

# July2017

/\ FINAL (!)

## Incubator PMC report for July 2017

### Timeline

Wed July 05	Podling reports due by end of day
Sun July 09	Shepherd reviews due by end of day
Sun July 09	Summary due by end of day
Tue July 11	Mentor signoff due by end of day
Wed July 12	Report submitted to Board
Wed July 19	Board meeting

### Shepherd Assignments

Drew Farris	DataFu		
Drew Farris	HAWQ		
John Ament	BatchEE		
Justin Mclean	MRQL		
P. Taylor Goetz	Myriad		
P. Taylor Goetz	ODF Toolkit		
P. Taylor Goetz	Superset		
Timothy Chen	Goblin		
Timothy Chen	HORN		
Timothy Chen	Traffic Control		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="703fd414-06a3-4ea4-9128-cdb6533dce90"><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="df19c388-2021-41ce-abb2-1a250741aaa2"><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="040a412b-2d63-460b-baa9-bcea1a0c7e1f"><ac:plain-text-body><![CDATA[	[none]	FreeMarker	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="74344c3f-dc51-4a7d-b864-be4333be99e2"><ac:plain-text-body><![CDATA[	[none]	Gossip	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="96552005-2297-4af2-96ae-0202a871c77d"><ac:plain-text-body><![CDATA[	[none]	Juneau	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a8f865d4-9eec-42cd-bb36-3749203803ca"><ac:plain-text-body><![CDATA[	[none]	Livy	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ff56ae53-cc37-4a3f-ab8b-53e5bbe396a9"><ac:plain-text-body><![CDATA[	[none]	MADlib	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="00f7e31a-bf59-4aad-99d0-f566ca495a16"><ac:plain-text-body><![CDATA[	[none]	MXNet	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="0a3ca7aa-6593-4509-84fc-b60152cddd75"><ac:plain-text-body><![CDATA[	[none]	Milagro	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e2f6de0a-9c6f-45b7-8e5f-72847f03828c"><ac:plain-text-body><![CDATA[	[none]	Myriad	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ec2d4703-c0d2-419e-8937-5c6a8363400d"><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="68bac588-7acf-4c30-96ad-c511903f90fe"><ac:plain-text-body><![CDATA[	[none]	Pulsar	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="0e99d674-6af2-4008-a890-224a679a9e96"><ac:plain-text-body><![CDATA[	[none]	RocketMQ	]]></ac:plain-text-body></ac:structured-macro>
David Fisher	Rya		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="a799b2ff-31de-49c5-8993-2de62e085db3"><ac:plain-text-body><![CDATA[	[none]	Sensio	]]></ac:plain-text-body></ac:structured-macro>

## Report content

Incubator PMC report for July 2017

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 57 podlings incubating. Three podlings joined us this month, two podlings have retired, one has graduated to a sub-project and the board has resolutions for one (three?) podlings to become TLPs this month. Podlings executed eight releases this past month. We have added two new IPMC members, neither an existing foundation member.

### \* Community

New IPMC members:

- Steve Blackmon
- Benjamin Young

### \* New Podlings

- Heron
- Livy
- Pulsar

### \* Podlings that failed to report, expected next month

- BatchEE - Low on list activity
- DataFu - No on list activity, will follow up with a status request and potentially retirement
- Gobblin - No on list activity, will follow up with a status request
- HORN - Actively voting on retirement, no report expected
- MRQL - Discussing retirement, single developer mode
- Netbeans - Moderate on list activity, no one stepped up to do report

### \* Graduations

The board has motions for the following:

- Fluo
- MADlib
- Streams

The following podlings graduated into subprojects:

- DistributedLog (BookKeeper)

### \* Retirements

The following podlings retired this month:

- Blur
- Sirona

### \* Releases

The following releases entered distribution during the month of June:

- 2017-06-09 Apache Weex 0.12.0
- 2017-06-13 Apache TrafficControl 1.8.1
- 2017-06-15 Apache Impala 2.9.0
- 2017-06-17 Apache RocketMQ 4.1.0
- 2017-06-18 Apache Tamaya 0.3
- 2017-06-22 Apache Fluo Recipes 1.1.0
- 2017-06-23 Apache Mnemonic 0.8.0
- 2017-06-28 Apache Juneau 6.3.0

### \* Infrastructure

- We need to continue to ensure that podlings are bootstrapped in proper sequence. Recent changes to Whimsy have caused new steps to be added, but not communicated. We will need infra help to support some of these new steps.

- Thanks to the infra team, the Incubator website has been migrated to a new technology and has received a needed facelift.

#### \* Miscellaneous

- With the website migrated, work will begin to refresh the Incubator documentation.

#### \* Credits

- Special thanks to Dave Fisher for rejoining the shepherd community

---

### Table of Contents

Airflow  
Annotator  
FreeMarker  
Gossip  
HAWQ  
Juneau  
Livy  
MADlib  
Milagro  
MXNet  
Myriad  
ODF Toolkit  
Pulsar  
RocketMQ  
Rya  
SensSoft  
Superset  
Traffic Control  
Weex

---

#### 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. We are working on our third apache release 1.8.2 to get more experience with the process
- 2.
- 3.

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

None

How has the community developed since the last report?

1. We had our second official release. 1.8.1 on May 9th 2017.
2. Since our last podling report 3 months ago (i.e. between Apr 1 and July 5, inclusive), we grew our contributors from 256 to 280
4. Since our last podling report 3 months ago (i.e. between Apr 1 and July 5, inclusive), we resolved 151 pull requests (currently at 1630 closed PRs)
5. A new meet-up group formed in Tokyo, Japan - they held their first meet-up on May 11.
6. Since being accepted into the incubator, the number of companies officially using Apache Airflow has risen from 30 to 98, 15 new from last podling report.

How has the project developed since the last report?

See above : 151 PR resolved, 24 new contributors, & 15 new companies officially using it.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☐ Working towards first release
- ☐ Community building
- ☒ Nearing graduation
- ☐ Other:

Date of last release:

2017-05-09

When were the last committers or PPMC members elected?

As mentioned on  
<https://cwiki.apache.org/confluence/display/AIRFLOW/Announcements#Announcements-Mar14,2017>  
Alex Guziel joined the Apache Airflow PPMC/Committer group.

Signed-off-by:

- ☐ (airflow) Chris Nauroth  
Comments:
- ☒ (airflow) Hitesh Shah  
Comments:
- ☐ (airflow) Jakob Homan  
Comments:

-----  
Annotator

Annotator provides annotation enabling code for browsers, servers, and humans.

Annotator has been incubating since 2016-08-30.

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

1. grow the participation of folks with commit bit
2. unblock community "waiting" on founders
3. ship some code

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

Not at this time.

How has the community developed since the last report?

May was a busy "in person" engagement month. Members of our community participated at <http://iannotate.org/> and ApacheCon NA 2017 in Miami--where there was an Annotator specific presentation: <https://apachecon2017.sched.com/mobile/#session:f8889fb0de17e3b2fec61e068889e06b> There was great interest in the project and it's future.

Additionally, documents were signed by members of The Hypothes.is Project <http://hypothes.is/> There is interest from Hypothesis in depending on and eventually contributing to Annotator.

We also added tbdinesh <http://people.apache.org/phonebook.html?uid=tbdinesh> as a core committer--he was in the list of initial committers in our original proposal.

How has the project developed since the last report?

We have successfully setup Gitbox (thanks Infra!), and moved our development to GitHub <https://github.com/apache/incubator-annotator> We have processed a couple PRs (from core committers), and intend to work in a PR-first, discuss, merge pattern to encourage more involvement from the community in code review.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☒ Initial setup

[X] Working towards first release  
[X] Community building  
[ ] Nearing graduation  
[ ] Other:

Date of last release:

XXXX-XX-XX - still coding toward a first release

When were the last committers or PPMC members elected?

We are still working to add all initial committers to the project.

Signed-off-by:

[\*](annotator) Nick Kew  
Comments: Activity (excluding f2f events referenced above) seems low but healthy.  
[ ](annotator) Brian McCallister  
Comments:  
[ ](annotator) Daniel Gruno  
Comments:  
[\*](annotator) Jim Jagielski  
Comments:

-----

FreeMarker

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

FreeMarker has been incubating since 2015-07-01.

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

We have a large user base but a rather small group of committers. This was actually expected, given the maturity (and topic) of the project. While the FreeMarker 3 branch, which was started 5 month ago, will be much more appealing for contributors, development and growth due to that will certainly take a long time. In other respects the project is mature and ready for graduation.

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

We are eager to graduate but we're afraid that the number of active contributors will prevent that in the foreseeable future. But remaining in the incubator hurts the project as well. One problem it causes is that many projects are stuck with the last non-Apache release of the project, from two years ago, since they won't depend on an artifact whose version contains "-incubating".

How has the community developed since the last report?

There were no changes since the last report.

How has the project developed since the last report?

Most activity was in the FreeMarker 3 branch; most notably, FreeMarker was modularized and switched to Gradle. The FreeMarker 2 branch was less active, but had some fixes and new features. During the last report period we have received a code donation, the source code of an online template evaluator service, which was since then cleaned up, and is now running on Apache infrastructure (<http://try.freemarker.org/>). The build of FreeMarker and of the template evaluator service was added to Travis CI.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☐ Working towards first release
- ☒ Community building
- ☒ Nearing graduation
- ☐ Other:

Date of last release:

2017-03-25

When were the last committers or PPMC members elected?

2017-02-13 Woonsan Ko, committer (non-PMC)

Signed-off-by:

- ☒ (freemarker) Jacopo Cappellato  
Comments:
- ☒ (freemarker) Jean-Frederic Clere  
Comments:
- ☒ (freemarker) David E. Jones  
Comments:
- ☐ (freemarker) Ralph Goers  
Comments:
- ☐ (freemarker) Sergio Fernández  
Comments:

-----  
Gossip

Gossip is an implementation of the Gossip Protocol.

Gossip has been incubating since 2016-04-28.

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

1. Increase the number of active committers/reviewers
2. Gossip used in a downstream apache application
3. Keep making releases

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?

We have two students in the GSOC program and some new active bodies. A majority of admin bandwidth was spent organizing and helping with proposals.

How has the project developed since the last report?

We have added new Crdt types. We have completed event listeners and callbacks. We made it easier to run the examples.

How would you assess the podling's maturity?

Software is solid. Current GSOC students will hopefully remain active after the GSOC period.

- ☐ Initial setup
- ☐ Working towards first release
- ☒ Community building
- ☐ Nearing graduation
- ☐ Other:

Date of last release:

2017-04-02

When were the last committers or PPMC members elected?

No

Signed-off-by:

[x](gossip) P. Taylor Goetz

Comments: Podling is doing well. The community is growing somewhat slowly but surely.

[x](gossip) Josh Elser

Comments: I think things are going well. The GSOC students seem to be progressing well which is nice. Hopefully they will stay involved after the end of the program.

[ ](gossip) Drew Farris

Comments:

-----  
HAWQ

Apache HAWQ is a Hadoop native SQL query engine that combines the key technological advantages of MPP database with the scalability and convenience of Hadoop. HAWQ reads data from and writes data to HDFS natively. HAWQ delivers industry-leading performance and linear scalability. It provides users the tools to confidently and successfully interact with petabyte range data sets. HAWQ provides users with a complete, standards compliant SQL interface.

HAWQ has been incubating since 2015-09-04.

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

1. Continue to improve the project's release cadence. To this end we plan on expanding automation services to support increased developer participation.
2. Continue to expand the community, by adding new contributors and focusing on making sure that there's a much more robust level of conversations and discussions happening around roadmaps and feature development on the public dev mailing list.
3. Starting with Apache HAWQ 2.2.0.0 release, the project will expand release artifacts to include the delivery of convenience binary artifacts. We expect the project to significantly refine the binary release process in several key areas. The team is limiting scope to HAWQ Core and PXF components. The appropriate LICENSE, DISCLAIMER and NOTICE files are included in the binary release artifacts.

FYI: The team has decided to move Ranger (optional component) support into a subsequent binary release. It remains supported in the source release.

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

Nothing urgent at this time.

How has the community developed since the last report?

1. Conference Talks (7):

- \* Introduction to latest technology in HAWQ 2.X, the 8th Database Technology Conference China (Speaker: Lili Ma, Pivotal, May 13, 2017)
- \* "Podling Shark" lightning session, ApacheCon North America 2017, (Speaker: Aleksandr Diachenko and Alexander Denissov, Pivotal, May 18, 2017)
- \* The Big Data Engine in Cloud Era, CSDN Cloud Computing

Technology Conference (Speaker: Zhenglin Tao, Oushu Inc, May 19, 2017)

- \* HAWQ Introduction. China International Big Data Industry Expo 2017 (Speaker: Lan Zhou, Oushu Inc, May 26, 2017)
- \* Extending Apache Ranger Authorization Beyond Hadoop: Review of Apache Ranger Extensibility Framework & Case Study for Integration with Apache Hawq, DataWorks Summit, <https://www.youtube.com/watch?v=6SE32zrgIAU> (Speaker: Ramesh Mani, Hortonworks & Alexander Denissov, Pivotal, June 13, 2017)
- \* Hawq Meets Hive - Querying Unmanaged Data, DataWorks Summit, <https://www.youtube.com/watch?v=slZJvHx1hM> (Speaker: Shivram Mani & Alex (Oleksandr) Diachenko, Pivotal, June 14, 2017)
- \* Apache HAWQ: Open Source MPP++ Database, The 12th China Open Source World Summit, (Speaker: Lei Chang, Oushu Inc, June 21, 2017)

2. Active contributions from approximately 20 different community contributors since the last report (April 2017).
3. The Open Source Greenplum Database community (<http://greenplum.org>) is reviewing opportunities to leverage Apache HAWQ components.

Here is the initial GitHub Pull Request:

<https://github.com/greenplum-db/gpdb/pull/2634>

How has the project developed since the last report?

1. Apache HAWQ 2.2.0.0 release has passed PPMC vote. This contains the initial convenience binary release. The release has been sent to IPMC for voting.

FYI: The project's Apache HAWQ 2.2.0.0 release manager (Ruilong Huo) represents Oushu Inc. (<http://www.oushu.io> - HAWQ++ is the world's first MPP SQL engine that can run native to the container cloud platform).

Release information:

- 1) Release page:  
<https://cwiki.apache.org/confluence/display/HAWQ/Apache+HAWQ+2.2.0.0-incubating+Release>

- 2) Issues/tasks fixed (80):  
<https://issues.apache.org/jira/issues/?filter=12339844>

2. Per project mentor guidance, the Apache Incubator logo has been applied to the project's website (<http://hawq.incubator.apache.org/>) and wiki (<https://cwiki.apache.org/confluence/display/HAWQ/Apache+HAWQ+Home>).
3. Per project mentor guidance, an incubator project logo was created and voted on. It has been applied to the projet's wiki: <https://cwiki.apache.org/confluence/display/HAWQ/Apache+HAWQ+Home>
4. To encourage more doc contributions, a discussions took place on dev/user mail lists to recommend moving the HAWQ doc repository (<https://github.com/apache/incubator-hawq-docs> master branch) into the HAWQ source repository (<https://github.com/apache/incubator-hawq>).

5. Project mail list activity:

Between Sat Apr 01 2017 and Fri Jun 30 2017:

dev@hawq.apache.org  
337 emails sent  
147 topics started  
87 participants



user@hawq.apache.org  
78 emails sent  
24 topics started  
35 participants

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☐ Working towards first release
- ☒ Community building
- ☒ Nearing graduation
- ☐ Other:

Date of last release:

2017-02-28, Apache HAWQ 2.1.0.0

When were the last committers or PPMC members elected?

Podling committers (1) added:  
Xiang Sheng (<https://github.com/stanlyxiang>), May 16, 2017

Signed-off-by:

- ☒ (hawq) Alan Gates  
Comments:
- ☐ (hawq) Konstantin Boudnik  
Comments:
- ☐ (hawq) Justin Erenkrantz  
Comments:
- ☐ (hawq) Thejas Nair  
Comments:
- ☐ (hawq) Roman Shaposhnik  
Comments:

-----  
Juneau

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

Juneau has been incubating since 2016-06-24.

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

1. Build up community with non-IBM and non-Salesforce contributors.
2. Solicit user feedback/usage
3. Grow awareness of the project

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 Juneau community continues collaborating with the Apache Streams team and have come up with significant new enhancements from the effort.

How has the project developed since the last report?

6.2.0 major release delivered on Apr 28, 2017.  
6.3.0 major release delivered on Jun 28, 2017.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☐ Working towards first release
- ☒ Community building
- ☐ Nearing graduation
- ☐ Other:

Date of last release:

2017-06-28

When were the last committers or PPMC members elected?

- Initial committers only.

Signed-off-by:

☒ (juneau) Craig Russell

Comments: I'm ambivalent about the use of Slack as a communications channel. While it's great for just-in-time discussion, it's not visible to the dev list (the digests are of no value). Juneau is mostly a one-man show and I'm hopeful that others will join the project.

☐ (juneau) Jochen Wiedmann

Comments:

☒ (juneau) John D. Ament

Comments: The podling is struggling a bit with community growth and public communication. While we've opened up Slack as a communication channel and it has led to some external contribution, we need to figure out how to get more communication happening on list and how to get more outside contributors.

-----  
Livy

Livy is web service that exposes a REST interface for managing long running Apache Spark contexts in your cluster. With Livy, new applications can be built on top of Apache Spark that require fine grained interaction with many Spark contexts.

Livy has been incubating since 2017-06-05.

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

N/A

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

N/A

How has the community developed since the last report?

The community is working on moving the project to Apache.

How has the project developed since the last report?

N/A

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☒ Initial setup
- ☐ Working towards first release
- ☐ Community building
- ☐ Nearing graduation
- ☐ Other:

Date of last release:

N/A

When were the last committers or PPMC members elected?

N/A

Signed-off-by:

[ ](lily) Bikas Saha  
Comments:  
[ ](lily) Brock Noland  
Comments:  
[ ](lily) Luciano Resende  
Comments:  
[X](lily) Jean-Baptiste Onofre  
Comments:

-----  
MADlib

Big Data Machine Learning in SQL for Data Scientists.

MADlib has been incubating since 2015-09-15.

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

1. Finalize trademark transfer from Pivotal to ASF.
2. Continue to produce regular Apache (incubating) releases.
3. Continue to execute and manage the project according to governance model of the "Apache Way".

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

1. The Apache MADlib Project is ready for graduation out of the incubator.  
Discussion by Project:  
<https://lists.apache.org/thread.html/070c6764fcd0448b2db8975936b52f7a28bd0e231c0e690288a6968e@%3Cdev.madlib.apache.org%3E>  
Vote by IPMC and community:  
<https://lists.apache.org/thread.html/733920464e8f8170d9cc831b701f275d757ee9448a7bfd05albf8dfd@%3Cgeneral.incubator.apache.org%3E>  
Trademark transfer from Pivotal to ASF is being tracked in:  
<https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-125>  
2. The resolution for graduation was tabled by the board last month due to trademark issue, and is now being re-submitted.

How has the community developed since the last report?

1. Some related events in Q2 2017:
  - \* May 25, 2017 - MADlib community call. Topic: New Features in Apache MADlib 1.11 (Frank McQuillan)
  - \* Jun 21, 2017 - Greenplum meetup in San Francisco. Topic: Apache Solr & MADlib (incubating): Enabling Massive Text Analytics In-Database (Bharath Sitaraman)
  - \* Jul 5-7, 2017 - PG Day Russia. Topic: Various on "Greenplum Day" Jul 5 including in-database analytics (Roman Shaposhnik and others)
  - \* Jul 25, 2017 (upcoming) - SF Bay ACM Chapter meetup. Topic: Advanced Analytics for Security: Lateral Movement Detection (Anirudh Kondaveti)
2. See material technical conversations on user/dev mailing lists and in the appropriate JIRAs and pull requests.

How has the project developed since the last report?

1. TLP readiness - maturity evaluation matrix <https://cwiki.apache.org/confluence/display/MADLIB/ASF+Maturity+Evaluation>
2. TLP readiness - graduation resolution <https://cwiki.apache.org/confluence/display/MADLIB/Graduation+Resolution>
3. TLP readiness - documented release process <https://cwiki.apache.org/confluence/display/MADLIB/Release+Process>
4. Active work in progress for 6th ASF release MADlib v1.12 scheduled for Jul/Aug 2017. Features include: more graph analytics (weakly connected components, breadth first search, all pairs shortest path, multiple graph measures), neural nets, stratified sampling, train-test split, improvements to decision tree & random forest, improvements to summary function
5. Mailing list activity in Q2: 295 postings to dev, 77 postings to user.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

[ ] Initial setup  
[ ] Working towards first release

- ☐ Community building
- ☒ Nearing graduation
- ☐ Other:

Date of last release:

MADlib v1.11 on 5/16/17.

When were the last committers or PPMC members elected:

Orhan Kislal on 9/7/16 and Nandish Jayaram on 9/7/16.

Signed-off-by:

- ☐ (madlib) Konstantin Boudnik  
Comments:
- ☒ (madlib) Ted Dunning  
Comments:
- ☐ (madlib) Roman Shaposhnik  
Comments:

-----  
Milagro

Distributed Cryptography; M-Pin protocol for Identity and Trust

Milagro has been incubating since 2015-12-21.

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

1. Improve project communication
2. Define clear commit guidelines
3. Building the Milagro community - engaging developers and cryptographers, raising awareness and helping to secure future of internet.

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

How has the community developed since the last report?

The community hasn't developed since the last report, due to lack of activity. Communication has picked up again so it's expected that the project will gather more interest soon.

How has the project developed since the last report?

In past few months, the Milagro community was focused on developing Milagro Crypto C Library.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☒ Working towards first release
- ☐ Community building
- ☐ Nearing graduation
- ☐ Other:

Date of last release:

n/a

When were the last committers or PPMC members elected?

n/a

Signed-off-by:

- ☐ (milagro) Sterling Hughes

Comments:

[ ](milagro) Jan Willem Janssen

Comments:

[\*](milagro) Nick Kew

Comments: Activity at Apache has not been progressing. A discussion has been opened, and team members appear keen to improve communication and deal with a backlog of activity.

IPMC/Shepherd notes:

jda: a discussion is active on the future of the project. will continue to monitor.

-----

MXNet

MXNet is an open-source deep learning framework that allows you to define, train, and deploy deep neural networks on a wide array of devices, from cloud infrastructure to mobile devices. It is highly scalable, allowing for fast model training, and supports a flexible programming model and multiple languages. MXNet allows you to mix symbolic and imperative programming flavors to maximize both efficiency and productivity. MXNet is built on a dynamic dependency scheduler that automatically parallelizes both symbolic and imperative operations on the fly. A graph optimization layer on top of that makes symbolic execution fast and memory efficient. The MXNet library is portable and lightweight, and it scales to multiple GPUs and multiple machines.

MXNet has been incubating since 2017-01-23.

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

1. Migrate code (GitHub) and website to Apache Infra.
2. Grow the community
3. Establish a reliable Release process consistent with Apache Way.

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

None

How has the community developed since the last report?

- a) On 5/27 MXNet published a comprehensive edit and makeover of the documentation including tutorials, how-to's, APIs and architecture guides. This was a broad effort that involved over 40 contributors.
- b) The PMC voted in a new committer who has been helping with the code migration and setup of the test infrastructure. We are making slow but steady progress towards getting the GitHub code migrated. The target date for migration is 7/17. Website migration will happen after that.
- c) Slack and dev@ are being used more actively.
- d) Two presentations/workshops on Apache MXNet at the O'Reilly AI Conf on 6/27 and 6/28
- e) A new blog post published on 6/23 showing users how to Build a Real-time Object Classification System with ApacheMXNet on Raspberry Pi.  
<https://aws.amazon.com/blogs/ai/build-a-real-time-object-classification-system-with-apache-mxnet-on-raspberry-pi/>

How has the project developed since the last report?

- a) Since the last report 42 authors have pushed 326 commits to master.
- b) Documentation- Architecture guides, How To's, Tutorials, and APIs have been improved.
- c) More features (e.g. operators) requested by the user community have been added.
- d) A new Perl language binding for MXNet was added.

How would you assess the podling's maturity?

Podling's still getting established in Apache - so maturity == Low.

Please feel free to add your own commentary.

☒ Initial setup  
☐ Working towards first release  
☐ Community building  
☐ Nearing graduation  
☐ Other:

Date of last release:

No Release yet, project is still getting established in Apache.

When were the last committers or PPMC members elected?

Ly Nguyen added as a committer and PPMC member in June 2017.

Signed-off-by:

☒ (mxnet) Sebastian Schelter  
Comments:  
☒ (mxnet) Suneel Marthi  
Comments:  
☒ (mxnet) Markus Weimer  
Comments:  
☒ (mxnet) Henri Yandell  
Comments:

-----  
Myriad

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

Myriad has been incubating since 2015-03-01.

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

1. Many existing committers finding it difficult to give dedicated time to the project as they are busy with other projects
2. Talk of retirement. No consensus
3. Need for a road map for 0.3 release

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

Discussions on whether the project should be retired. No consensus.  
Interest from committers to vet releases (multiple), contribute features (at least couple of committers), return to active contributions (at least couple of committers)

How has the community developed since the last report?

\* dev@ mailing list has seen renewed activity due to retirement discussions.  
No consensus.

How has the project developed since the last report?

\* Little to no activity outside the mailing list

How would you assess the podling's maturity?

Please feel free to add your own commentary.

☐ Initial setup  
☐ Working towards first release  
☒ Community building

[ ] Nearing graduation  
[ ] Other:

Date of last release:

2016-06-20 myriad-0.2.0-incubating

When were the last committers or PPMC members elected?

2016-06-16 Yuliya Feldman

Signed-off-by:

[ ](myriad) Benjamin Hindman  
Comments:  
[ ](myriad) Danese Cooper  
Comments:  
[X](myriad) Ted Dunning  
Comments:  
[ ](myriad) Luciano Resende  
Comments:

IPMC/Shepherd notes:

jda: manual reminder sent. active discussion around retirement.

-----  
ODF Toolkit

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

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

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

1. Attract more developers
2. Have frequent releases
3. Fulfil the 1.0 version

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?

Patches had been provided by new developers.

How has the project developed since the last report?

After years, we finally made a new release our last supporting JDK6, continuing on JDK 8 including new RDF  
feature.

We already received patches from developers we haven't known before, but not yet sufficient patches to elect  
new committers.

We discussed ways to clean the code base for a 1.0 version of this project.

Finally, Svante Schubert is continuing to work on his branch for a major collaboration feature, certainly to  
attract more developers.

Date of last release:

2017-04-10

When were the last committers or PPMC members elected?

2012-10-29

Signed-off-by:

[ ](odftoolkit) Nick Burch  
Comments:

[ ](odftoolkit) Yegor Kozlov  
Comments:  
[x](odftoolkit) Tom Barber  
Comments: Slow but steady progress as usual. That said I know Svante has some stuff up his sleeve for later this year to drive development further which would hopefully attract new committers.

-----

Pulsar

Pulsar is a highly scalable, low latency messaging platform running on commodity hardware. It provides simple pub-sub semantics over topics, guaranteed at-least-once delivery of messages, automatic cursor management for subscribers, and cross-datacenter replication.

Pulsar has been incubating since 2017-06-01.

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

- 1.Finish onboarding the PPMC. We have 11 of 16 fully in place. With two more close to setup.
- 2.Make an Apache release
- 3.Grow the community with new Committers/PPMC members.

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. All but one of the 16 initial Committers have participated in one way or another. The existing Pulsar Github was moved to an Apache Github/GitBox setup. Development discussions are happening on the dev list.

How has the project developed since the last report?

The project cut a v1.18 release with the old organization just before conversion. The SGA has been received, the code has switched to org.apache.pulsar and the license/copyright and headers modified to Apache.

10 authors have pushed 82 commits to the project in the last 30 days.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- [x] Initial setup
- [x] Working towards first release
- [ ] Community building
- [ ] Nearing graduation
- [ ] Other:

Good progress for the first month.

Date of last release:

2017-06-17, v1.18

When were the last committers or PPMC members elected?

No one has been added in the last 30 days.

Signed-off-by:

- [X](pulsar) Dave Fisher  
Comments: The project is getting started. Very good so far.
- [x](pulsar) Jim Jagielski  
Comments:
- [x](pulsar) P. Taylor Goetz  
Comments: The project is off to a good start.
- [x](pulsar) Francis Liu  
Comments:



-----

RocketMQ

RocketMQ is a fast, low latency, reliable, scalable, distributed, easy to use message-oriented middleware, especially for processing large amounts of streaming data.

RocketMQ has been incubating since 2016-11-21.

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

1. Grow the community
2. Vote for new PMC member
3. Near graduation

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 mailing list activity: solutions and milestone plan are discussed and decided in dev mailing list. While more and more questions and asking for support for RocketMQ are occurred in user mailing list.
- \* Organized the meetup in community, register guys are more than 700, attending guys are more than 350+. we invite not only rocketmq committer, but also storm committer, rocketmq contributor from community to speak at the meetup.
- \* Star has increased from 854 to 1660+ since the latest report, and RocketMQ website average weekly PV are over 11000+.
- \* Community project RocketMQ-Jms, RocketMQ-Console and RocketMQ-Flume release 1.0.0 version. Some new projects are tending to release the first version, including RocketMQ-MySQL and RocketMQ-Spark. Contributors in this repository have over 20.
- \* Polish RocketMQ website 2 round, adding some example for new feature, version compatible explanation and many trivial update.
- \* Up to now, 80% of integration projects are from community, we have been encouraging and helping community contributions around the world.

How has the project developed since the last report?

- \* Up to now, 183+ issues have been reported on JIRA and 98 have been resolved or closed. Since Now, 128 pull requests have been created and 93 pull requests have been merged or closed.
- \* Teams are sparing no effort in next version 4.2.0 development, which will support TLS, Log4j2, and implements OpenMessaging Specification.
- \* Community released the second version, including many important features, such as batched message, message filtering based on SQL 92, log appender and so on.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

[X] Initial setup - The initial setup went pretty well with assistance from John Ament and his Incubator documentation improvement efforts.

[X] Working towards first release - The podling has performed its first release as noted below

[X] Community building - Some contributors are beginning to participate. More users have been encouraged to resort to user mail list to asking for help or support, and more and more users are on the way.

[X] Nearing graduation - Making progress, not there yet.

[ ] Other:

Date of last release:

2017-06-02

When were the last committers or PPMC members elected?

- \* Last elected committers as following:
  - \*\* 2017.7.2 - Jaskey
- \* Last elected PMC member as following:
  - \*\* 2017.7.5 - dongeforever

Signed-off-by:

[X](rocketmq) Bruce Snyder

Comments:  
[ ](rocketmq) Brian McCallister  
Comments:  
[ ](rocketmq) Willem Ning Jiang  
Comments:  
[ ](rocketmq) Luke Han  
Comments:  
[X](rocketmq) Justin McLean  
Comments:

-----  
Rya

Rya (pronounced "ree-uh" /r/) is a cloud-based RDF triple store that supports SPARQL queries. Rya is a scalable RDF data management system built on top of Accumulo. Rya uses novel storage methods, indexing schemes, and query processing techniques that scale to billions of triples across multiple nodes. Rya provides fast and easy access to the data through SPARQL, a conventional query mechanism for RDF data.

Rya has been incubating since 2015-09-18.

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

1. Have more releases as part of the Apache Foundation
2. Increase diversity of contributors.
3. Continue to harden and develop core Rya features to improve user experience

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?

- \* Adina Crainiceanu presented two Rya talks at Apache Big Data North America on May 16-18 in Miami:
  1. a 10 minutes presentation at the Podling Shark Tank, and a
  2. full presentation on Apache Rya - A Scalable RDF Triple Store.
- \* Pull Requests and issue reports from non-committers which are integrated into the repository continue to be received

How has the project developed since the last report?

- \* Resolved a handful of issues that users have encountered
- \* GeoTemporal Indexing
- \* Expanded reasoning capabilities by implementing owl:hasValue inference
- \* Added construct query support, aggregation support (sums, counts, averages, max, min), and geo temporal filtering to Rya Fluo Incremental Updater
- \* Mongo backed entity centric index
- \* Merged pull request for singleton AccumuloMinicluster to speed up integration tests

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- [ ] Initial setup
- [ ] Working towards first release
- [X] Community building
- [ ] Nearing graduation
- [ ] Other:

Date of last release:

2016-10-28

When were the last committers or PPMC members elected?

- \* New committer and PPMC member David Lotts elected on Oct 25, 2016
- \* New PPMC member Caleb Meier elected on Jan 3rd, 2017

Signed-off-by:

[x](rya) Josh Elser  
Comments: Relatively quiet, not much happening on the community or code side. I see some contributions from non-committers in the past month. Need to push podling towards growing their member-base, and set some goals.  
[ ](rya) Edward J. Yoon  
Comments:  
[ ](rya) Venkatesh Seetharam  
Comments:  
[x](rya) Billie Rinaldi  
Comments:

IPMC/Shepherd notes:

David Fisher: From <http://rya.incubator.apache.org/download/> there is no release. No way for potential users to easily find the October release.

-----  
SensSoft

SensSoft is a software tool usability testing platform

SensSoft has been incubating since 2016-07-13.

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

1. Moving towards the first Incubating release of the source code and other release artifacts to NPM/Bower, etc.
2. Grow the Apache SensSoft (Incubating) community.
3. Complete the issues highlighted at the SensSoft Roadmap  
<https://cwiki.apache.org/confluence/display/SENSOFT/Roadmap>

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 community has been following a scrum method, sprinting towards the generation of the first release candidate for UserALE.js. This is showing a solid basis for growth and is working well. Josh Poore (project lead) recently demo'd SensSoft to a NASA Earth Science Data Systems Working Group focused on Search Relevancy. SensSoft is actively being evaluated and used within several prototypes at NASA JPL.

How has the project developed since the last report?

The project is making sustainable progress. Project members are driving towards generation of UserALE.js 0.1.0 which is great.

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

[X] Initial setup  
[X] Working towards first release  
[X] Community building  
[ ] Nearing graduation  
[ ] Other:

Date of last release:

XXXX-XX-XX

When were the last committers or PPMC members elected?

Arthi Vezhavendan was added to PPMC and Committer base on 2017-01-24

Signed-off-by:

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

-----  
Superset

Superset is an enterprise-ready web application for data exploration, data visualization and dashboarding.

Superset has been incubating since 2017-05-21.

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

1. Grow the community and enroll new committers.
2. Move the website to apache.org
3. Plan and execute our first Apache release

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

- \* The team is working on finishing up all the items on the Incubation Checklist.
- \* We're working on finding a way to link Jira Issues to GitHub Issues as the team has a preference for using GitHub Issues for tracking.

How has the community developed since the last report?

- \* Organic growth of our Github contributors (140), forks (2,066), watchers (743) and stars (14,773)

How has the project developed since the last report?

- \* The code repository has been migrated from <https://github.com/airbnb/superset> to <https://github.com/apache/infra/superset> on Apache's infrastructure
- \* A healthy, constant flow of bug fixes, quality improvements and new features, take a look at the project's Pulse on Github for more details

How would you assess the podling's maturity?

Please feel free to add your own commentary.

[X] Initial setup  
[ ] Working towards first release  
[ ] Community building  
[ ] Nearing graduation  
[ ] Other:

Date of last release:

- \* No official release yet since being voted into Apache Incubation. (Planning for the first Apache release in Q3, 2017)

When were the last committers or PPMC members elected?

- \* Project is still functioning with the initial set of committers

Signed-off-by:

[ ](superset) Ashutosh Chauhan  
Comments:  
[ ](superset) Luke Han  
Comments:  
[x](superset) Jim Jagielski  
Comments:

-----

## Traffic Control

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

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

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

1. Enhance automation to facilitate committer voting on new releases.
2. Enhance documentation to ease ramp-up time for new community members.
3. Grow the community.

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?

Traffic Control held a summit with Traffic Server during ApacheCon North America. This helped bring more attention to our project and encouraged cross-pollination between the two projects. In total there were 45 confirmed attendees with 23 signing up as having interest in Traffic Control.

Traffic Control added 5 new contributors (went from 65 to 70 total unique contributors).

How has the project developed since the last report?

Since the last report (April 2017), we have

- \* Merged 225 Pull Requests with 714 commits from 27 contributors
- \* Opened 189 JIRA issues
- \* Closed 97 JIRA issues
- \* 393 emails sent by 44 people on the dev@ list
- \* 15 emails sent by 9 people on the users@ list

How would you assess the podling's maturity?  
Please feel free to add your own commentary.

- ☐ Initial setup
- ☐ Working towards first release (wip)
- ☒ Community building (wip)
- ☐ Nearing graduation
- ☐ Other:

Date of last release:

Traffic Control 1.8.1 was released on 6/13/2017.  
Traffic Control 2.0 has passed the project team and IPMC vote and a release announcement should be sent out this week.

When were the last committers or PPMC members elected?

No new committers or PPMC members yet.

Signed-off-by:

- [X](trafficcontrol) Phil Sorber  
Comments:
- [X](trafficcontrol) Eric Covener  
Comments:
- [ ](trafficcontrol) Daniel Gruno  
Comments:
- [ ](trafficcontrol) J. Aaron Farr  
Comments:

IPMC/Shepherd notes:

johndament: The podling has grown quite a bit, once they start growing outside of their initial committer set they should think about graduation.

-----

Weex is a framework for building Mobile cross-platform high performance UI.  
Weex has been incubating since 2016-11-30.

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

1. Move the whole workflow and infra into Apache (not the codebase only).
2. Develop more contributors and committers.
3. First Release while in Incubator.
4. Improve the developer's activity in JIRA.

How has the community developed since the last report?

- We voted a new committer: Jianfeng Li
- Our Github has growth contributors (110), forks (517), watchers (222) and stars (3858)
- We have encouraged people to discuss more in the JIRA and dev mail list. Up to now, 21 issues are reported on JIRA and 15 of them are resolved.
- WebSite Trends Month Active , PV(295160) , UV(39778) , IP(36988), New Unique Visitor(20700) , Session(68747)

How has the project developed since the last report?

- Excluding merges, 19 authors have pushed 33 commits to master and 240 commits to all branches.
- This month we have successfully released the 0.12 version.

How would you assess the podling's maturity?

Please feel free to add your own commentary.

- [ ] Initial setup
- [ ] Working towards first release
- [x] Community building
- [ ] Nearing graduation
- [ ] Other:

Date of last release:

2017-06-08

When were the last committers or PPMC members elected?

Jianfeng Li, 16 June 2017

Signed-off-by:

- [ ](weex) Luke Han  
Comments:
- [ ](weex) Willem Jiang  
Comments:
- [ ](weex) Stephan Ewen  
Comments:
- [x](weex) Raphael Bircher  
Comments:  
I started during the first release, and at this time there was nearly no communication on the list. We increased the list traffic last month over 100% but there is still a load of work to do.

IPMC/Shepherd notes:

David Fisher: I want to note that Raphael is stepping up efforts as Mentor to really help the Weex community. Seems to be effective.

jda: manual reminder sent. report drafted on dev list