

July2018

Incubator PMC report for July 2018

FINAL

Timeline

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

Shepherd Assignments

Dave Fisher	Gossip		
Dave Fisher	ODF Toolkit		
Drew Farris	Livy		
Drew Farris	ServiceC omb		
John Ament	NetBeans		
John Ament	PLC4X		
Justin Mclean	S2Graph		
P. Taylor Goetz	BatchEE		
P. Taylor Goetz	Milagro		
P. Taylor Goetz	Traffic Control		
Timothy Chen	Annotator		
Timothy Chen	HAWQ		
Timothy Chen	SensSoft		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b27193b2-877f-4d09-8bf2-a2bafec5a4a9"><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="e4d90618-5ecd-4bf3-b6c0-0c04e1b925ad"><ac:plain-text-body><![CDATA[[none]	Amaterasu]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="f683a35b-0a54-4677-bb54-f2470b67b0b1"><ac:plain-text-body><![CDATA[[none]	Goblin]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="be6d17be-5f9b-45d0-85a8-fa55802de089"><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="2d25e768-eb0b-4773-a907-17ac8fed729f"><ac:plain-text-body><![CDATA[[none]	Rya]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7e2b2ffd-1bba-4362-9218-eab56517f9a7"><ac:plain-text-body><![CDATA[[none]	Weex]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for July 2018

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 presently 52 podlings incubating. During the month of June, podlings executed 6 distinct releases and there are several ongoing votes and releases for July. We added no new IPMC members. No podlings are graduating this month but there's discussion about HAWQ graduating, and many are heading towards graduation. We had one project, AriaTosca,

retire to the attic. There were no IP clearances.

A number of podlings didn't report and that may be due to less reminders sent out this month. The three podlings that failed to report last month also failed to report this month. Two (Druid and Spot) to why last month, lack of activity and response from Pony Mail is concerning.

Of note in June on the list there was discussions on

- retirement and what to do with gitHub/gitbox repos
- Gossip exiting the incubator

* Community

New IPMC members:

- None

People who left the IPMC:

- None

* New Podlings

- Doris

* Podlings that failed to report, expected next month

- Airflow
- Annotator
- BatchEE
- Druid
- Gossip
- Milagro
- Pony Mail
- S2Graph
- SensSoft
- Spot

* Graduations

- None

* Releases

The following releases entered distribution during the month of June:

- 2018/07/02 Goblin 0.12.0
- 2018/06/22 Servicecomb 1.0
- 2018/06/17 Pulsar 1.22.1
- 2018/06/15 Singa 1.2.0
- 2018/06/07 Traffic Control 2.2.0
- 2018/06/05 Dubbo 2.6.2

The Traffic Control release seems odd as they are a TLP. I'll follow up on that.

* IP Clearance

None

* Legal / Trademarks

None

* Infrastructure

None

* Miscellaneous

None

* Credits

Table of Contents

Amaterasu
Gobblin
Gossip
HAWQ
Livy
MXNet
NetBeans
ODF Toolkit
PLC4X
Rya
ServiceComb
Traffic Control
Weex

Amaterasu

Apache Amaterasu is a framework providing configuration management and deployment for Big Data Pipelines.

It provides the following capabilities:

Continuous integration tools to package pipelines and run tests.
A repository to store those packaged applications: the applications repository.
A repository to store the pipelines, and engine configuration (for instance, the location of the Spark master, etc.): per environment - the configuration repository.
A dashboard to monitor the pipelines.
A DSL and integration hooks allowing third parties to easily integrate.

Amaterasu has been incubating since 2017-09.

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

1. Prepare the first release
2. Grow up user and contributor communities
3. Prepare documentation

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?

- * Two conference talks have been delivered (PyCon 11 and SDP)
- * Initial documentation has been created, targeted for Amaterasu's next release

How has the project developed since the last report?

- * since the last report 4 release candidates have been built and voted upon, at the time of this report the last RC (4) has passed the vote on in the general@incubator mailing list
- * Two additional contributors started contributing to the code base
- * One more organization we are aware of have started a POC with Amaterasu

Date of the last release:

N/A

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](amaterasu) Jean-Baptiste Onofré

Comments:

[](amaterasu) Olivier Lamy

Comments:

[X](amaterasu) Davor Bonaci

Comments: Activity has picked up recently, and the first release is coming. That said, I think the community is not developing fast enough to suggest a good fit for the Foundation. I have general concerns whether the project can graduate and how quickly.

IPMC/Shepherd notes:

Gobblin

Gobblin is a distributed data integration framework that simplifies common aspects of big data integration such as data ingestion, replication, organization and lifecycle management for both streaming and batch data ecosystems.

Gobblin has been incubating since 2017-02-23.

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

1. Make frequent releases

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?

* Various major components and futuristic features are being driven by the community

(non-committers) thus building a very healthy pool of contributors that can be voted

in as committers.

* Continued growth in engagement over Gitter IRC, and mailing lists.

* 79% of commits (a record in Gobblin community) were from non-committer contributors.

* Email stats since last report:

user@gobblin.incubator.apache.org : 44 dev@gobblin.incubator.apache.org : 200

* There have been 66 Commits since last report:

git log --format='%ci' | grep -cE '(2018-0(4|5|6|7))'

* 52 ie. 79% of those commits were by non-committers:

git log --format='%ae %ci' | grep -E '(2018-0(4|5|6|7))' | cut -d ' ' -f 1
| sort | uniq -c | sort -n

* Recurring video conference based meetup has been happening every month with a

healthy attendance.

* Gobblin was presented and well received in various meetups / conferences around the world (independently by Apache community members).

How has the project developed since the last report?

* Major progress in Gobblin's evolution as Platform-as-a-Service - being driven

by couple of non-committers.

* Comprehensive work being driven by a non-committer for stability of Gobblin

cluster at extreme scale.

* Enhancements to key integrations such as Salesforce, Couchbase, Kafka, etc.

* Addition of features for compliance and security. Increased adoption in this

area by the community (for critical use-cases such as GDPR).

* Release 0.12.0 done.

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:

2018-07-02

When were the last committers or PPMC members elected?

Joel Baranick in December, 2017.

Signed-off-by:

[X](gobblin) Jean-Baptiste Onofre
Comments:
[X](gobblin) Olivier Lamy
Comments:
[X](gobblin) Jim Jagielski
Comments:

IPMC/Shepherd notes:

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.
- 2.
- 3.

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

We opened up a discussion thread to move Gossip out of the incubator. It was initially -1 ed by a committer. There is a separate discussion on the main incubator list about our lack of timely reporting. After a short discussion we agreed we not want to be an undue burden and wanted to give the IPMC time back so they can follow up with the other podlings. We will take an official vote on Monday to close down the podling.

How has the community developed since the last report?

How has the project developed since the last report?

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:

XXXX-XX-XX

When were the last committers or PPMC members elected?

Signed-off-by:

[](gossip) P. Taylor Goetz
Comments:
[x](gossip) Josh Elser
Comments: Still trying to encourage retirement. The ASF and its
processes are more of a burden than a boon at this point.
[](gossip) Drew Farris
Comments:

IPMC/Shepherd notes:

There is a discussion on dev@ about retirement of this podling.
<https://lists.apache.org/thread.html/8acab107a15a6bc50253cb20aa633a3ac36b70364294f149660fc6c4@%3Cdev.gossip.apache.org%3E>
Dave Fisher

HAWQ

HAWQ is an advanced enterprise SQL on Hadoop analytic engine built around a robust and high-performance massively-parallel processing (MPP) SQL engine evolved from Greenplum Database.

HAWQ has been incubating since 2015-09-04.

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

Nothing at this time.

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 :

* HAWQ for data scientists, Shanghai CIO forum (Speaker: Lei Chang, June 23, 2018)

2. "Graduate Apache HAWQ (incubating) as a TLP" discussion was started:
<https://lists.apache.org/thread.html/67a2d52ef29cbf9e93d8050ed0193cc110a919962dd92f8436b343b7@%3Cdev.hawq.apache.org%3E>

The followings is the content of the thread:

With the 2.3.0.0-incubating release officially out, the Apache HAWQ community and its mentors believe it is time to consider graduation to the TLP:
<https://lists.apache.org/thread.html/b4a0b5671ce377b3d51c9b7ab00496a1eebfcbf1696ce8b67e078c64@%3Cdev.hawq.apache.org%3E>

Apache HAWQ entered incubation in September of 2015, since then, the HAWQ community learned a lot about how to do things in Apache ways. Now we have a healthy and engaged community, ready to help with all questions from the HAWQ community. We delivered four releases including two binary releases, now we can do self-driving releases in good cadence. The PPMC has demonstrated a good understanding of growing the community by electing 12 individuals as committers and PPMC members. The PPMC addressed the maturity issues one by one followed by Apache Project Maturity Model, currently all the License and IP issues are resolved. This demonstrated our understanding of ASF's IP policies.

All in all, I believe this project is qualified as a true TLP and we should recognize this fact by formally awarding it such a status. This thread means to open up the very same discussion that we had among the mentors and

HAWQ community to the rest of the IPMC. It is a DISCUSS thread so feel free to ask questions.

To get you all going, here are a few data points which may help:

Project status:

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

Project website:

<http://hawq.incubator.apache.org/>

Project documentation:

<http://hawq.incubator.apache.org/docs/userguide/2.3.0.0-incubating/overview/HAWQOverview.html>

<http://hawq.apache.org/#download>

Maturity assessment:

<https://wiki.apache.org/confluence/display/HAWQ/ASF+Maturity+Evaluation>

DRAFT of the board resolution is at the bottom of this email

Proposed PMC size: 45 members

Total number of committers: 45 members

PMC affiliation (* indicated chair):

Pivotal (20)

* Oushu (7)

Amazon (3)

Hashdata (2)

Autonomic (1)

Confluent (1)

Datometry (1)

Hortonworks (1)

Microsoft (1)

PETUUM (1)

Privacera (1)

Qubole (1)

Snowflake (1)

State Street (1)

Unifi (1)

Visa (1)

ZEDEDA (1)

1549 commits on develop

1375 PR"s on GitHub

63 contributors across all branches

1624 issues created

1350 issues resolved

dev list averaged ~53 msgs/month over last 12 months

user list averaged ~6 msgs/month over last 12 months

129 unique posters

committer affiliations:

active

pivotal.io

oushu.io

hashdata.cn

occasional

amazon.com

autonomic.ai

confluent.io

datometry.com

hortonworks.com

microsoft.com

petuum.com

privacera.com

qubole.com

snowflake.net

statestreet.com
unifisoftware.com
visa.com
zededa.com

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:

2018-03-12, Apache HAWQ 2.3.0.0

When were the last committers or PPMC members elected?

- 1) Lav Jain: April 5, 2018
- 2) Shubham Sharma: April 5, 2018

Signed-off-by:

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

IPMC/Shepherd notes:

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:

1. Grow the community to have more activities.
2. Grow more contributors and committers.
3. Regular release.

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

None

How has the community developed since the last report?

Activity on the dev and user mailing lists as well as the rate of reported JIRA issues has stayed steady.

How has the project developed since the last report?

Two security issues are fixed and support version is upgraded to Spark 2.3.1.

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:

2018-02-06

When were the last committers or PPMC members elected?

2017-09-18

Signed-off-by:

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

IPMC/Shepherd notes:

MXNet

Apache MXNet is a lean, flexible, and ultra-scalable deep learning framework that supports state of the art in deep learning models, including convolutional neural networks (CNNs) and long short-term memory networks (LSTMs). The framework has its roots in academia and came about through the collaboration and contributions of researchers at several top universities. It has been designed to excel at computer vision, speech, language processing and understanding, generative models, concurrent neural networks, and recurrent neural networks. MXNet allows you to define, train, and deploy networks across a wide array of use cases from massive cloud infrastructure to mobile and connected devices. It provides a very flexible environment with support for many common languages and the ability to utilize both imperative and symbolic programming constructs. MXNet also very lightweight. This allows it to scale across multiple GPUs and multiple machines very efficiently, which is beneficial when conducting training on large datasets in the cloud.

Apache MXNet has been incubating since 23-Jan, 2017.

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

1. Establish a predictable release process consistent with Apache Way -- ESTABLISHED.
2. Grow the community -- ONGOING.
3. Bring website up to Apache standard - COMPLETED.
4. Identify remaining ICLAs or SGAs that need signing - NEAR COMPLETION.

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) Various Slack channels, dev@ mailing lists, and user discussion forums (<http://discuss.mxnet.io>) are being used actively. The contributors have been working on having all discussions on the public dev@ mailing list as much as possible. At least one of the discussions infringed on Apache code of conduct and mentors had to step in and provide guidance to the community.

b) After lengthy discussion the community agreed to establish user@mxnet.incubator.apache.org with the goal to increase users as well as grow the contributor community.

c) Community events:

- In April the community organized the first MXNet meetup in Seattle. We had over 100 attendees, 15% from outside Amazon.
- In May the community organized a virtual hangout with two sessions to collect feedback and ideas from the community.
- Korean MXNet community is growing well and most of contributors are outside of Amazon:
<https://www.meetup.com/ko-KR/MXNet-Korea-User-Group/events/248531181/> . The Facebook group has 530 members in <https://www.facebook.com/groups/mxnetkr>
- Community members presented at https://www.meetup.com/Artificial-Intelligence-in-Practice/events/250882931/?_xtid=gqFyqDI5NDE5MTUyoXCkaXBhZA&from=ref
- Community members presented at meetups in Vancouver (link see under Webinars)
- Community members presented Intro to Amazon SageMaker with a sentiment analysis demo using MXNet and Gluon at AWS Loft ML Day in SF on 6-19 to ~150 attendees
- Community members presented Model serving with Model Server for MXNet at AWS Loft ML Day in SF on 6-19 to ~150 attendees
- A committer presented Distributed inference with Spark and MXNet at Apache Roadshow in Berlin, 13-14 June (talk abstract - <https://foss-backstage.de/session/distributed-inference-using-apache-mxnet-and-apache-spark>)
- A committer presented Distributed inference with Spark and MXNet at Spark+AI in SF, ~100 attendees (<https://databricks.com/speaker/naveen-swamy>)
- Community members presented Model serving with MXNet at dotAI conference in Paris, France, ~300 attendees, May 30 (<https://www.dotconferences.com/2018/05/hagay-lupesko-model-serving-for-deep-learning>)
- Community members presented ONNX with MXNet demo at prepareAI conference in St Louis, May 8, ~100 attendees (<http://prepare.ai/conference/conference-agenda-details/>)
- Community members presented Model serving with MXNet at prepareAI conference in St Louis, May 8, ~100 attendees (<http://prepare.ai/conference/conference-agenda-details/>)
- Community members presented Introduction to deep learning with MXNet and Gluon, distinguished lecture, at BGU university in Israel, April 10, ~100 attendees (<https://www.cs.bgu.ac.il/~frankel/Lupesko/Lupesko.pdf>)
- Community members presented Model serving with MXNet at AI IL meetup in Tel Aviv, Israel, ~50 attendees (<https://www.meetup.com/artificial-intelligence-il/events/249312879/>)

d) Blogs about MXNet: We established MXNet blog:

- <https://medium.com/apache-mxnet>. As of today, the blog has 634 followers and 16 blogs published. Additionally, blogs have been published related or about MXNet at:
- <https://aws.amazon.com/blogs/machine-learning/category/artificial-intelligence/apache-mxnet-on-aws/>
 - <https://blogs.technet.microsoft.com/machinelearning/tag/mxnet>
 - <https://zh.mxnet.io/blog/>
 - <https://blogs.apache.org/mxnet>
 - <https://github.com/chinakook/Awesome-MXNet>
 - <https://medium.com/mlreview/10-deep-learning-projects-based-on-apache-mxnet-8231109f3f64>
 - <https://medium.com/datreio/training-with-keras-mxnet-on-amazon-sagemaker-43a34bd668ca>
 - <https://medium.com/@julsimon/a-quick-look-at-the-swish-activation-function-in-apache-mxnet-1-2-79d9ff9d1673>
 - <http://gigasquidsoftware.com/blog/2018/06/03/meet-clojure-mxnet-ndarray/>
 - <https://dzone.com/articles/ingesting-apache-mxnet-gluon-deep-learning-results>
 - <https://becominghuman.ai/an-introduction-to-the-mxnet-api-part-1-848febdcf8ab>
 - <https://cosminsanda.com/posts/counting-object-with-mxnet-and-sagemaker/>
 - CVPR 2018 paper - Relation Networks for Object Detection - <https://arxiv.org/abs/1711.11575> and

<https://github.com/msracver/Relation-Networks-for-Object-Detection>
 - Gluon-CV related blogpost -
<https://dzone.com/articles/using-apache-mxnet-gluoncv-with-apache-nifi>
 -
<https://aws.amazon.com/blogs/machine-learning/maximize-training-performance-with-gluon-data-loader-workers/>
 -
<https://medium.com/apache-mxnet/accelerating-deep-learning-on-cpu-with-intel-mkl-dnn-a9b294fb0b9>
 -
<https://medium.com/apache-mxnet/mxnet-for-pytorch-users-in-10-minutes-a7353863406a>
 -
<https://medium.com/apache-mxnet/gluoncv-deep-learning-toolkit-for-computer-vision-9218a907e8da>
 -
<https://medium.com/apache-mxnet/mxboard-mxnet-data-visualization-2eed6ae31d2c>
 - <https://medium.com/apache-mxnet/mxnet-gluon-in-60-minutes-3d49eccaf266>
 - <https://medium.com/apache-mxnet/announcing-apache-mxnet-1-2-0-d94f56da154b>
 -
<https://medium.com/apache-mxnet/mxnet-1-2-adds-built-in-support-for-onnx-e2c7450ffc28>
 -
<https://aws.amazon.com/blogs/machine-learning/the-importance-of-hyperparameter-tuning-for-scaling-deep-learning-training-to-multiple-gpus/>
 -
<https://medium.com/apache-mxnet/image-classification-with-mxnet-scala-inference-api-8ab6ce1bbccf>
 -
<https://medium.com/apache-mxnet/object-detection-with-mxnet-scala-inference-api-9049230c77fd>
 -
<https://medium.com/apache-mxnet/scala-api-for-deep-learning-inference-now-available-with-mxnet-v1-2-bcb13235db95>
 -
<https://medium.com/apache-mxnet/train-using-keras-mxnet-and-inference-using-mxnet-scala-api-49476a16a46a>
 - <https://medium.com/apache-mxnet/page-segmentation-with-gluon-dcb4e5955e2>
 - <https://medium.com/apache-mxnet/announcing-keras-mxnet-v2-2-4b8404568e75>
 - <https://devblogs.nvidia.com/tensor-core-ai-performance-milestones/>
 -
<https://aws.amazon.com/blogs/machine-learning/use-pre-trained-models-with-apache-mxnet/>

e) Webinars, Technical talks and lectures about MXNet:
 - Thomas Delteil presented at meetups in Vancouver:
https://www.youtube.com/watch?v=RgIa3_BjGyk&t=163s ,
<https://www.youtube.com/watch?v=mN15vKIyfoA> and
<https://www.youtube.com/watch?v=Kl20xBnY6OA>
 - Gluon debugging: https://www.youtube.com/watch?v=6-d0oJVw9_0
 - TVM stack: <https://www.youtube.com/watch?v=DaCPJrTwT00>
 - https://mxnet.incubator.apache.org/tutorials/vision/cnn_visualization.html
 - visual search: <https://www.youtube.com/watch?v=9a8MATfFVwI>
 - Cifar 10 super convergence: <https://www.youtube.com/watch?v=00XTkQPkUio>
 - Sparse Tensors: <https://www.youtube.com/watch?v=smZfsYhDFkY>
 - Mixed precision training with MXNet:
<https://www.youtube.com/watch?v=pR4KMh1lGC0>
 - See all videos on <https://www.youtube.com/apachemxnet> -
 - Thomas Delteil presented a lecture on MXNet Gluon and Deep Learning at the Machine Learning Summer School in Buenos Aires in front of 170 students.
<http://mlss2018.net.ar/>

How has the project developed since the last report?

a) The community released MXNet 1.2 with significant feature enhancements:
 1. Scala Inference API
 2. ONNX model import
 3. Model Quantization with calibration
 4. MKL-DNN Integration
 5. Improved exception handling for operators

6. Enhanced FP16 support
7. Profiling enhancements
Detailed release notes are provided on Apache Wiki and Github.
See
<https://cwiki.apache.org/confluence/display/MXNET/Apache+MXNet+%28incubating%29+1.2.0+Release+Notes>

b) Github: In April the project had 523 contributors. Excluding merges, 62 authors have pushed 199 commits to master and 258 commits to all branches. In May the project had 531 contributors. Excluding merges, 55 authors have pushed 141 commits to master and 177 commits to all branches. In June the project had 550 contributors. Excluding merges, 66 authors have pushed 171 commits to master and 221 commits to all branches. We are working on encourage more contributors to the project.

b) The community voted to adopted Jira for issue management. The change is methodology is partially adopted. On June 8th a committer started a vote to stop using Jira which created a passionate debate and is still ongoing.

c) List of design proposals published on MXNet Apache Wiki : All design proposals have been or are being discussed on dev@mxnet.apache.org. Four design proposals originated from external contributors. See
<https://cwiki.apache.org/confluence/display/MXNET/Design+Proposals>.

d) Ecosystem development:

1. MXNet Model Serving:
- Repo: <https://github.com/aws-labs/mxnet-model-server>
-
<https://medium.com/apache-mxnet/model-server-for-apache-mxnet-adds-support-for-gluon-models-dd398a710f94>

2. ONNX support
- Repo: <https://github.com/onnx/model-zoo>
- Repo: <https://github.com/onnx/models>
-
<https://medium.com/apache-mxnet/mxnet-1-2-adds-built-in-support-for-onnx-e2c7450ffc28>

3. Keras: The community release MXNet backend for Keras, a high level API for deep learning.
- Repo - <https://github.com/aws-labs/keras-apache-mxnet/>
-
<https://medium.com/@julsimon/apache-mxnet-as-a-backend-for-keras-2-9993f97843e7>
- <https://medium.com/apache-mxnet/announcing-keras-mxnet-v2-2-4b8404568e75>
-
<https://aws.amazon.com/blogs/machine-learning/apache-mxnet-incubating-adds-support-for-keras-2/>
-
<https://medium.com/apache-mxnet/keras-gets-a-speedy-new-backend-with-keras-mxnet-3a853efcd75>

e) Documentation: We continue to improve the documentation on Architecture guides, How To's, Tutorials, and APIs continue to be improved. 43 new public tutorials related and about MXNet have been published YTD.
- https://thomasdelteil.github.io/CNN_NLP_MXNet/
- <https://gluon-crash-course.mxnet.io/>
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/python/types_of_data_augmentation.md
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/python/data_augmentation_with_masks.md
-
<https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/gluon/datasets.md>
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/onnx/inference_on_onnx_model.md
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/onnx/finetuning_gluon.md
-

<https://aws.amazon.com/blogs/machine-learning/speeding-up-apache-mxnet-using-the-nnpack-library/>
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/python/data_augmentation.md
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/gluon/data_augmentation.md
-
<https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/gluon/naming.md>
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/gluon/custom_layer.md
- https://github.com/ThomasDelteil/VisualSearch_MXNet
-
https://github.com/apache/incubator-mxnet/blob/master/docs/tutorials/gluon/pretrained_models.md
- https://mxnet.incubator.apache.org/tutorials/gluon/save_load_params.html
- <http://mxnet.incubator.apache.org/tutorials/python/profiler.html>
-

How would you assess the podling's maturity?

Podling is still having difficulties to grow the contributor and committer community. Maturity == Low to Medium.

Please feel free to add your own commentary.

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

Date of last release: (latest to oldest releases)

1. Apache MXNet-incubating 1.2.0 (major release) was published on May 21 2018 - <https://github.com/apache/incubator-mxnet/releases/tag/1.2.0>

When were the last committers or PPMC members elected?

During the reporting period 1 contributor was elected as committer. The PPMC elected Jim Jagielski as mentor.

Signed-off-by:

[X](mxnet) Sebastian Schelter
Comments:
[](mxnet) Suneel Marthi
Comments:
[](mxnet) Markus Weimer
Comments:
[](mxnet) Henri Yandell
Comments:
[X](mxnet) Jim Jagielski
Comments:

IPMC/Shepherd notes:

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. Brand and domain donation.

2. 3rd donation, consisting of documentation, e.g., tutorials and related images.

3. Handing over of netbeans.org to Apache and related rerouting.

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

None.

How has the community developed since the last report?

- Currently 588 (557 in last report) on Apache NetBeans users mailing list and 422 (406 in last report) on Apache NetBeans dev mailing list.
- New splash screen discussed and accepted.
- NetBeans represented at FossBack 2018:
<https://blogs.apache.org/netbeans/entry/apache-netbeans-at-foss-backstage>
- NetBeans represented at COFAN 2018:
<https://blogs.apache.org/netbeans/entry/netbeans-for-javaee-project-productivity>

How has the project developed since the last report?

- May 2018: Apache NetBeans (incubating) 9.0 RC1 released:
<https://netbeans.apache.org/download/index.html>
- Hats off to Emilian Bold, the first non-Sun/Oracle release manager, who put the RC1 mentioned above together.
- 2nd Oracle donation consisting of 1.5M lines of code grant accepted and code handed over:
<https://github.com/apache/incubator-netbeans/tree/2ndDonation>

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:

2018-05-28 Apache NetBeans (incubating) 9.0 RC1

When were the last committers or PPMC members elected?

- * PPMC member Laszlo Kishalmi elected on June 12, 2018
- * PPMC member Reema Taneja elected on June 29, 2018

Signed-off-by:

[](netbeans) Ate Douma
Comments:
[X](netbeans) Bertrand Delacretaz
Comments:
[](netbeans) Daniel Gruno
Comments:
[](netbeans) Mark Struberg
Comments:

IPMC/Shepherd notes:

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. Improve code generation from ODF XML grammar to reduce complexity (sponsorship by German Government next Winter)
2. Attract more developers - (collaboration scheduled for the end of October)
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?
Several bug reports yet no new committers.

How has the project developed since the last report?
Svante Schubert won for a second time sponsorship by German government's prototypefund.org for next Winter.
His previous collaboration feature runs finally full-featured with green regression tests on his GitHub fork, and was not yet merged into the main branch.

Merge needs to be done before October, so the new sponsorship can focus on source code generation from the ODF XML grammar loaded into Graph Database (Apache Tinkerpop).

Date of the last release:

2017-04-10

When were the last committers or PMC members elected?

2017-08-02 - Michael Stahl

Signed-off-by:

[](odftoolkit) Nick Burch
Comments:
[](odftoolkit) Yegor Kozlov
Comments:
[X](odftoolkit) Tom Barber
Comments:

The podling continues good development under Svante's guidance, but is certainly lacking other committers. That said, there is an active community with people responding to questions and comments on the mailing list. I hope the move to Gitbox helps.

IPMC/Shepherd notes:

ODF Toolkit has moved source code development to GitBox/GitHub. The website remains in SVN as it is tied together with Apache CMS.
Given that this pooling has been in Incubation for seven years it is worth considering retirement. Just yesterday there was a question about how ready it is for production given that it has "--incubating" in the package names.
Dave Fisher

PLC4X

PLC4X is a set of libraries for communicating with industrial programmable logic controllers (PLCs) using a variety of protocols but with a shared API.

PLC4X has been incubating since 2017-12-18.

Most important issues to address while moving towards graduation:

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

1. Building the community: The PPMC and committer group has a large percentage of codecentric employees, we have been recruiting people from other companies, but will have to continue these efforts for establishing a healthy Apache community.

2. Onboarding of new committers: With PLC4X several people on the team are not very familiar with the Apache Way. We have started and will continue our efforts on this onboarding.

3. Make our first release

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

Not at the moment.

How has the community developed since the last report?

Christofer has invested most of his time in May and June spreading the word about PLC4X.

There have been several conference talks on PLC4X:

- * Building IoT 2018 in Cologne: Christofer Dutz - Industrial IoT und Edge Computing mit Apache PLC4X und Apache Edgent
- * @Vanced Network Meetup in Solingen: Christofer Dutz - Open-Source in der Industrie mit Apache PLC4X
- * Apache Roadshow EU in Berlin: Christofer Dutz - Revolutionizing the Production Industry with Apache PLC4X

A blog post in a major corporate tech-blog has been published with the content of the JavaSPEKTRUM article on Apache Edgent and Apache PLC4X: <https://blog.codecentric.de/2018/06/edge-computing-industrial-iot-apache-edgent-apache-plc4x/>

As a direct result on a talk in Berlin about mentoring new committers, we have started creating "low-hanging-fruit" tagged issues and have started promoting them. As a result we have 2 potential contributors interested in tackling these. Hopefully this will result in continued activity from their side.

We have reached out to a developer offering Modbus, EthernetIP and OPC-UA drivers on his Github page to get him to join our efforts.

The dev-list subscriptions has gone up by 7 (now 29) from the last report.

We now have a Twitter account: <https://twitter.com/ApachePlc4x> as another channel of communicating interesting PLC4X facts. After tweeting a lot about our project, the number of followers to that has drastically increased to 41

How has the project developed since the last report?

We have setup a VPN for the CI server and other committers to be able to communicate with real PLCs, however this feature has not proven to improve the involvement of other committers so far.

We have recently found a set of git repos with Apache licensed code for accessing Modbus, Ethernet/IP and OPC-UA devices. We have reached out to the original developer to maybe get him on board with our team.

How would you assess the podling's maturity?

We have a mix of new participants and experienced Apache people involved. So far, the new participants have shown great willingness and success in adopting the Apache Way. However, we still need to continue: the on-boarding increasing the diversity of the team Also, will we need to decide and establish all the processes involved in releasing software at Apache

Please feel free to add your own commentary.

- [] Initial setup
- [X] 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:

☐ (plc4x) Greg Trasuk
Comments:
☒ (plc4x) Justin Mclean
Comments:
☒ (plc4x) Stefan Bodewig
Comments:

IPMC/Shepherd notes:

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. Increase diversity of committers

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?

- * Moving towards graduation

How has the project developed since the last report?

- * Started major version change Apache Rya 4.0.0 (incubating)
- * Migrated from using older OpenRDF Sesame libraries to newer RDF4J
- * Fixed various bugs related to the use of RDF4J libraries (e.g. language support for literals)
- * Added tests for RyaStreams join iterator
- * Moved Geo Indexing specific maven repositories into the geoindexing profile, so geo-related jars are only fetched when geoindexing profile is used
- * Repackaged reusable Mongo DB test code in a new project rya.test.mongo
- * Implemented Kafka Connect Sink implementations for Accumulo and MongoDB backed Rya

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:

2018-03-04

When were the last committers or PPMC members elected?

* PPMC member Caleb Meier elected on Jan 3rd, 2017

Signed-off-by:

[](rya) Josh Elser
Comments:
[](rya) Edward J. Yoon
Comments:
[](rya) Venkatesh Seetharam
Comments:
[x](rya) Billie Rinaldi
Comments:

IPMC/Shepherd notes:

ServiceComb

ServiceComb is a microservice framework that provides a set of tools and components to make development and deployment of cloud applications easier.

ServiceComb has been incubating since 2017-11-22.

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

1. Community building
2. Project visibility
3. Trademark transfer

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?(2018-04)

- * The community get the second release out:
service-center 1.0.0-m2, java-chassis 1.0.0-m2 and saga 0.2.0
- * The community hold a meetup in Beijing at Jun 27.

<http://servicecomb.incubator.apache.org/docs/apache-servicecomb-incubating-day-report/>

* There are about 6 departments of Huawei are using ServiceComb and other 8 companies are using ServiceComb.

How has the project developed since the last report? (2018-04)

- * Java Chassis merged with 102 PRs by 16 contributors.
- * Service Center merged with 58 PRs by 6 contributors .
- * Saga merged with 48 PRs by 14 contributors.
- * 70 mails send by 19 people on the dev@ list (2018-04)
- * 156 mails sent by 17 people on the dev@ list 2018-05
- * 97 mails send by 23 people on the dev@ list (2018-06)

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:

2018-06-22 Java Chassis 1.0.0-m2
2018-06-22 Service Center 1.0.0-m2
2018-06-22 Saga 0.2.0

When were the last committers or PPMC members elected?

2018-06-08 Zheng yangyong was elected as committer.

2018-06-11 Li Dagang(Eric) was elected as committer.

Signed-off-by:

[X](servicecomb) Roman Shaposhnik

Comments:

[X](servicecomb) Jean-Baptiste Onofre

Comments:

[](servicecomb) Timothy Chen

Comments:

IPMC/Shepherd notes:

Weex

Weex is a framework for building high-performance mobile applications with modern web development experience.

Weex has been incubating since 2016-11-30.

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

1. Develop more non-Chinese contributors and committers, to bring diversity to the community.

2. Improve the developer's activity in the mailing list.

3. Create regularly release schedule.

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?

* Forward to the second release, it's under RC2 now.

* Set up a regular monthly meeting between Alibaba, Tencent, and PayTM developers.

* 493(213 new issues since the last report) issues have been reported on JIRA, and 321(79 since the last report) of them are resolved or closed.

* Our Github repo has growth in contributors 154(9 new since the last report), forks 1321(153 new), watchers 497(46 new) and stars 9985(1255 new)

How has the project developed since the last report?

* Excluding merges, 20 authors have pushed 77 commits to master and 93 commits to all branches. (June 4, 2018 - July 4, 2018)

* Refactor the thread model of JavascriptCore and integrate WeexCore call to iOS.

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?

None.

Signed-off-by:

[](weex) Luke Han

Comments:

[X](weex) Willem Jiang

Comments:

[](weex) Stephan Ewen

Comments:

[](weex) Raphael Bircher

Comments:

IPMC/Shepherd notes: