body><![CDATA[[none]	TinkerPop]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a7707b10-8a28-400a-925a-5a2d10fbd22d"><ac:plain-text-body><![CDATA[[none]	Twill]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="eee7a379-5982-4230-a662-53ef3b13a5c6"><ac:plain-text-body><![CDATA[[none]	Unomi]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b7fc8bf5-4f72-4c9f-96ea-56fcb16db435"><ac:plain-text-body><![CDATA[[none]	iota]]></ac:plain-text-body></ac:structured-macro>
--	--------	------	--

Report content

Incubator PMC report for May 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 54 podlings currently undergoing incubation. We have two podlings planning to graduate this month, had one podling retire and added one podling to the roster. April was a generally quiet month, losing a single IPMC member, adding none, and completing nine releases.

* Community

New IPMC members:

- None

People who left the IPMC:

- Rob Vesse

* New Podlings

- Gossip

* Retired Podlings

- Concerted

* Graduations

The board has motions for the following:

- TinkerPop
- Zeppelin

* Releases

The following releases entered distribution during the month of April:

- 2016-04-07 Apache MADlib 1.9-incubating
- 2016-04-11 Apache Eagle 0.3.0-incubating
- 2016-04-12 Apache Metron 0.1BETA-incubating
- 2016-04-15 Apache Tamaya 0.2-incubating
- 2016-04-19 Apache Singa incubating-0.3.0.tar.gz
- 2016-04-20 Apache TinkerPop 3.1.2-incubating
- 2016-04-20 Apache TinkerPop 3.2.0-incubating
- 2016-04-20 Apache Apex Core v3.2.1-incubating
- 2016-04-21 Apache Geode 1.0.0-incubating.M2

* Infrastructure

- Kudos to the infra team for assisting with troubleshooting the new Marvin Jr. Podling reminder email script.

* Credits

- Report Manager: Marvin Humphrey/John Ament

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

* Still getting started at the Incubator

- Airflow
- Beam

- Gearpump
- Mnemonic
- Quickstep
- Tephra

* Not yet ready to graduate

No release:

- Fineract
- Guacamole
- iota
- S2Graph
- Toree

Community growth:

- Blur
- Eagle
- Joshua
- Kudu
- OpenAz
- Quarks
- Slider
- SystemML
- Tamaya
- TinkerPop
- Twill
- Unomi

* Ready to graduate

The Board has motions for the following:

- TinkerPop
- Zeppelin

* Did not report, expected next month

- Climate Model Diagnostic Analyzer

Table of Contents

Airflow
 Beam
 Blur
 Eagle
 Fineract
 Gearpump
 Guacamole
 iota
 Joshua
 Kudu
 Mnemonic
 OpenAz
 Quarks
 Quickstep
 S2Graph
 Slider
 SystemML
 Tamaya
 Tephra
 TinkerPop
 Toree
 Twill
 Unomi

Airflow

Airflow is a workflow automation and scheduling system that can be used to author and manage data pipelines.

Airflow has been incubating since 2016-03-31.

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

1. Continued releases
2. Grow up user and contributor communities
3. Improve and extend documentation and samples on the 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?

- * Since we applied to the incubator in March 17 (proposal email), we grew our contributors from 100 to 137
- * Since we applied to the incubator in March 17 (proposal email), we accepted/merged 119 PRs
- * Dev mailing lists activity increased

How has the project developed since the last report?

- * All resources have been created (website, Jira, git & github mirror, ...) but not all migrations are complete
- * The code donation (and Pull Request migration) is in progress
- * The website is in progress
- * The previous wiki has been migrated to confluence, Jira is now used instead of the old Github issues, and the old google group is being EOLd on May 2 as users migrate to the dev mailing list

Date of last release:

N/A - we are still working out releases as we are new to Apache.

When were the last committers or PMC members elected?

N/A - we are still new to Apache and have not gotten to the point of electing new committers/PMC members.

Signed-off-by:

[](airflow) Chris Nauroth
[X](airflow) Hitesh Shah
[x](airflow) Jakob Homan

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). Dataflow pipelines simplify the mechanics of large-scale batch and streaming data processing and can run on a number of runtimes like Apache Flink, Apache Spark, and Google Cloud Dataflow (a cloud service). Beam also brings DSL 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. Continued releases
2. Grow up user and contributor communities
3. Improve and extend documentation and samples on the 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?

- * Both user and dev mailing lists activity increased
- * We sustain a high level of activity on the pull request cycles (submit, review, ...)

How has the project developed since the last report?

- * All resources have been created (website, Jira, git & github mirror, ...)
- * The code donation has been completed
- * The website has been published, we are still in the process of donated documentation and samples resources
- * We renamed all package to match the Apache convention
- * We started the re-organization and refactoring of the project structure (isolating and moving some modules)

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](beam) Jean-Baptiste Onofre
[](beam) Jim Jagielski
[X](beam) Venkatesh Seetharam
[](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[+5]; dev@ - 78[+6]
- The community involvement has not really changed over the past few months.

How has the project developed since the last report?

- Software bug fixes and performance enhancements.
- An announcement was made of a production release of Apache Blur (very cool!)
http://mail-archives.apache.org/mod_mbox/incubator-blur-user/201604.mbox/3CCAGW2whTUEFLYCUeyAdT3_KH9FcZdrBJ6SLV3YMuUSA-4XBfR%3DA%40mail.gmail.com%3E

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

Concerted

Shepherd/Mentor notes:

Julian Hyde (jhyde):

Concerted voted to retire, and on 2016-04-25 the IPMC voted to retire Concerted.

Retirement steps are complete;
see <https://issues.apache.org/jira/browse/INFRA-11752>

Eagle

Eagle is a Monitoring solution for Hadoop to instantly identify access to sensitive data, recognize attacks, malicious activities and take actions in real time.

Eagle has been incubating since 2015-10-26.

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

1. Release of Apache Eagle 0.3v
2. Growing the community by adding contributors from different areas
3. We are working on making changes to the public site to include more developer friendly information

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?

We have companies in banking, gov and other domains starting to use the product and are actively contributing.

- Presented at QCon Software Conference Beijing 2016
- Shared at Shanghai Java User Group

How has the project developed since the last report?

We presented at the Hadoop Summit in Dublin and garnered more interests from different companies. We are having more outreach by participating in conferences and we are presenting in Apache Con, Strata + Hadoop World.

Date of last release:

2016-03-28

When were the last committers or PMC members elected?

NIL

Signed-off-by:

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

Shepherd/Mentor notes:

Julian Hyde (jhyde):

The project made its first release this quarter, which is a great achievement. As such, release 0.3 should no longer be listed as a goal, but nevertheless it is good that they are aiming to show that the release process is repeatable.

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. Finalizing the initial release
2. Addressing with the VP of Legal our Hibernate dependency and how to remedy (removal, exception, etc.)
3. Continue to encourage participation in Fineract collaboration and communication infrastructure by bringing offline interactions online and directing collaboration to occur in right channels.

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?

- Participated in and selected one intern for Google Summer of Code working on the India Unified Payments Interface.
- Community members attended LinuxFest Northwest in Bellingham, WA and encouraged developers to get involved in the Apache Fineract community.
- Discussion of requirements is taking place via mailing lists for major enhancements community is anticipating over next quarter.
- Jim Jagielski has joined the ranks of Fineracts mentors.

How has the project developed since the last report?

More communications are happening on the appropriate lists as we have been directing conversations to happen in the public mailing lists and outside of previous community silos.

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
[](fineract) Jim Jagielski

Shepherd/Mentor notes:

Gearpump

Gearpump is a reactive real-time streaming engine based on the micro-service Actor model.

Gearpump has been incubating since 2016-03-08.

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

1. Make initial Apache branded release and release on a regular schedule.
2. Build community interest and support.
3. Complete apache branded website with compelling examples and roadmap.

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?

- High amount of activity and discussion on dev mailing list (april 604)
- Increased participation and interest from developers to provide new sources and examples

How has the project developed since the last report?

- All outstanding issues from old github repo have been reopened as JIRAs.
- All prior merges in old repo have been converted to PR's and merged to master within apache gearpump repo.
- Branded apache website is completed and reviewed. Publishing to <https://s.apache.org/T86l>
- INFRA tickets required to publish to repository.apache.org completed, travis build scripts updated

Date of last release:

No release yet.

When were the last committers or PMC members elected?

No new committers or PMC members elected yet.

Signed-off-by:

```
[X](gearpump) Andrew Purtell
[ ](gearpump) Jarek Jarcec Cecho
[ ](gearpump) Reynold Xin
[ ](gearpump) Todd Lipcon
[ ](gearpump) Xuefu Zhang
```

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. Finalizing the move to Apache infrastructure
2. Making the first Guacamole release under the Apache Incubator
3. Accepting additional committers

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

The PPMC has had difficulty fully engaging mentors to finish the migration of Apache Guacamole to Apache infrastructure and processes.

Though the project has successfully moved to Apache-hosted git repositories, the processes for importing our existing JIRA issues and adding our first non-founding committer (Frode Langelo) were stagnated until very recently.

How has the community developed since the last report?

While our community of users continues to grow, this is mainly the established pattern of growth that existed prior to acceptance in the Apache Incubator. We are expecting this to change once the Apache Guacamole website is up, public announcement of the move to the Incubator is made, and users of the old SourceForge forums are directed to the mailing lists.

How has the project developed since the last report?

The project has accepted its first committer outside the original founding developers, Frode Langelo, and development has moved onto Apache-hosted git repositories. Development of new Apache Guacamole website is nearing completion.

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

Shepherd/Mentor notes:

Jean-Baptiste Onofre (jbonofre):

Frode's account should be OK now. For the Jira, the difficulty was that the project wanted to import existing Jira content. Not sure it's a good idea as bunch of users don't exist in the Apache Jira and there's lot of "noise". I would recommend to start by creating the highest priority tasks and issues in the Apache Jira.

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. Transitioning code, website and Jira issues to the Apache repositories.
2. Continuing iota website and code development using the Apache infrastructure.
3. Building an iota community via a variety of means including presentations at series of conferences beginning with ApacheCon followed by 7x24 Spring Conference, Data center Dynamics and others throughout the summer. We will work with our mentors to find other means of community building.

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

This project has been slower to start than we anticipated. Why? Based on experience in the field with we questioned the scalability of our

original architecture. We made the difficult decision that we needed a more scalable architecture and we on the design of a baseline architecture. WE now feel we have a good starting point and with the help of the iota community we can move forward with.

How has the community developed since the last report?

None

How has the project developed since the last report?

None

Date of last release:

None

When were the last committers or PMC members elected?

None yet.

Signed-off-by:

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

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?

Joshua is being represented at ApacheCon. Currently one presentation is taking place in addition to a meetup.

How has the project developed since the last report?

The project team have been working on development code as well as official branding. All infrastructure has been migrated over to Apache. We are aiming for our first release some time after ApacheCon.

Date of last release:

N/A

When were the last committers or PMC members elected?

Kellen Sunderland and Felix Hieber joined the PPMC and as committers on April 11th, 2016.

Signed-off-by:

[](joshua) Paul Ramirez

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

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?

The PPMC invited Binglin Chang as a new committer and PPMC member. Binglin has continued to increase his involvement and some of his colleagues are now becoming more involved as well.

Since the last report, we had contributions from several new contributors, including five new contributors from outside Cloudera (the initial code grantor). Of those, two were substantial pieces of new functionality: a new sink to ingest data from Apache Flume, and an improved integration with Apache Spark. We hope that the authors of these contributions will continue to be involved in the community and grow into committers.

In March, one of our mentors approached the development community with some ideas to make the project more approachable for developers. As a result of this, we've been working on a few action items:

- 1) We started a project blog and have been posting weekly updates on development progress as well as more detailed posts on new features or project news. We've heard from community members that this is an easy way to track the project without subscribing to high-volume dev lists.
- 2) Our dev@ list is currently very noisy with a lot of code review comments. For more casual contributors this can be overwhelming as code review is intermingled with higher level project discussions. The community discussed the issue and decided to split the code review traffic to a new mailing list (currently awaiting infra to create it)
- 3) We've started an informal policy of emailing the dev list with a heads-up whenever any new larger feature first shows up on gerrit or JIRA, to make it easier for people to give input on large items without having to see every bug fix or small improvement. For example, whenever design discussion starts on a new feature, the contributor sends a 'heads-up' email to dev@.

Development activity continues to be healthy. Mailing list traffic for the dev list in April and May was similar to the preceding months. Patch contribution level is also fairly stable (50-80 commits/month) with lulls around releases as people focus on release testing, documentation, etc.

User mailing list traffic is steadily increasing in quantity and diversity: 108 messages in April vs 51 in March, 40 in Feb. April saw 26 distinct people contribute to the user list discussion vs 15 in Feb. Website traffic also indicates growing interest, with April having 30% more traffic than Feb.

In the last two months, Kudu-related talks were delivered in San

Francisco, San Jose, Beijing, Boston, DC, Detroit, London, and Dublin. More talks are scheduled for the upcoming months at ApacheCon Big Data, Strata/Hadoop World, Berlin Buzzwords, and various local meetups.

How has the project developed since the last report?

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

- We released Apache Kudu 0.7.1-incubating on 3/9/2016 to fix a few issues found in the 0.7.0 release.
- We released Apache Kudu 0.8.0-incubating on 4/10/2016 with various new features and improvements.

Date of last release:

April 10, 2016: Apache Kudu 0.8.0-incubating

When were the last committers or PMC members elected?

- Binglin Chang was added as a committer and PPMC member on April 4, 2016.

Signed-off-by:

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

Mnemonic

Mnemonic is a Java based non-volatile memory library for in-place structured data processing and computing.

Mnemonic has been incubating since 2016-03-03.
Mnemonic code was moved into ASF on 3/24.

Various metrics for April are as follows:

Metric	counts
Non Merge Commits	16
Contributors	4
Jira New Issues	20
Resolved Issues	24
Pull Requests merged	6
Pull Requests proposed	7

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

1. Apache releases: we are preparing 0.1.0 release
2. Expand and build strong community, keep active dev list discussions, adding new contributors
 - a. Mnemonic website is up: <http://mnemonic.incubator.apache.org/>
 - b. Mnemonic cwiki is up: <https://cwiki.apache.org/confluence/display/MNEMONIC/Apache+Mnemonic+Home>
3. Develop new features: We are working in Native computing layer design.

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

Everything is on track, nothing urgent at this time.

How has the community developed since the last report?

1. We completed setting up infrastructure. Both project website and cwiki are up running. We are working on adding contents and documentation to the sites so new contributors and community members can get better view of what this project is and how to get involved.
2. We realized, unlike other application level projects, library (or API) projects traditionally had less contributors until certain amount of adopters successfully integrated those libraries. We plan to help some integrations, such as HBase, Arrow, etc. after our first official code release.
3. We submitted 2 proposals for Strata Hadoop World NYC.
4. We will present Mnemonic project in upcoming ApacheCon Big Data conference on 5/11.

How has the project developed since the last report?

We are designing the native computing infrastructure, built native computing interface code.

In preparing first official release:

1. Defined coding guidelines for the project. Reformatted all code based on defined coding style.
2. Added the checkstyle to automate the checking process.
3. Formalized the names of annotations, classes and methods.
4. Removed the dependency of property mnemonic.core.version.
5. Added a new profile for artifacts releasing and signing.
6. Added release plugin and configuration for releasing.

Date of last release:

No release yet, working on 0.1.0 release

When were the last committers or PMC members elected?

No election, since the project is still in early stage of developing.

Signed-off-by:

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

 OpenAz

OpenAz is a project to create tools and libraries to enable the development of Attribute-based Access Control (ABAC) Systems in a variety of languages. In general the work is at least consistent with or actually conforming to the OASIS XACML Standard.

OpenAZ has been incubating since 01-2015.

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

1. Increasing community involvement: We are focusing on getting information published on the website to aid the acquisition and orientation of new contributors, as well increasing the number of committers, in order to reduce individual member workload, reduce reliance on any particular member, and to increase project velocity. David Ash was voted in as a committer last month. There are several other committers that may be added in the near future.
2. Merging the source code from AT&T, JP Morgan and Oracle and creating POMs. AT&T and JP Morgan source code has been merged and POMs created, it needs uploading into the Apache repository. AT&T and JP Morgan ICLAs were received, but Oracle is taking more time. The IP clearance is on it's way.

3. Get our first release done.

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

None. The efforts to revive the project seem to be paying off and discussions are gaining traction.

How has the community developed since the last report?

Pam Dragosh, David Ash and Ajith Nair are the active committers at this point. The project intends to reboot the PPMC which may include removing some of the initial committers who are no longer active. There are, however, several contributors who could be potentially added as committers in the future.

How has the project developed since the last report?

1. Ongoing collective effort to improve project documentation, marketing material for better visibility and adoption.
2. A pending contribution towards PEP implementation is being reviewed and we plan to merge in these changes soon.
3. We are looking into design feedback/comments from Nextlabs - a XACML service provider. We hope to incorporate many of these in the coming weeks.

Date of last release:

N/A

When were the last committers or PMC members elected?

One new committer was elected Tuesday, May 3rd.

We have several interested folks in working on the project:

- Yannick Grignon
- Farasath Ahamed
- Inge Olaisen
- Carlos Perez
- Remon Sinnema

Signed-off-by:

[](openaz) Emmanuel Lecharny
[](openaz) Colm O Heigeartaigh
[](openaz) Hadrian Zbarcea

Quarks

Apache Quarks is a programming model and micro-kernel style runtime that can be embedded in gateways and small footprint edge devices enabling local, real-time, analytics on the continuous streams of data coming from equipment, vehicles, systems, appliances, devices and sensors of all kinds (for example, Raspberry Pis or smart phones). Working in conjunction with centralized analytic systems, Apache Quarks provides efficient and timely analytics across the whole IoT ecosystem: from the center to the edge.

Quarks has been incubating since 2016-02-29.

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

1. Create and expand a diverse community of contributors and committers around Quarks project
2. Create the first Apache release of Quarks
3. Complete migration of website to host all content at Apache (javadoc is still on Github).

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

Nothing jumps to mind.

How has the community developed since the last report?

We have added a new committer and PPMC member, Cazen Lee. Cazen started making valuable contributions almost immediately when Quarks went into incubation. The Quarks community is still small, as Quarks has just recently incubated, therefore we aim to continue growing our list in the near future.

How has the project developed since the last report?

- * According to JIRA, the project has added 59 new issues in April, and 34 issues have been resolved.
- * On April 13th, the committers hosted an "Introduction to Quarks" webcast where they presented a high-level overview of the project and demonstrated how to write, monitor and debug Quarks applications.

Date of last release:

Quarks is still new, so we haven't created an Apache release as of yet.

When were the last committers or PMC members elected?

In April we added a new committer and PPMC member, Cazen Lee.

Signed-off-by:

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

Quickstep

Quickstep is a high-performance data processing platform based on a core relational kernel.

Quickstep has been incubating since 2016-03-29.

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

1. Find, nurture, and collaborate with an expanded cadre of contributors from different organizations, electing the most dedicated as additional committers.
2. Demonstrate open collaboration and decision making among contributors, committers and the PPMC in The Apache Way
3. Create one or more usable official releases of Apache Quickstep.

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?

This is the first report for this project.

How has the project developed since the last report?

This is the first report for this project.

Date of last release:

N/A as the project was just incubated.

When were the last committers or PMC members elected?

April 2016.

Signed-off-by:

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

Shepherd/Mentor notes:

Julian Hyde (jhyde):

Most of the initial committers have submitted CLAs and joined the mailing lists, and activity on the dev list is looking healthy.

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?

- * There has been concerns about s2graph community, and we have decided that committers in the same company will avoid face-to-face discussion and start out on the dev mailing list. we opened up 11 discussions on the dev mailing list afterwards.
- * There are two new contributors, one is working on packaging and other is working on bug fix.
- * We will present about Apache S2Graph in apache-big-data-north-america.

How has the project developed since the last report?

- * We started out the discussion on first release plan and project road map.
- * We opened up the discussion on the project logo.
- * 45 issues are created, 31 issues are resolved, 34 commits merged.
- * We deprecated our old github repository under kakao and made it point to ASF github repository.

Date of last release:

No yet

When were the last committers or PMC members elected?

No

Signed-off-by:

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

Slider

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

Slider has been incubating since 2014-04-29.

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

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

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

No.

How has the community developed since the last report?

We continue to see regular interaction with users who are trying out Apache Slider (incubating). Questions range from introductory "I'm trying to do \$x" to bug reports "I did \$a and encountered an issue \$b". The committers try their best to keep people involved and steer them toward contributions. The strong majority of contributions continue to come from existing podling members.

How has the project developed since the last report?

The IP clearance for the Apache Kafka Slider "app-package" process was addressed (after discussion on general@i.a.o) and this codebase was merged in. This was a great contribution by a few external developers and very much appreciated.

Work continues on "phase 2" of some long-term work to include support for running Docker containers in Slider. The intent behind this work is that we will make it easier for (prospective) users to build applications that can run on Slider. There is a tentative patch on SLIDER-906 and some related documentation updates.

We have agreed upon a next release (0.91) happening in May, but have not made any significant progress towards that release yet.

We had a discussion put forth about the future of Slider by a podling member. At the inception of Slider, Apache Hadoop's YARN was not ready to support long-lived services "natively". While Slider has been in incubation, YARN has continued to grow towards treating these services natively, partly as a result of requirements necessary of Slider to work. As such, we discussed whether it would make sense to ask the YARN project if they would be interesting in combining efforts, moving the core of Slider into YARN itself. The Slider community reacted positively, noting that YARN is likely a better long-term home for the core of Slider; however, there were a few points raised that would need to be addressed before any action was actually taken. The mail discussion can be found at <https://s.apache.org/3ywX>

Date of last release:

2016-01-09 slider-0.90.2-incubating

When were the last committers or PMC members elected?

2015-07-07: Yu (Thomas) Liu

Signed-off-by:

[](slider) Arun C Murthy
[](slider) Devaraj Das
[X](slider) Jean-Baptiste Onofré
[](slider) Mahadev Konar

Shepherd/Mentor notes:

Drew Farris (drew)
- No apparent mentor activity on the dev/private mailing lists over the past few months for the Slider podling. Active committers and a reasonable amount of

JIRA and mailing list activity.

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 February through April had 199 messages involving topics such as algorithms, DML functionality, usability, and bug fixes. In addition, we have had many discussions on our JIRA site and in pull request conversations. Fred Reiss presented at Spark Summit East on February 17 about SystemML internals. Berthold Reinwald spoke at the Spark Technology Center on March 9 about scalable machine learning with SystemML. Niketan Pansare spoke on April 28 at Datapalooza in Austin and on April 29 at Rice University about declarative machine learning at scale with SystemML. Researchers in Germany are working to add Flink as an additional SystemML backend. On GitHub, the project has been starred 267 times and forked 92 times. We have gained two new contributors since February 1st.

How has the project developed since the last report?

We produced our first Apache release, version 0.9.0-incubating. Numerous additions have been made to the project, including core functionality, usability improvements, and documentation. The project has had 204 commits since February 1. In the same time frame, 155 new issues have been reported on our JIRA site and 77 issues have been resolved. 114 pull requests opened since February 1 have been closed. We are preparing to produce our second incubator release.

Date of last release:

2016-02-15 (version 0.9.0-incubating)

When were the last committers or PMC members elected?

NONE

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. We released Tamaya 0.2-incubating on 2016-04-06.
2. Attract more community members.
3. Create a fully ASF compliant release (both source and binary, since the source one was good this time around)

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

There are no such items.

How has the community developed since the last report?

- We elected a new committer.
- Some people expressed their interest to use Tamaya in their projects.

How has the project developed since the last report?

- Release 0.2-incubating has been released.
- New features based on community input have been added to the roadmap.

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
[](tamaya) Mark Struberg
[](tamaya) Gerhard Petracek
[](tamaya) David Blevins

Tephra

Tephra is a system for providing globally consistent transactions on top of Apache HBase and other storage engines.

Tephra has been incubating since 2016-03-07.

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

1. Complete project setup
2. Regular releases
3. Improve community engagement

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

- None at this time.

How has the community developed since the last report?

- Mailing lists have been created.
- Working on the project page.

How has the project developed since the last report?

- Project is in initial setup phase.

- SGA has been submitted, and git repo is set up.
- Waiting for other bootstrapping tasks like JIRA import to be done.
- Working towards the first release.

Date of last release:

None

When were the last committers or PMC members elected?

- None since coming to incubation

Signed-off-by:

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

TinkerPop

TinkerPop is a graph computing framework written in Java

TinkerPop has been incubating since 2015-01-16.

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

The outstanding issues have been addressed.

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

Nothing at this time.

How has the community developed since the last report?

We have 4 new committers: Ketrina Yim, Pieter Martin, Jean-Baptiste Musso, and Michael Pollmeier

How has the project developed since the last report?

Since the last status report, 172 new JIRA issues have been created and 126 have been resolved. Open discussions of the project related activities and decisions continue on the TinkerPop dev mailing list. The committer count has been growing and various other TinkerPop community members have shown an increased amount of participation on tickets and in discussion. Excellent work is being performed by a variety of folks on a consistent basis that has resulted in two releases since the last report.

Over the past few months, work was performed to stand up a new DataStax hosted website, Planet TinkerPop. This raised concerns in the Apache community and it was determined that the best course of action was to request that Planet TinkerPop be taken down, and some of the vendor-neutral TinkerPop content be donated to TinkerPop for usage on the TinkerPop incubator site. Efforts were made to avoid running afoul of Apache guidelines prior to the Planet TinkerPop launch, but the PPMC did not have a complete understanding of Apache brand management expectations. Over the past weeks, the transition of DataStax donated content to the TinkerPop site has been completed. In the process, the PPMC has received guidance and clarification from mentors and the wider Apache community. We feel that the Apache TinkerPop site is stronger now than ever and will continue to serve as a great vendor neutral resource going forward. The lessons learned throughout this process will serve the TinkerPop PPMC well going forward.

Since the last status report, 166 new JIRA issues have been created and 127 have been resolved. Open discussions of project related activities and decisions continue on the TinkerPop dev mailing list.

Date of last release(s):

- * 2016-04-21 - Apache TinkerPop 3.1.2
- * 2016-04-21 - Apache TinkerPop 3.2.0

When were the last committers or PMC members elected?

- * Ketrina Yim - committer (2016-03-14)
- * Pieter Martin - committer (2016-02-21)
- * Jean-Baptiste Musso - committer (2016-03-28)
- * Michael Pollmeier - committer (2016-03-26)

Please see our Project Status page as we have been diligent to update it accordingly.

<http://incubator.apache.org/projects/tinkerpop.html>

Signed-off-by:

```
[ ](tinkerpop) Rich Bowen
[X](tinkerpop) Daniel Gruno
[ ](tinkerpop) Hadrian Zbarcea
[ ](tinkerpop) Matt Franklin
[ ](tinkerpop) David Nalley
```

Shepherd/Mentor notes:

Daniel Gruno (humbadooh):

I am satisfied that the issues raised last month have been amicably resolved.

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: nearing completion
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

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. Substantial effort in our build system to make Apache compliant

releases.

3. 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

Twill

Twill is an abstraction over Apache Hadoop® YARN to reduce the complexity of developing distributed applications.

Twill has been incubating since 2013-11-14.

Top three items to resolve before graduation:

- More engagement from the community.
- Improve documentations and examples.
- Increase adoption.

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

- None at this time.

How has the community developed since the last report?

- Fluo, being proposed for Apache incubation uses Twill
- Twill talk scheduled in Apache Big Data Conference 2016
- 7 new JIRA issues filed since last report
- 2 JIRA issues resolved since last report
- 1 new contributors submitted patches

How has the project developed since the last report?

- Version 0.8.0-incubating is being worked upon

Date of last release:

- 2016-01-26: 0.7.0-incubating

What are the plans for the next period?

- Encourage contributions from active users
- Identify potential committers
- Engage more on social channels (IRC and Twitter)

When were the last committers or PMC members elected?

- August 4, 2015 : Henry Saputra

Signed-off-by:

[](twill) Vinod K
[](twill) Arun C Murthy
[](twill) Tom White
[X](twill) Patrick Hunt

[](twill) Andrei Savu

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.

How has the project developed since the last report?

The first Apache Unomi 1.0.0-incubating has been released.

We are now preparing a serie of dependency updates, heading to the 1.1.0-incubating release.

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
[](unomi) Chris Mattmann