

November2016

/\ FINAL (!)

Incubator PMC report for November 2016

Timeline

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

Shepherd Assignments

Drew Farris	Predictio nIO		
Drew Farris	Toree		
John Ament	DataFu		
John Ament	Fluo		
Justin Mclean	Unomi		
P. Taylor Goetz	AriaTosc a		
P. Taylor Goetz	Beam		
P. Taylor Goetz	Fineract		
Timothy Chen	Eagle		
Timothy Chen	Slider		
Timothy Chen	iota		
Jia Zhai	Distribut edLog		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="82e09d25-7c17-431b-9ba3-11c69135b25b"><ac:plain-text-body><![CDATA[[none]	Annotator]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="86157d24-7faa-4413-b96c-b764e9e2274b"><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="8a1d5378-3af7-4f95-b3d0-74fb3a05fca6"><ac:plain-text-body><![CDATA[[none]	CarbonData]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a14cd28d-51fd-40ef-aad0-e647c49d6d27"><ac:plain-text-body><![CDATA[[none]	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="77ccf93f-7b9a-4677-a84e-5e6820eb87eb"><ac:plain-text-body><![CDATA[[none]	Edgent]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="acbd2474-cd73-4d3b-91a1-0ddf3f219a6f"><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="75159745-6a60-4909-9747-c376294ef84e"><ac:plain-text-body><![CDATA[[none]	Hivemall]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ab1d3a22-acc0-44fa-a72d-7eef3329730d"><ac:plain-text-body><![CDATA[[none]	Impala]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="60bf77c3-37b1-4caf-8c7e-062d65fd2425"><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="f85561c9-36b7-4f21-912a-3a8837f80bf1"><ac:plain-text-body><![CDATA[[none]	NetBeans]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7a07d5fc-04e3-4b92-aeca-112ea8cea522"><ac:plain-text-body><![CDATA[[none]	S2Graph]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="4243a08e-bc75-4cc7-b890-ee632c3623ea"><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="c50f05fa-636c-4b7c-b13d-9746afaad0b8"><ac:plain-text-body><![CDATA[[none]	Spot]]></ac:plain-text-body></ac:structured-macro>

<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a622c79d-c2e6-41e6-8b74-417b1be21fe5"><ac:plain-text-body><![CDATA[[none]	Streams]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="de87880f-9f72-481e-8e96-b15a82902db5"><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="e6188334-1b8d-4758-b0d4-7df1d204c58a"><ac:plain-text-body><![CDATA[[none]	Tamaya]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for November 2016

The Apache Incubator is the entry path into the ASF for projects and codebases wishing to become part of the Foundation's efforts.

There are currently 63 podlings incubating - no change from last month. The Geode podling has proposed itself for graduation. There were two new IPMC members added this past month. The IPMC approved 17 releases this month as well.

* Community

New IPMC members:

- Felix Meschberger
- Stephan Ewen

People who left the IPMC:

- None

* New Podlings

- None

* Graduations

The board has motions for the following:

- Geode

* Releases

The following releases entered distribution during the month of October:

- Apache Unomi 1.1.0-incubating 2016-10-03
- Apache Impala 2.7.0-incubating 2016-10-04
- Apache Fluo 1.0.0-incubating 2016-10-05
- Apache Htrace 4.2.0-incubating 2016-10-07
- Apache PredictionIO 0.10.0-incubating 2016-10-07
- Apache Pirk 0.2.0-incubating 2016-10-09
- Apache HAWQ 2.0.0.0-incubating 2016-10-10
- Apache CarbonData 0.1.1-incubating 2016-10-11
- Apache Metron 0.2.1BETA-incubating 2016-10-13
- Apache Geode 1.0.0-incubating 2016-10-15
- Apache Mnemonic 0.3.0-incubating 2016-10-21
- Apache Juneau 6.0.0-incubating 2016-10-24
- Apache Fineract 0.4.0-incubating 2016-10-25
- Apache Fluo Recipes 1.0.0-incubating 2016-10-27
- Apache Rya 3.2.10-incubating 2016-10-28
- Apache S2Graph 0.1.0-incubating 2016-10-30
- Apache Beam 0.3.0-incubating 2016-10-31

* IP Clearance

- None

* Legal / Trademarks

- Incoming podlings are pushing more for GitHub as master approaches.

- In addition, the usage of GitHub Issues continues to be asked.
- Better alignment between all impacted parties is needed for both of these.

* Infrastructure

- While the loss of a server is never convenient, the report manager gives huge kudos to the infra team for replacing the moin-moin wiki server quickly, avoiding much impact on this months report.

* Miscellaneous

- None

* Credits

- Report Manager: John D. Ament

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

* Still getting started at the Incubator

- AriaTosca
- Guacamole
- Hivemall
- iota
- NetBeans
- Spot
- Toree

* Not yet ready to graduate

Stagnant:

- Sirona

No release:

- Blur
- Climate Model Diagnostic Analyzer
- DistributedLog
- Edgent
- Joshua

Community growth:

- DataFu
- Fineract
- Fluo
- Impala
- PredictionIO
- S2Graph
- Streams
- Unomi

* Potentially Ready to Graduate

- BatchEE
- Beam
- CarbonData
- Eagle
- Slider
- SystemML
- Tamaya

* Did not report, expected next month

- Annotator

Table of Contents

AriaTosca
Beam
Blur
CarbonData
Climate Model Diagnostic Analyzer
DataFu
DistributedLog
Eagle
Edgent
Fineract
Fluo
Guacamole
Hivemall
Impala
iota
Joshua
NetBeans
PredictionIO
S2Graph
Sirona
Slider
Spot
Streams
SystemML
Tamaya
Toree
Unomi

AriaTosca

ARIA TOSCA project offers an easily consumable Software Development Kit(SDK) and a Command Line Interface(CLI) to implement TOSCA(Topology and Orchestration Specification of Cloud Applications) based solutions.

AriaTosca has been incubating since 2016-08-27.

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

1. Finalize migration of existing ARIA code, including Parser-NG with TOSCA and TOSCA for NFV profiles, and workflow engine to the ASF repo.
2. Create and publish release process for ARIA-TOSCA Project, release process will include Parser, TOSCA profiles, Workflow Engine, CLI.
3. Publish "How To Contribute" guide at ARIA-TOSCA Website.

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

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

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

How has the community developed since the last report?

Preparing TOSCA DSL training materials
Confluence space initialized

How has the project developed since the last report?

1. Created master branch
2. CI on travis enabled for the project
3. Initial workflow engine code migrated to ASF repo
4. Preparing migrating Parser code to ASF repo
5. Various task executors have been implemented

Date of last release:

None

When were the last committers or PMC members elected?

Project is being established in incubator with the proposed initial set of committers.

Signed-off-by:

[X](ariatosca) Suneel Marthi
[X](ariatosca) John D. Ament
[](ariatosca) 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). 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 Apex, Apache Spark, and Google Cloud Dataflow. Beam also brings SDKs in different languages, allowing users to easily implement their data integration processes.

Beam has been incubating since 2016-02-01.

The most important issue to address in the move towards graduation:

1. Make it easier for the Beam community to learn, use, and grow by expanding and improving the Beam documentation, code samples, and 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?

- * 441 closed/merged pull requests
- * High engagement on dev and user mailing lists (742 / 179 messages)
- * Several public talks, articles, and videos including:
 - @Scale San Jose ("No shard left behind: APIs for massive parallel efficiency in Apache Beam")
 - Strata + Hadoop World NYC ("Learn stream processing with Apache Beam")
 - Paris Spark Meetup ("Introduction to Apache Beam")
 - Hadoop Summit Melbourne ("Stream/Batch processing portable across on-prem (Spark, Flink) and Cloud with Apache Beam")
 - Hadoop User Group Taipei ("Stream Processing with Beam and Google Cloud Dataflow")
 - Data Science Lab London ("Apache Beam: Stream and Batch Processing; Unified and Portable!")

How has the project developed since the last report?

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

* Second and third incubating release (0.2.0 and 0.3.0) and a release guide [1]
* New DirectRunner support for testing streaming pipelines[2]
* Continued improvements to the Flink, Spark, and Dataflow runners
* Added support for new IO connectors, including MongoDB, Kinesis, and JDBC with Cassandra, MQTT support pending in pull requests
* Addition of the Apache Apex runner on a feature branch, and continued work on the Apache Gearpump runner and Python SDK feature branches. [3]
* Continued reorganization and refactoring of the project
* Continued improvements to documentation and testing

[1]: <http://beam.incubator.apache.org/contribute/release-guide/>

[2]: <http://beam.incubator.apache.org/blog/2016/10/20/test-stream.html>

[3]: <http://beam.incubator.apache.org/contribute/work-in-progress/#feature-branches>

Dates of last releases:

- * 2016/08/07 - 0.2.0-incubating
- * 2016/10/31 - 0.3.0-incubating

When were the last committers or PMC members elected?

The following committers were elected on 2016/10/20:

- * Thomas Weise
- * Jesse Anderson
- * Thomas Groh

Signed-off-by:

[X](beam) Jean-Baptiste Onofré
[](beam) Venkatesh Seetharam
[](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@ - 62[+1]; dev@ - 78[-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
[](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 a couple of new releases
2. Increase the communities
3. Prepare 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?

The community activity increased: many new users started to use and test CarbonData, we had more than 300 issues created till Nov; Two new finance enterprises have formally deployed CarbonData to their business system, and the query performance speeded up 10-70 times in comparison to old system (both are bank enterprise in China).

We finished 2nd Meetup in Beijing on 29th Oct, and CarbonData has increased 10+ contributors in last month.

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, and preparing org website.

We did the 2nd release (0.1.1-incubating) in Oct and we are preparing a new one(0.2.0) in Nov.

We have finished 2 technical talks in Bay area with Databricks, Alluxio in last month for discussing ecosystem integration with Spark and Alluxio.

Date of last release:

2016-10-10

When were the last committers or PMC members elected?

We elected a new committer Kumar Vishal on 2016-10-15.

Signed-off-by:

[X](carbondata) Henry Saputra
[X](carbondata) Jean-Baptiste Onofré
[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.

The achievements in October of 2016 are:

1. Drafted the Apache CMDA project web page
2. Developed a mechanism to keep track of user id in each service call
3. Designed a mechanism to define a workflow, to execute a workflow, to collect the result and provenance from the workflow execution.

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

1. To facilitate workflow scenarios in which a scientist can chain two or more services;
2. To develop a suite of science use cases;

3. To demonstrate the use of provenance in recommendation of data, users, and services.

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 climate scientist Frank Lee has joined the project to develop science use cases. The priority has been shifted to helping him in adding new source data files and modifying the software to accommodate these new data files.

How has the project developed since the last report?

The issues that we found from our 2016 Caltech/JPL climate summer school have been addressed. Examples are: a mechanism to capture and store user id for provenance data collection, and removal of some functionalities that are not meant for users that do not have login accounts on our website.

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
```

Shepherd/Mentor notes:

Chris Mattmann:

I'm worried still about this podling. Discussion seemed to still happen offlist and to be summarized to the list, instead of actively involving others in the community. Also, it says the website has been drafted - where is it? Where are the corresponding JIRA issues? Where are the public discussions about the site. The podling is/was on a short leash since the [DISCUSS] Retirement vote. I'd say give it one more month, and would appreciate others chiming in here. The other thing I'd say is that there is definitely mentor attrition here. I don't think we have enough active mentors reinforcing and sending the message (myself included) day to day. It's definitely more reactive now than proactive. The podling could use some help in that area.

DataFu

DataFu provides a collection of Hadoop MapReduce jobs and functions in higher level languages based on it to perform data analysis. It provides functions for common statistics tasks (e.g. quantiles, sampling), PageRank, stream sessionization, and set and bag operations. DataFu also provides Hadoop jobs for incremental data processing in MapReduce.

DataFu has been incubating since 2014-01-05.

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

1. Resolve NOTICE and LICENSE issues for binary distributions
2. Continued releases
3. Increased committer 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?

- * No updates

How has the project developed since the last report?

- * Released 1.3.1. Now using ASF-associated signing key. Feedback from previous release addressed.
- * Website updated alongside 1.3.1 release.
- * Cleaned up release instructions.

Date of last release:

2016-08-10

When were the last committers or PMC members elected?

July 2016 (Eyal Allweil)

Signed-off-by:

[](datafu) Ashutosh Chauhan
[X](datafu) Roman Shaposhnik
[](datafu) Ted Dunning

Shepherd/Mentor notes:

(rvs) pushing this community towards graduation is pretty high on my TODO list. I think they are as ready as they are ever going to be.

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:

- 1.Continue to grow the community, and increase diversity of community.
- 2.Improve documentation, including documentation of project and processes.
- 3.Successful 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?

1. Increase in contributions from community.
 - 8 created and 4 resolved issues in community JIRA in October.
2. Lots of engagement on documentation.
 - Enhance existing documents,
 - Setup guides for developers and committers.
3. Increased traffic on the mailing list, in particular, due to committers engaging more actively with contributors.
 - we have 35 people subscribed mail list.
 - 125 messages to distributedlog mail list in October.

How has the project developed since the last report?

1. Documentation has improved for project build and project deployment. Added more information on community

page.

New added pages:

<https://cwiki.apache.org/confluence/display/DL/Developer+Guide>

<https://cwiki.apache.org/confluence/display/DL/Committer+Guide>

2. Involved more discussion of new ideas and bring in new features. Include major discussions like 'transaction support', 'batch commit', "EventStore" . etc.

3. First release expected on November, and repackaging of the project under apache namespace are being discussed.

- Pull requests for repackaging.

- The major blocker is <https://issues.apache.org/jira/browse/DL-2>. We were expecting to let DL depends on an official bk version.

Date of last release:

NA

When were the last committers or PMC members elected?

NA

Signed-off-by:

[](distributedlog) Flavio Junqueira

[](distributedlog) Chris Nauroth

[X](distributedlog) Henry Saputra

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. Jira, PR are all populated with good description so that people can reference in the future.
2. Massive improvement on unit test, now it becomes stable and build status is closed monitored in README
3. Community has one more discussion about graduation and all are with positive feedback, and we are going through the graduation steps.

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 conference: QCon Shanghai
- More active contributors from YHD.com participate in large feature development

How has the project developed since the last report?

- 0.5 version is being in active development. This version will include big improvement in application management and alert engine.
- Job monitoring feature is ready including map/reduce and spark job
- Standalone alert engine which is highly scalable and user friendly in that user does not need code to configure alerting rules
- UI is completely rewritten to expose metadata to end user and user can define data source, stream and policy etc.
- Cluster health monitoring is ready, including hdfs health indicator, hbase health indicator etc.

Date of last release:

2016-07-19, and 0.5 release is being prepared

When were the last committers or PMC members elected?

- Jinhu Wu 2016-09-10

Signed-off-by:

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

Edgent

Edgent is a stream processing programming model and lightweight runtime to execute analytics at devices on the edge or at the gateway.

Edgent 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 the Edgent project
2. Create the first Apache release of Edgent
3. Document a repeatable release 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?

* A well attended meetup in Plano, TX about Apache Edgent was held on September 9th
"Sensor Data Analytics Acceleration with Apache Edgent"
<http://www.meetup.com/DFW-Sensor-Technology/events/233150402/>

* A presentation about Apache Edgent was presented on October 21st @ Open Source India 2016
Approximately 400 attendees were present.

* A presentation about Apache Edgent, Streaming Analytics and Weather Company Data was presented at World of Watson in Las Vegas on October 26th.

* A demo about Apache Edgent was presented at World of Watson in Las Vegas on October 24th.
"IoT Device Events to Streaming Analytics in 15 Minutes with Bluemix"

* For the months of September and October four new community members either subscribed to the edgent-dev mailing list, or opened JIRAs.

How has the project developed since the last report?

* A new build system was created using Gradle. The initial contribution for this work came from two community contributors. The follow up work was done by committers.

According to JIRA, the project has added the following:

* September: 11 new issues; 15 issues resolved.

* October: 46 new issues; 23 issues resolved.

* October activity was almost exclusively due to readying for the first Edgent release.

Date of last release:

* We have not created an Apache release, but are working on one to be ready soon. During the September and October time period we made build changes and handled licensing work and are poised to create our first release.

When were the last committers or PMC members elected?

* In May, we added two new committers and PPMC members, Kathy Saunders and Queenie Ma.

Signed-off-by:

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

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. With the first official incubation release out, we want to work on getting frequent successful incubation releases shipped to the community on a monthly basis.
2. Adding new committers and contributors to the project along with empowering volunteers and support partner organizations in the community to contribute to the codebase while using the project infrastructure.
3. Resolving the questions around introduction of a rework of the original code - in particular repository /infrastructure/release questions which result from the microservice architecture

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?

- * With the first official incubation release out, we look add several new committers who have worked on modules and features that will soon be merged into the release.
- * The community is more involved and engaged in asking questions on the developer and user mailing lists. We are working to ensure that all inquiries are actively addressed and more discussions regarding functionality and design occur on the mailing lists.

How has the project developed since the last report?

- * Our first release was made on October 25, 2016. We received 4 binding votes from the community.
- * Significant work was led by the Mifos core development team to address the licensing issues with a dependency on Hibernate and replacing it with OpenJPA - details on the code changes made and the challenges overcome can be found at <https://youtu.be/C4lXtXY-MY4>

Date of last release:

2016-10-25

When were the last committers or PMC members elected?

2016-05-05

Signed-off-by:

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

Shepherd/Mentor notes:

(rvs) next item on the immediate TODO list for the community is to master releases of binary convenience artifacts

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 and users

2. Additional releases
- 3.

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?

- * Two talks at the Accumulo Summit mentioned Fluo:
 - * <http://accumulo-summit.com/program/talks/tips-for-writing-fluo-applications/>
 - * <http://accumulo-summit.com/program/talks/indexing-strategies-for-searching-semantic-networks/>
- * Hadoop Weekly Issue #191 mentioned Fluo's release
- * Website traffic is up after the first release according to Google Analytics
- * Created Twitter Account : <http://twitter.com/apache-fluo>
- * Recent release of Rya which uses Fluo

How has the project developed since the last report?

- * Successfully made first release and two more releases :
 - * <https://fluo.apache.org/release/fluo-recipes-1.0.0-incubating/>
 - * <https://fluo.apache.org/release/fluo-1.0.0-incubating/>
 - * <https://lists.apache.org/thread.html/0d02022b544fef972dabcfbbf59381fc811fe3d0843f7052568cd56e@%3Cgeneral.incubator.apache.org%3E>
- * Created Fluo Tour (easy, hands on introduction) : <https://fluo.apache.org/tour/>
- * Completed Podling NameSearch : <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-109>
- * fluo.io now redirects to fluo.apache.org
- * The fluo-io GitHub org was renamed to astralway

Date of last release:

2016-10-28

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: The podling is definitely finding their legs in creating and releasing software, but very little progress has been made on growing the community (no prospective members).
I am also happy with the state of trademarks for the podling.

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?

It is clear that the continuing lack of a release is a major obstacle to community development, especially given outstanding pull requests which

cannot be merged due to pre-release code freeze.

We believe that the source and documentation are finally up-to-date with respect to Apache and the Incubator, and that we are ready to move forward with the procedures surrounding the release. Any assistance in navigating our first release would be greatly appreciated.

How has the community developed since the last report?

Community participation has remained roughly the same as last report. Mailing list participation is active but unchanged. Since last report, we have received an additional 3 pull requests, and have engaged the contributors for code review. Code looks good, and response to feedback has been professional, but merge is blocked until we can get the first release out of the way.

How has the project developed since the last report?

All outstanding issues which were aimed at the project's first release (0.9.10-incubating) have been completed, as have all bugs discovered during testing. We have made preparatory changes to the project website with the release in mind.

The project's old SourceForge forums have finally been closed, replaced by the mailing lists. The forums have been left in read-only mode for the sake of reference, with a stickied announcement notifying users of the move.

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 Onofré
[](guacamole) Daniel Gruno
[](guacamole) Olivier Lamy
[X](guacamole) Jim Jagielski
[X](guacamole) Greg Trasuk

Hivemall

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

Hivemall has been incubating since 2016-09-13.

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

1. Create the first Apache release
2. Community growth
3. IP clearance

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?

- * Presented a talk in Hadoop Summit Tokyo on Oct 26.
- * Still in progress at migrating the repository/community to ASF infra on <https://issues.apache.org/jira/browse/HIVEMALL-8>

How has the project developed since the last report?

- * We setup JIRA and managed issues on <https://issues.apache.org/jira/browse/HIVEMALL/>
- * We build the project site on <http://hivemall.incubator.apache.org/>
- * Documentation updates has been made <http://hivemall.incubator.apache.org/userguide/>

Date of last release:

No release yet

When were the last committers or PMC members elected?

None

Signed-off-by:

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

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. Community growth
2. Transition of user documentation to Apache hosting
3. Migration of pre-commit continuous integration testing to publicly-available infrastructure

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

No

How has the community developed since the last report?

Our last report was in August. Since then, we have five new contributors who have authored patches, while two relatively recent contributors who were active before August have continued their involvement by authoring new patches. Traffic to our developer mailing list has grown by about 60%.

How has the project developed since the last report?

There have been 241 commits since the last report.

Our status website now has 16 of the 17 listed work items complete. We had our first Apache release and have a wiki page describing how to perform the release in detail. We scrubbed our code using the RAT tool for copyright notices not compliant with the ASF rules. We wrote guidelines for contributors on how to become a committer and added a new committer. All developer documentation has now moved to the Apache-hosted wiki.

Date of last release:

2016-10-05

When were the last committers or PMC members elected?

2016-08-18

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 the developer community
2. Outreach at events outside the Apache ecosystem to inform and invite participation in the project
3. Getting an alpha release out by the end of this year based on the current code base.

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

None

How has the community developed since the last report?

Increased user interest with 4 individuals that are contributing in spurts of activity. This needs to become consistent.

How has the project developed since the last report?

Some improvements in the core engine, performers and the build process

Date of last release:

XXXX-XX-XX

When were the last committers or PMC members elected?

None but we are working on a proposal to add a new committer (a contributor that has made significant contributions in the last 6 months)

Signed-off-by:

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

Shepherd/Mentor notes:

Justin Mclean:

I have some concerns re progress of this poddling:

- The on list PPMC activity is low and there may be discussions happening off list.
- There is little progress towards making an Apache release.
- I'm not sure there are 3 active PPMC members.

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. Creating our first release.
2. Continue to build the community
3. Identify specific users and use cases.

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

None.

How has the community developed since the last report?

We have added a few new members.

How has the project developed since the last report?

A new release is imminent; we just need to pull the trigger. We have put together over sixty "language packs" that will be released in a no-dependency version and a Docker container.

Date of last release:

Forthcoming.

When were the last committers or PMC members elected?

John Hewitt (August 13, 2016)
Max Thomas (pending)
Michael Hedderich (pending)

Signed-off-by:

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

NetBeans

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

NetBeans has been incubating since 2016-10-01.

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

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

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

None.

How has the community developed since the last report?

Mailing lists and Wiki set up.

Discussions about voting, the CCLA/SGA, are taking place in the private mailing list, all other discussions in dev mailing list.

We want as many discussions as possible in dev mailing list, i.e., as public as possible.

How has the project developed since the last report?

Apache NetBeans Proposal has been published and accepted into incubation.

Apache Transition plan, listing everything needing to be done, including proposed milestones, is being worked on: <https://cwiki.apache.org/confluence/display/NETBEANS/Apache+Transition>

CCLA and SGA have been provided in draft form by Oracle and have been approved by ASF. Currently they're in the process of being approved and signed by Oracle.

Once that's done, experiments with migrating hg.netbeans.org/releases to Apache Git can proceed.

Date of last release:

No releases yet.

When were the last committers or PMC members elected?

No one has been elected so far.

Signed-off-by:

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

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. The initial code for PredictionIO was granted on 2016-06-16. A second grant of PredictionIO templates and SDKs was granted on 2016-09-20.

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

1. Establish a formal release schedule and process, allowing for dependable release cycles in a manner consistent with the Apache way.
2. Grow the community to establish diversity.
3. Migrate the remainder of former 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?

1. Both user and development mailing list are seeing increased activity.
2. Users requesting features are coming forward with code contributions. There have been discussions regarding future roadmap and development on the mailing lists.
3. The ecosystem around engine templates is slowly gaining traction again on Apache infrastructure. There are discussions around engine templates.

How has the project developed since the last report?

1. The first Apache release has been released on 2016-10-17.
2. The second software grant has been issued and recorded by ASF. Seven templates and five SDKs are now transferred to Apache.

Date of last release:

Apache PredictionIO 0.10.0-incubating on 2016-10-17

When were the last committers or PMC members elected?

Paul Li was elected as committer and PMC member on Aug 30, 2016.

Signed-off-by:

[X](predictionio) Andrew Purtell
[](predictionio) James Taylor

[](predictionio) Lars Hofhansl
[](predictionio) Luciano Resende
[](predictionio) Xiangrui Meng
[X](predictionio) Suneel Marthi

Shepherd/Mentor notes:

Drew Farris (shepherd): Two Mentors active on the mailing lists. Healthy activity observed on lists.

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 two sessions on S2Graph: usecases
- * Devview: <https://devview.kr/2016/schedule#session/160>
- * TechPlanet: http://techplanet.skplanet.com/eng/speaker_track2.html#spk_trk2_5
- * Started to discuss overhauling the website layout and contents

How has the project developed since the last report?

- * The first Apache release has been released on 2016-11-01.
- * 30 issues are created, 32 issues are resolved.

Date of last release:

Apache S2Graph 0.1.0-incubating on 2016-11-01.

When were the last committers or PMC members elected?

No

Signed-off-by:

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

Sirona

Monitoring Solution.

Sirona has been incubating since 2013-10-15.

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

1. Increase community/visibility
2. Get a bit more dynamic
- 3.

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

No

How has the community developed since the last report?

Not much

How has the project developed since the last report?

Some enhancements around the javaagent for short time living JVM use cases.

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

Shepherd/Mentor notes:

struberg: Cool project, but sadly we have a very low activity

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 had 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, the Slider Core module (including Application Master, Client and Java unit tests) has been rolled into an Apache Hadoop YARN feature branch as a hadoop-yarn-slider-core module. This effort is being captured in: YARN-5079 : Native YARN framework layer for services. All the migrations are happening in a branch (namely yarn-native-services) created off Hadoop trunk. 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>

For the benefit of those who would like to use Slider Core from the Hadoop codebase, but continue to use classic Slider Agent and legacy app-packages, a new branch (yarn-native-services) has been created in classic Slider repo. This branch retains only the Agent (python) code and the App Packages. It adds Slider Core module as a Hadoop dependency. This allows users to consume the latest of Slider Core. Future classic Slider releases can be made purely off of this branch.

Additionally an effort is being made to create an agent-less provider for legacy app-packages, which could help users to completely migrate to Hadoop codebase for creating long-running services with Slider. Expectation from app-package owners would be to make minor modification of their packages. Primary changes would be to shed the dependency on python Agent code, effectively making packages much simpler than what they look like today in classic Slider.

How has the project developed since the last report?

There has been fewer bug fixes on the classic Slider side as most of the focus has been on developing features on native YARN services in the Hadoop codebase. Work continues on support for complex services (assemblies) and agent-less applications in Slider. The efforts on Slider Core will continue in the new services branch created in YARN. Few key features which were merged to classic Slider as well were - SLIDER-875 - Ability to create an Uber application package with capability to deploy and manage as a single business app, and SLIDER-1107 - Generate app configuration files in AM. Several issues identified by Coverity scans were resolved as well. Additionally few patches were contributed by the community fixing functional and performance issues. Slider community plans to ship a release in the next quarter with all these features and bug fixes.

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

Spot

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

Spot has been incubating since 2016-09-23.

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

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

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

We currently have a hard dependency via our LDA implementation that requires a GPL license. A rewrite is in progress to move the code from LDA-C to Spark LDA.

How has the community developed since the last report?

This is first report after accepting Apache Spot to incubator and we still have not transitioned everything to ASF. We do however see increased interest in the project, primarily on our Slack channel.

How has the project developed since the last report?

This is first report.

Date of last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

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

Streams

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

Streams has been incubating since 2012-11-20.

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

1. Community growth and PMC maturity.
2. Demonstrate a consistent release schedule.
3. Participation of project community within related standards-bodies and Apache projects.

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

The project is now on a monthly reporting schedule to monitor effective progress in growing the community and active participation.

How has the community developed since the last report?

dev@streams.incubator.apache.org
115 emails (up >100%) sent by 16 people (+1), divided into 36 topics (up 400%).

<http://streams.incubator.apache.org>
196 Sessions (up ~60%), 247 Users (up ~35%), 1110 Pageviews (up >100%).

How has the project developed since the last report?

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

Excluding merges, 4 authors (+3) have pushed 32 commits (up ~375%) to master.
On master, 22 files (up 250%) have changed.

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

Excluding merges, 7 authors (+6) have pushed 106 commits (up >1000%) to master.
On master, 868 files (up >1000%) have changed.

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

Excluding merges, 3 authors (+2) have pushed 7 commits (+1) to master.
On master, 24 files have changed.

52 Issues closed for the upcoming 0.4-incubating release
35 new Issues opened (+28)

Date of last release:

2016-10-03 : 0.3-incubating release

When were the last committers or PMC members elected?

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

Signed-off-by:

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

Shepherd/Mentor notes:

Suneel Marthi: Healthy community activity, following the 0.3 release on Oct 3, 0.4 release is presently in the works and work has been scoped out for the 0.5 release.

Ate Douma: the Streams podling is getting back on track, making good progress:

- * new community participants
- * good mailing list discussions
- * elected a new committer and pPMC member (Joey)
- * a new release candidate 0.4-incubating is in progress

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 August through October had 375 messages on a wide range of topics. We have gained 4 new contributors to the main project since August 1st. Our website has been redesigned with the help of several design engineers and we have commits from 3 new contributors to the website project. On GitHub,

the project has been starred 417 times and forked 156 times.

Niketan Pansare gave a talk with the title "Apache SystemML - Declarative Machine Learning at Scale" on October 7th in the CS graduate seminar at UC Merced. Matthias Boehm gave a talk on "Compressed Linear Algebra for Large-Scale Machine Learning" at TU Dresden on August 30th. We presented the papers "Compressed Linear Algebra for Large-Scale Machine Learning" (research paper + poster) and "SystemML: Declarative Machine Learning on Spark" (industry paper) at VLDB'16. The "Compressed Linear Algebra for Large-Scale Machine Learning" paper won the VLDB 2016 Best Paper Award. We gave two 90 minute tutorials at the BOSS'16 workshop, co-located with VLDB'16, and our paper "SPOOF: Sum-Product Optimization and Operator Fusion for Large-Scale Machine Learning" has been accepted at CIDR'17.

How has the project developed since the last report?

The main project has had 213 commits since August 1. The website project has had 51 commits since August 1. Since August 1, 241 issues have been reported on our JIRA site and 137 issues have been resolved or closed. 79 pull requests have been created since August 1, and 72 pull requests 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. Extend Community
2. Improve Documentation and release base components
3. Graduate als TLP in 2017 after next releases

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

After our next releases and improving our homepage we'd like to graduate as a TLP in 2017.

How has the community developed since the last report?

We have increasing feedback of use cases to be covered and also a few bug reports, so Tamaya is used by developers and companies. Additionally we are doing regular bi-weekly hangouts and have defined a well defined roadmap to release and restructure our project. Apart from that we gathered feedback from JavaOne 2016: official Configuration JSR proposal and Anatole had a talk there.

The conflicts that arose in the past seem to be solved as some of the mentors and early committers have left the project. Since we've established our hangouts we have the feeling of beeing more connect and

able to work towards a common goal.

How has the project developed since the last report?

Apart from the main project we established new repositories to hold extensions and sandbox modules. In order to get rid of the problematic webpage generation we decided to start from scratch.

In addition to making the API more smooth, we try to keep the base repository/project as lean as possible and prepare a next release. Due to various technical issues (infrastructure behaves differently from local builds and checkouts) we did not meet our deadline on getting a release out by the end of October.

Date of last release:

2016-04-06

When were the last committers or PMC members elected?

Phil Ottlinger at 2016-04-24.

Signed-off-by:

[X](tamaya) John D. Ament
[](tamaya) David Blevins

Shepherd/Mentor notes:

johndament:

While the podling has had issues in the past around community growth, I am seeing them operate more consistently, with clear open discussions. I believe they're operating well and should be ready to graduate soon, if not already.

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. Make a release: The community is working on RC from 0.1.x branch. Master has moved to start support for Spark 2.0. Currently working on RC3.
2. 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) In progress. Project elected new PPMC member.

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

NONE, previous issue with LGPL dependency has been RESOLVED as the JeroMQ dependency has now been released as MPL license.

How has the community developed since the last report?

1. Active communication in mailing list and gitter with early adopters
2. No worries about transitioning from old Spark Kernel code.
3. More external contributions being made. Mainly focused on master to stabilize Spark 2.0 support

How has the project developed since the last report?

1. Working on 1st release. Got an RC3 going on for a vote
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:

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

Shepherd/Mentor notes:

Drew Farris (shepherd): Two Mentors active on the mailing lists. Healthy activity and progress towards release observed on the mailing lists.

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?

We are still 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. Some small improvements have been performed in that way.

How has the project developed since the last report?

The second Apache Unomi 1.1.0-incubating has been released.

We updated key dependencies and provided new features in addition of the bug fixes. We also fixed the LGPL issue in the binary distribution. The Unomi rules and conditions engine has been improved as well.

Date of last release:

2016-10-03

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](unomi) Jean-Baptiste Onofré
[](unomi) Bertrand Delacretaz
[](unomi) Chris Mattmann

Shepherd/Mentor notes:

Chris Mattmann:

I am going to humbly request being removed as a mentor for this podling. I haven't

had a chance to send it to the list but I don't have the bandwidth to actively provide oversight here.