

June2019

FINAL DO NOT EDIT

NOTE that the content for this month report is in markdown format- please do not remove any formatting characters when adding your reports.

It is also best to:

- Keep all lines under 76 characters long.
- All content under the ### headings should be indented by two spaces.
- Please don't change the text in the headings or add new ones.
- Include a space after a bullet point or full stop on a numbered list.
- Use [X] (no spaces) to sign off reports.

A script will be run on this report before it is submitted to the board. If your lines are too long it may insert line breaks in unexpected places.

Reflow

```
fold -w 76 -s $1 | cat -s
```

Timeline

Wed June 05	Podling reports due by end of day
Sun June 09	Shepherd reviews due by end of day
Sun June 09	Summary due by end of day
Tue June 11	Mentor signoff due by end of day
Wed June 12	Board meeting

Shepherd Assignments

Dave Fisher	SAMOA
Dave Fisher	TVM
Drew Farris	Druid
Drew Farris	IoTDB
Justin Mclean	Taverna
P. Taylor Goetz	BRPC
P. Taylor Goetz	Myriad
P. Taylor Goetz	Tephra
Timothy Chen	Nemo
Timothy Chen	Spot
Timothy Chen	Warble
	Craill
	DLab
	Daffodil
	DataSketches

Daniel Dai	Hivemall
	Iceberg
	Marvin-AI
	Milagro
	Omid
	OpenWhisk
	Pony Mail
	SINGA
	ShardingSphere
	Superset
	Tuweni
	Zipkin

Report content

Incubator PMC report for June 2019

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

At the request of the board, this monthly report is in markdown so that it's easier to read. Some podlings struggled to get their content in the correct format. If you are viewing the report not in that format it can be seen here: <https://cwiki.apache.org/confluence/display/INCUBATOR/June2019>

There are presently 49 podlings incubating. During the month of May, podlings executed 12 distinct releases. We added one new IPMC member and three IPMC members retired.

We have no new podlings this month. One project graduated last month Dubbo and a couple of other podlings are heading towards graduation in the next few months.

Milagro which failed to report twice, reported this month (to be confirmed) and is again trying to reboot themselves, they seem a bit more serious this time. Toree's report, also missing last month, did report. 4 Podlings BRPC, DLab, and Iceberg didn't submit a report and have been asked to do so next month.

Samoa is considering retiring.

There were no IP clearances.

The IPMC requires clarity from the board if podlings need to follow ASF's release policy and distribution policy while in incubation.

Some Background

What is happening in practice is that the IPMC allows releases with minor issues to go ahead and assumes they will be fixed in a later release (before graduation) and stops releases with any serious problem from going ahead (with an IPMC member or two voting -1). A -1 vote is not a veto but usually stops a release, there's only been a handful of successful releases with -1 votes. I don't think anyone thinks that allowing releases with minor issues to be released is an incorrect approach even if it doesn't follow policy, even some people have expressed an opinion that release must be "legal". Serious problems, such as; including GPL licensed software, including compiled code, or copyright violations, in a release is currently seen as a reason to vote -1 on a release. Podlings do try and do the right thing but about 1 in 5 podling releases has a problem like this. Historically there's been a few releases with serious issues that have got legal and incubator VP approval to make a release on a one by one, once-off basis.

There is a clear consensus that the IPMC want to make this somehow easier for podlings.

Current Situation

Recently it has been suggested by several people (including ex-board members) that serious issues should be allowed in an IPMC release, and those issues only need to be fixed before graduation. We've had several long conversations about this and this seems to be unresolvable. I'll note there have been previous conversations along the same lines and going way back, for instance in 2013, that also failed to get consensus.

Some people believe that podling releases do not have to follow policy and most things are allowed until graduation and others say that ASF policy can be ignored but releases must be legal, (which they may or may not realise is stricter than what happens currently). This is damaging to the incubator and causing confusion among podlings. This issue also seems to attract commentary by non-IPMC members whose opinions, not in line with policy, are spread outside the IPMC list. One podling is likely to leave the incubator, in part, because of this situation.

Actions

Initially, this report contained a proposal, to try and improve this situation, however, as it didn't have the consensus of the IPMC, so it has been removed.

From this discussion, The IPMC has identified a few things that can be done, like improving the incubating DISCLAIMER, and coming up with clearer guidance with the the help of the legal committee on exactly what a podling release should follow.

But either of those solves the underlying issue. We have been given contradictory guidance from the board in the past, which is making this situation difficult. If the board could provide some clarity on what extent that podling releases need to follow ASF policy, the IPMC would appreciate it.

Current Challenges

These policies live outside the incubator, and it's put in the difficult position of ensuring that podlings (via their mentors) are following them by the time they graduate. In some cases it seems that these are not been clearly passed on, this may be due to time poor mentors, lack of policy knowledge, overzealous automation, and a poor fit of ASF policies with modern development workflows. Modern development workflows seem to bring up more issues for the IPMC, with it's use of multiple GitHub repositories, wider looser communities (causing possible branding and trademark issues), and expected use of alternative means of distribution. While work has been done, and continues on how these practices fit into the ASF landscape/policy, some podlings are struggling with inconsistent guidance and lack of documentation on this. The IPMC also need to be clear up front that some podlings may need to be major changes to how they do things.

A while back the IPMC identified all of the mentors that were missing, and not performing their duties, and ask them if they wanted to continue, most of them stood down from the IPMC. During this process, it was found there was a strong correlation between not signing off reports and missing mentors. Looking at the reports being signing off there are currently 13 possibly inactive mentors. This number is a massive improvement over the last time where we had over 70 inactive mentors and is a good indicator that mentors are more engaged.

Some of the incubator wiki pages were missed in the transfer and infra are looking into why.

Podlings were reminded that they to put their logos on <http://www.apache.org/logos/> about 1/3 of them responded and added their logo.

A large number of Podlings didn't have their PPMC members or mentors signed up to their private list. A few mentors and podlings have corrected this and reminders to fix this situation has been sent to 20 podlings.

Community

New IPMC members:

- Arpit Agarwal

People who left the IPMC:

- Andy Seaborne
- Mark Thomas
- Stefan Bodewig

New Podlings

None

Podlings that failed to report, expected next month

- BRPC
- DLab
- Iceberg
- Marvin-AI

Graduations

- Dubbo

The board has motions for the following:

- None

Releases

The following releases entered distribution during the month of May:

- [Apache Druid 0.14.1](#)
- [Apache Druid 0.14.2](#)
- [Apache MXNet 1.4.1](#)
- [Apache Omid 1.0.1](#)
- [Apache OpenWhisk Runtimes v1.13.0](#)
- [Apache Ratis 0.3.0](#)
- [Apache Weex 0.24.0](#)
- [Apache Zipkin 2.13.0](#)
- [Apache Zipkin 2.14.0](#)
- [Apache Zipkin API 0.2.1](#)
- [Apache Zipkin Dependencies 2.2.0](#)
- [Apache Zipkin Layout Factory 0.0.5](#)

IP Clearance

None

Legal / Trademarks

None

Infrastructure

None

Miscellaneous

- [Discussion of the use of stack overflow resulted in a legal FAQ item being added](#)
- [Page speed set up their web site](#)
- [Discussion on SGA/CCLA and when they are needed](#)
- [One incubating proposal was turned down as it didn't align with ASF aims](#)
- [New moderators add for general@ list](#)

Credits

Table of Contents

[Crail](#)
[Daffodil](#)
[DataSketches](#)
[Druid](#)
[Hivemall](#)
[IoTDB Marvin-AI](#)
[Milagro](#)
[Myriad](#)
[Nemo](#)
[Omid](#)
[OpenWhisk](#)
[PonyMail](#)
[SAMOA](#)
[ShardingSphere](#)
[SINGA](#)
[Spot](#)
[Superset](#)
[Taverna](#)
[Tephra](#)
[Toree](#)
[Tuweni](#)
[TVM](#)
[Warble](#)
[Zipkin](#)

Crail

Crail is a storage platform for sharing performance critical data in distributed data processing jobs at very high speed.

Crail has been incubating since 2017-11-01.

Three most important unfinished issues to address before graduating:

1. Community building: Increase on number of contributors and users.
2. Further increase project visibility by establishing a convincing use cases of Crail.
3. Make Crail better consumable / cooperate with other projects on Crail integration.

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?

- Apache Crail was presented at 2019 Strata data conference in San Francisco. The presentation is online available.
- We added one new blog posts to the projects Apache home page. It covers deployment options for tiered storage disaggregation.
- A paper on an Apache Crail based, unified ephemeral data storage architecture got conditionally accepted at the USENIX ATC'19 Conference in July.
- The @ApacheCrail twitter account has now 164 followers. We use it to regularly tweet the availability of new blog posts, code releases etc.

How has the project developed since the last report?

- New CrailConfiguration API to allow setting Crail configuration dynamically instead of loading it from a file.
- NVMe tier: It is now possible to set Client host NQN to support restricted access on target side. If host NQN is not specified, a random UUID is still generated.
- Started design of fault tolerance/robustness extensions to Crail namenode implementation.
- Started integration of Crail as an ephemeral data store for serverless frameworks. Publicly discussed integration options at a presentation on 2019 OpenFabrics workshop.
- Continued native C++ implementation of Crail tailored to Machine Learning framework integration.

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:

Nov-26-2018

When were the last committers or PPMC members elected?

December 4th, 2018

Have your mentors been helpful?

Our mentors are very helpful.

Signed-off-by:

- (crail) Felix Cheung Comments:
- (crail) Julian Hyde Comments: Dev list is pretty quiet; a few commits per month, not much discussion. So, activity level is a concern. However, the project members seem to be doing all they can as regards outreach (conference, twitter).
- (crail) Luciano Resende Comments:

IPMC/Shepherd notes:

Daffodil

Apache Daffodil is an implementation of the Data Format Description Language (DFDL) used to convert between fixed format data and XML/JSON.

Daffodil has been incubating since 2017-08-27.

Three most important unfinished issues to address before graduating:

1. Increase community growth and participation (high priority)
2. Establish a frequent release schedule
3. Work with other Apache projects where Daffodil could provide extra functionality

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?

- Continued community support of user developing Daffodil/DFDL training materials
- Continued community support of Apache PLC4X to develop DFDL schemas
- Presented talk at Apache Roadshow DC 2019
- Presented talk at Cross Domain Technical Forum 2019
- Submitted a talk for ApacheCon North America 2019
- Continued contributions from 1 new contributor
- Lost a mentor but got a new one
- Increased discussions/proposals with community for new Daffodil features (e.g. lookahead, enumerations, character encoding)

How has the project developed since the last report?

- 29 commits merged from 4 different contributors
- 55 issues created, 52 issues resolved
- Greatly improved interoperability with IBM DFDL implementation, primarily related to separator behavior
- Merged in new enumeration and TypeValueCalc feature
- Many miscellaneous bug fixes
- The next release is likely within the next month or two

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:

2019-02-26

When were the last committers or PPMC members elected?

None, same as project incubation

Have your mentors been helpful?

Yes, no issues.

Signed-off-by:

- (daffodil) Dave Fisher Comments: Podling is trying to grow the community.
- (daffodil) Christofer Dutz Comments: Great seeing new people being voted in. And wasn't Brandon Sloane voted in to the PPMC some time around 2019-05-31?

IPMC/Shepherd notes:

DataSketches

DataSketches is an open source, high-performance library of stochastic streaming algorithms commonly called "sketches" in the data sciences. Sketches are small, stateful programs that process massive data as a stream and can provide approximate answers, with mathematical guarantees, to computationally difficult queries orders-of-magnitude faster than traditional, exact methods.

DataSketches has been incubating since 2019-03-30.

Three most important unfinished issues to address before graduating:

1. Finish code migration
2. Set up automated builds
3. Establish code review practices

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

aware of

No

How has the community developed since the last report?

We are still in the process of setting up permissions and figuring out Apache environment.

How has the project developed since the last report?

Most DataSketches repos have been moved to Apache repos.

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:

No releases yet

When were the last committers or PPMC members elected?

We have just signed up our initial committers

Have your mentors been helpful?

Yes, very helpful.

Signed-off-by:

- (datasketches) Liang Chen Comments:
- (datasketches) Kenneth Knowles Comments:
- (datasketches) Furkan Kamaci Comments:

IPMC/Shepherd notes:

Druid

Druid is a high-performance, column-oriented, distributed data store.

Druid has been incubating since 2018-02-28.

Three most important unfinished issues to address before graduating:

1. Move the website to Apache infrastructure.
2. Fix potential errata in the website reported by Apache Podling Website Checks.

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 healthy, constant flow of bug fixes, quality improvements and new features are still ongoing at <https://github.com/apache/incubator-druid>.

How has the project developed since the last report?

- Since the last report there have been 493 commits from 86 contributors.
- We have released 4 versions, i.e., 0.13.0, 0.14.0, 0.14.1, and 0.14.2. 0.13.0 was the first Apache release.
- We have code frozen 0.15.0 branch and are preparing for the first release candidate.
- The web page migration from druid.io to druid.apache.org is still in progress.

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:

Druid 0.14.2-incubating on 2019-05-29

When were the last committers or PPMC members elected?

The Druid PPMC elected 10 committers to the project on November 20, 2018.

Have your mentors been helpful?

They have been very helpful.

Signed-off-by:

- (druid) Julian Hyde Comments: In my opinion Druid is overdue to graduate. There seems to be off-list consensus that graduation is near but I would like to see more conversation on-list. I don't agree that mentors "have been very helpful"; for the record, the other two mentors have been absent, including from tight release votes. I suspect that they are fatigued.
- (druid) P. Taylor Goetz Comments:
- (druid) Jun Rao Comments:

IPMC/Shepherd notes:

Justin Mclean: Elected 10 committers at once? I can see a more recent PPMC addition in March.

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 unfinished issues to address before graduating:

1. Community growth (committers and users)
2. One or more Apache Releases as an Incubator project
3. Documentation improvements

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?

- Github watchers/stars are gradually increasing: 219 stars as of June 3 (was 199 on Mar 1)
- Twitter account @ApacheHivemall followers are gradually increasing: 196 followers as of June 3 (was 189 on Mar 1)
- PPMC is invited for a private tech talk about Hivemall at Indeed

How has the project developed since the last report?

- Preparing the third ASF release, v0.6.0, by the end of June

Since the last report, we have

- In the last 3 months, we opened 14 JIRA issues and closed 15 JIRA issues as seen in <https://goo.gl/QFQEF5> (as of June 3) Created Resolved
Mar 2019 2 7 Apr 2019 9 8 May 2019 3 0
- Created 8 Pull Requests and closed 8 Pull Requests between Mar 1th and May 31th.

<https://github.com/apache/incubator-hivemall/pulls?utf8=%E2%9C%93&q=is%3Apr+created%3A2019-03-01..2019-05-31>

<https://github.com/apache/incubator-hivemall/pulls?utf8=%E2%9C%93&q=is%3Apr+closed%3A2019-03-01..2019-05-31>

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:

2018-12-03

When were the last committers or PPMC members elected?

- Elected Jerome Banks as a committer on April 2, 2018.

Have your mentors been helpful?

- Koji is active at mentoring.

Signed-off-by:

- (hivemall) Xiangrui Meng Comments:
- (hivemall) Daniel Dai Comments:
- (hivemall) Koji Sekiguchi Comments:

IPMC/Shepherd notes:

IoTDB

IoTDB is a data store for managing large amounts of time series data such as timestamped data from IoT sensors in industrial applications.

IoTDB has been incubating since 2018-11-18.

Three most important unfinished issues to address before graduating:

1. Learn to how to release a formal version;
2. Declare a clear release roadmap;
3. Attract more contributors joining the community.

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

aware of?

No.

How has the community developed since the last report?

We can find that IoTDB is becoming a new choice when users need a time series database. New users from UCI (IoTDB-104), Sangfor Inc., University of Liverpool (IoTDB-94) etc. are testing and proposing issues for IoTDB.

How has the project developed since the last report?

1. IoTDB is more stable than before. Now it works well on a VM provided by Apache (iotdb-vm.apache.org).
2. A distributed version is almost done (in the 'cluster branch').
3. In the last 3 months, we opened 78 JIRA issues and resolved/closed 40 JIRA issues. Created Resolved Mar 2019 42 26 Apr 2019 15 6 May 2019 21 8
4. In the last 3 months, we merged/closed 107 PRs in Github (seen at Jun

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

Have your mentors been helpful?

Yes

Signed-off-by:

- (iotdb) Justin Mclean Comments: Needs to be more on-list communication. There report mentions new users, where is this communication taking place?
- (iotdb) Christofer Dutz Comments: I agree that there is almost no communication on the list. I have mentioned this quite a number of times before, communication tends to increase for a few days but drops to almost 0 a few days later. Also I would have liked to read about how the COMMUNITY has developed, not the project.
- (iotdb) Willem Ning Jiang Comments: We could encourage the developer to share their ideas through the mailing list.
- (iotdb) Kevin A. McGrail Comments: Would like to see more discussion on the lists and people suggesting bug fixes being asked to join the community. Like the list of the issues preventing release.

IPMC/Shepherd notes:

Marvin-AI

Marvin-AI is an open-source artificial intelligence (AI) platform that helps data scientists, prototype and productionalize complex solutions with a scalable, low-latency, language-agnostic, and standardized architecture while simplifies the process of exploration and modeling.

Marvin-AI has been incubating since 2018-08-21.

Three most important unfinished issues to address before graduating:

1. Move development infrastructure to Apache
2. Build Community
3. Build 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?

- Brought in a new PPMC member
- Participation in the local PyData event with a presentation of the project by Zhang Yifei.

How has the project developed since the last report?

- Community is engaged on the development for R support.

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?

07-January-2019 - Wei Chen

Have your mentors been helpful?

- The mentors are responsive when the community needs it.

Signed-off-by:

- (marvin-ai) Luciano Resende Comments: William Colen Comments:
-

Milagro

Distributed Cryptography; M-Pin protocol for Identity and Trust

Milagro has been incubating since 2015-12-21.

Three most important unfinished issues to address before graduating:

1. Growing the community of contributors and getting additional/new mentors
2. Getting the first release out and establishing a frequent release schedule
3. Work with other Apache projects where Milagro could provide extra functionality

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

aware of?

Milagro needs additional mentors.

How has the community developed since the last report?

We have brought on 3 new contributors in the last month.

How has the project developed since the last report?

1. Updates and bug fixes to the core Milagro crypto library.
2. 3 new committers - Giorgio Zoppi, Chris Morris, and Alessandro Budroni and 2 contributors - John McCane Whitney and Howard Kitto who we hope will become contributors. There is an active development stream now occurring.
3. 1 new PPMC - Giorgio Zoppi
4. Setup a Milagro Confluence site where all contributors and PPMC members can create content at will.
5. JIRA is now setup for committers and PPMC members so that we can collaborate at will and efficiently. Also setup the Kanban board and cleaned out the dead/redundant/duplicate issues.

<https://issues.apache.org/jira/projects/MILAGRO/issues/MILAGRO-18?filter=all> openissues

To Do:

1. File Podling report for June, due in a few days (done).
2. Move existing website on SVN to new githubpubsub 'asf-site' branch publishing system. Will file JIRA request with Infra.
3. Travis - CI for milagro-crypto-c and milagro-javascript repos. Will file JIRA request with Infra.

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:

Milagro is immature but now has some momentum. It is critical the momentum is continued to get the first release out.

Date of last release:

XXXX-XX-XX

When were the last committers or PPMC members elected?

04-June-2019

Have your mentors been helpful?

Requests for help in establishing new committer accounts or just general advice for how to get things done could be delivered in a more timely matter (or if at all). Big thank you to Justin and Craig for the guidance this month.

Signed-off-by:

- (milagro) Nick Kew Comments: This podling has long been moribund, but a sudden flurry of activity suggests potential revival. I am just now (hopefully) catching up with the renewed activity.

IPMC/Shepherd notes:

Justin Mclean: Please ask on the incubator genral@ list for new mentors.

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 unfinished issues to address before graduating:

1. MYRIAD-279: Reworking a new UI design
2. MYRIAD-280: Improve overall stability
3. Rethink a new roadmap for empowering the project

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

aware of?

We need new mentors. The PPMC number is a problem too, we are only two PPMCs/committers, so we have problems with the minimum number of votes for consensus.

How has the community developed since the last report?

We don't have activity in our community for this period.

How has the project developed since the last report?

After a few years of inactivity the project is slowly getting activity, at least from the PPMC point of view. The initial PPMCs are not active anymore, nevertheless the current PPMCs (only two right now) are able of continuing the project from a formal point of view, we are able of creating new releases and to put under control the ASF resources available for Myriad, so the current PPMCs has the knowledge of creating new releases from end-to-end.

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: Reorganize the main goals of the project for getting new comers.

Date of last release:

Release 0.3.0 on March, 25, 2019.

When were the last committers or PPMC members elected?

2018-05-22 Committer/PMC member Juan P. Gilaberte. 2018-03-28 Committer/PMC member Javi Roman.

Have your mentors been helpful?

No answer.

Signed-off-by:

- (myriad) Benjamin Hindman Comments:
- (myriad) Ted Dunning Comments:

IPMC/Shepherd notes:

Ted Dunning: I am much too busy lately to be an effective mentor. This increases the urgency to find new helpers, especially during a project reboot.

Justin Mclean: Please ask on the incubator genral@ list for new mentors. How can you make release with only 2 PPMC members? Are there any committers or other people you can vote in as PPMC members?

Nemo

Nemo is a data processing system to flexibly control the runtime behaviors of a job to adapt to varying deployment characteristics.

Nemo has been incubating since 2018-02-04.

Three most important unfinished issues to address before graduating:

1. Grow the community (committers, contributors, users)
2. Add and improve features
3. Create more 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?

- Committers actively involved in sending PRs and doing code reviews
- GSoC students actively contributing to Nemo
- Engaging with developers in open source venues in June
Beam Summit @ Berlin, KRnet in Korea

How has the project developed since the last report?

- Improved and stablized Nemo pass interfaces
- Developed more formal treatment of expressing optimization policies in Nemo
- Added a Machine-Learning-based policy (e.g., xgboost)
- Added offheap memory store and improved shuffle
- Added a lambda pass to enable running Nemo jobs in AWS Lambda
- A paper describing Nemo got accepted at USENIX ATC 2019, which will be presented in July.

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:

Release 0.1 on Dec. 31, 2018

When were the last committers or PPMC members elected?

None yet.

Have your mentors been helpful?

Mentors have been responsive on answering the questions the community have had. They guided the community well on how to engage with developers in the open source community.

Signed-off-by:

- (nemo) Davor Bonaci Comments: The project is on track. I'm starting to be worried about the decreased activity lately.
- (nemo) Hyunsik Choi Comments:
- (nemo) Byung-Gon Chun Comments:
- (nemo) Jean-Baptiste Onofre Comments:
- (nemo) Markus Weimer Comments:

IPMC/Shepherd notes:

Omid

Omid is a flexible, reliable, high performant and scalable ACID transactional framework that allows client applications to execute transactions on top of MVCC key/value-based NoSQL datastores (currently Apache HBase) providing Snapshot Isolation guarantees on the accessed data.

Omid has been incubating since 2016-03-28.

Three most important unfinished issues to address before graduating:

1. Document additional features added for the integration with Apache Phoenix, as well as, the low latency algorithm.
2. Get positive feedback from other projects currently in Apache.
3. Do podling name search.

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?

1. Lars Hofhansl returned as a mentor.
2. A new release 1.0.1 was done that contains optimizations and bug fixes for the integration with Apache Phoenix.

How has the project developed since the last report?

Release 1.0.1 was published and includes optimization to Omid that inferred from the integration with Phoenix. These features mostly improve the performance of Omid, however, some are bug fixes that revealed by the integration with Phoenix while pushing Omid to its corner cases.

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:

2019-05-21

When were the last committers or PPMC members elected?

2018-09-25

Have your mentors been helpful?

Our mentors were amazingly helpful, they help us finding performance and correctness bugs, as well as helped in the release process by evaluating the release and vote internally and externally.

Signed-off-by:

- (omid) Alan Gates Comments:
- (omid) James Taylor Comments:
- (omid) Lars Hofhansl Comments:

IPMC/Shepherd notes:

OpenWhisk

Distributed Serverless computing platform.

OpenWhisk has been incubating since 2016-11-23.

Three most important unfinished issues to address before graduating:

1. None. The PPMC has engaged the community in initiating the graduation 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?

1. Strong dev-list activity with good mix of technical discussion, project direction, highlighting of significant PRs, and release coordination. For last 3 months, we are averaging over >150 messages per month sent by >28 individuals.
2. incubator-openwhisk Github stars: 4,032 (+207 since last report)
3. incubator-openwhisk GitHub forks: 779 (+41 since last report)
4. Slack community 1,243 members (+84 from last report) and >12,000 messages since the last report. Actively used by users for support.
5. Received 6 new ICLAs since the last report.

How has the project developed since the last report?

1. Have now released multiple versions of virtually all OpenWhisk software components. Nearing completion of project's first "unified release" wave.
2. Active experimentation with KNative and OpenWhisk combinations to understand how best to leverage both technologies to support Serverless computing for our user community.
3. Formal legal transfer of the OpenWhisk trademark from IBM to ASF has been completed.

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:

2019-05-13 (released second versions of 9 OpenWhisk runtime components)

When were the last committers or PPMC members elected?

Dominic Kim was elected as committer and PPMC member on 4/12/2019

Have your mentors been helpful?

Yes. Mentors have been helpful.

Signed-off-by:

- (openwhisk) Bertrand Delacretaz Comments: Expecting graduation soon indeed!
- (openwhisk) Matt Sicker Comments:
- (openwhisk) Krzysztof Sobkowiak Comments:

IPMC/Shepherd notes:

Justin Mclean: Just a minor thing. I see a number of conversations happening on the private list that probably should be occurring on the public list.

PonyMail

Pony Mail is a mail-archiving, archive viewing, and interaction service, that can be integrated with many email platforms.

Pony Mail has been incubating since 2016-05-27.

Three most important unfinished issues to address before graduating:

1. Grow the community
2. Lower the bar to participation
3. Get more releases out

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

aware of?

No answer.

How has the community developed since the last report?

- A new release was be made on 20-04-2019

How has the project developed since the last report?

- Security fix done and CVE issued
- New release made

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:

2019-04-20

When were the last committers or PPMC members elected?

No answer.

Have your mentors been helpful?

No answer.

Signed-off-by:

- (ponymail) John D. Ament Comments:
- (ponymail) Sharan Foga Comments: Good to see there has been some activity this quarter with a new release out and a security fix in place.

IPMC/Shepherd notes:

SAMOA

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

SAMOA has been incubating since 2014-12-15.

Three most important unfinished issues to address before graduating:

1. Revitalize the project by resuming development
2. Enlarge the user base and contributing 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 list activity:
- @dev: 8 messages

How has the project developed since the last report?

- 1 new PRs created
- Working on code for supporting Apache Heron

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:

2016-09-30

When were the last committers or PPMC members elected?

January 2018

Have your mentors been helpful?

Yes, the mentors have been helpful and responsive.

Signed-off-by:

- (samoa) Alan Gates Comments: Dave Fisher (shepherd for this reporting cycle) recently suggested that this podling consider retirement as there is very low activity and it has been incubator for quite some time. Whether and how to attempt to restart engagement is being discussed.
- (samoa) Ashutosh Chauhan Comments:

IPMC/Shepherd notes:

ShardingSphere

Sharding-Sphere is an ecosystem of transparent distributed database middleware, focusing on data sharding, distributed transaction and database orchestration.

ShardingSphere has been incubating since 2018-11-10.

Three most important unfinished issues to address before graduating:

1. Continuing to grow the community
2. Make several 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?

The community has been growing. Some volunteers take the issues and some contributors want to donate UI to ShardingSphere. We have finished the first apache release on 21 Apr.

How has the project developed since the last report?

There have a new landscape v0.1 release for ShardingSphere, to split core features and technical layers. Technical layers consist of Interface Layer, SQL Layer, Transaction Layer, Storage Layer and Orchestration Layer, each layer can combine lots 3rd party components. We just try to develop ShardingSphere as a platform which can integrate more and more products.

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:

2019-04-21

When were the last committers or PPMC members elected?

2019-01-21

Have your mentors been helpful?

Mentors are responsive and helpful.

Signed-off-by:

- (shardingsphere) Craig L Russell Comments:
- (shardingsphere) Benjamin Hindman Comments:
- (shardingsphere) Willem Ning Jiang Comments: Could you share information about the mailing list traffic and potential committers?
- (shardingsphere) Von Gosling Comments:

IPMC/Shepherd notes:

SINGA

SINGA is a distributed deep learning platform.

SINGA has been incubating since 2015-03-17.

Three most important unfinished issues to address before graduating:

1. Choosing the project PMC chair.
2. Finalize the Graduation Resolution.
3. Discuss and Vote for graduation with mentors and team members.

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?

- SINGA developers participated EU FOSSA Apache Hackathon, and discussed with other Apache projects (e.g., PLC4X) on collaboration and community building. Multiple new contributors from the Hackathon have made pull requests.
- SINGA committers are collaborating with SpamAssassin on a GSoC2019 project.
- Currently, we have 43 contributors, 1678 stars and 416 forks on github.

How has the project developed since the last report?

- Version 2.0 has been released.
- We have fixed some issues for graduation including the logo, Apache Website Navigation Links Policy, etc.
- There are about 90 commits (including Rafiki) since last report.
- There are 165, 125 and 92 emails on dev@ list for March, April and May respectively.

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:

2019-04-20

When were the last committers or PPMC members elected?

2018-12-21

Have your mentors been helpful?

Yes, mentors have been helpful. We will need more suggestions from them soon when we start the graduation process.

Signed-off-by:

- (singa) Alan Gates Comments:
- (singa) Ted Dunning Comments:
- (singa) Thejas Nair Comments:

IPMC/Shepherd notes:

Spot

Apache Spot is a platform for network telemetry built on an open data model and Apache Hadoop.

Spot has been incubating since 2016-09-23.

Three most important unfinished issues to address before graduating:

1. Revive community activity (Discussion in mailing lists, increase frequency of commits)
2. Make it easier for Devs to contribute to the project (e.g. documentation, framework).
3. Developing a workflow in Spot that allows for intuitive analytics.

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?

Unfortunately there has not been much more community activity. However, we are need of more committers as there are only a couple of active committers at this time. This creates a bottleneck which makes it hard for the community to grow in interest.

How has the project developed since the last report?

There's been discussions on the dev list to merge the existing SPOT-181 ODM branch into the master branch so that its easier to use the Open Data Model within the existing architecture. This should avoid a lot of confusion among new community members in the future as most of the recent work as been done in the SPOT-181 branch. A series of PRs have also been evaluated and we're waiting for more votes before starting the merge process on those PR, and also the PR for moving SPOT-181 into master.

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:

2017-09-08

When were the last committers or PPMC members elected?

2018-01-18

Have your mentors been helpful?

Yes our single mentor is responsive, however we only have one and may need 1 more to help.

Signed-off-by:

- (spot) Uma Maheswara Rao G Comments: I don't see much progress in community building. Very low activities on mailing lists. Also this project needs additional mentors.

IPMC/Shepherd notes:

Justin Mclean: Please ask for more mentors the incubators general@ list.

Superset

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

Superset has been incubating since 2017-05-21.

Most important issues to address in the move towards graduation:

Plan and execute our first Apache release. While pushing three release candidates this far, we ironed out many issues around licensing, copyright, and the existence of binary files in the repository / source code releases. Address remaining project operations issues. See below.

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

aware of? For project operations, we have made the following progress and need the following support: Trademark: No progress has been made here. No one volunteered to get this done. Some attempts have been made but the documentation and deliverable are confusing / unclear.

How has the community developed since the last report?

- Organic growth of our Github contributors (343->368), forks (4300->4633), watchers (1170->1166) and stars (23,326->24,442)

How has the project developed since the last report?

- Added Ville Brofeldt as a new committer
- Promoted Beto de Almeida from committer to PMC member
- Promoted Krist Wongsuphasawat from committer to PMC member on 2019-04-23

- In general, there has been a healthy, constant flow of bug fixes, quality improvements and new features. Please see project's commit log on Github for more details. Some highlights:

Release-related:

3 RCs were brought up for vote (2 failed, one vote in-progress), issues were raised and addressed, the release process and release validation process is becoming more automated through the use of Docker, which allows for a reproducible reference implementation Moved example-related data to live outside of the repository and be downloaded upon running "superset load_examples" Got rid of an image with an unidentified license Removed .mo files (translation-related binaries)

New features / Progress

- feat: Scheduling queries from SQL Lab (#7416) (#7446)
- feat: Live query validation in the SQL Lab UI
- Progress on visualizations as plugins and embeddable components
- propagate color mapping from dashboard to charts
- view presto row objects in data grid

Bug fixes, code quality and housekeeping

- Improved stability / usability of SQL Lab
- Improved stability / usability of geospatial visualizations
- More repo automation with issue stale bots
- Fix control validation handling

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

No official release yet since being voted into Apache Incubation. (Still planning for the first Apache release in Q2, 2019)

When were the last committers or PPMC members elected?

- Ville Brofeldt - Committer (2019-05-07)
- Beto de Almeida - PMC member (2019-05-20)
- Krist Wongsuphasawat PMC member (2019-04-23)

Signed-off-by:

- (superset) Ashutosh Chauhan Comments:
- (superset) Jakob Homan Comments: Just starting to mentor; looking at various growth and ASF Way areas to improve.
- (superset) Alan Gates Comments:

Taverna

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

Taverna has been incubating since 2014-10-20.

Three most important unfinished issues to address before graduating:

1. Re-engage PPMC members
2. Release IP-fixed repositories
3. Graduate!

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

aware of?

Whimsy detected some PPMC members where not signed up to private@taverna list, as well as false negatives on emiratus mentors and alternative addresses. Missing members have been requested added.

(our private list is quiet except for PPMC voting)

How has the community developed since the last report?

Active in March/April, quiet in May.

How has the project developed since the last report?

Release candidate of Taverna Mobile was in progress based on GitHub release mechanism, but this became challenging as its Android source code tend to include a Gradle binary and don't have good release-helpers like in Maven.

This meant the Android release had to be delayed to be done "manually".

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:

2018-01-18

When were the last committers or PPMC members elected?

2018-02-26 (PPMC)

Have your mentors been helpful?

Mentor was not too attentive in April - sorry.

Signed-off-by:

- (taverna) Stian Soiland-Reyes Comments: The community is slow, but figuring out things together

IPMC/Shepherd notes:

Tephra

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

Tephra has been incubating since 2016-03-07.

Most important unfinished issues to address before graduating:

1. Improve community engagement
2. Increase adoption

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

aware of?

- None at this time.

How has the community developed since the last report?

- 4 new JIRAs filed since the last report

How has the project developed since the last report?

- Working on 0.16.0-incubating release
- Working on HBase 1.5 and 2.1 support

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:

2018-09-04

When were the last committers or PPMC members elected?

- None since coming to incubation

Have your mentors been helpful?

- Mentors are helpful and responsive.
- We do only have 2 mentors.

Signed-off-by:

- (tephra) Alan Gates Comments:
- (tephra) James Taylor Comments:

IPMC/Shepherd notes:

Toree

Toree provides applications with a mechanism to interactively and remotely access Apache Spark.

Toree has been incubating since 2015-12-02.

Three most important unfinished issues to address before graduating:

1. Active community
2. Increase active contributors

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?

After the release of Toree 0.3.0 the community slowed down, but we are starting to see asks for supporting Spark 2.4.x which requires releasing Scala 2.12 artifacts and this should spark some community momentum.

How has the project developed since the last report?

Very little activity on the last three/four months after the Toree 0.3.0 release

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:

2018-11-13

When were the last committers or PPMC members elected?

Ryan Blue was added to the PMC on 2017-04-03

Have your mentors been helpful?

There was nothing requiring mentor intervention on the last quarter.

Signed-off-by:

- (toree) Luciano Resende Comments:
- (toree) Julien Le Dem Comments:
- (toree) Ryan Blue Comments:

IPMC/Shepherd notes:

Tuweni

Tuweni is a set of libraries and other tools to aid development of blockchain and other decentralized software in Java and other JVM languages.

Tuweni has been incubating since 2019-03-25.

Three most important unfinished issues to address before graduating:

1. Make a first ASF release.
2. Build a community of developers and implement project governance.
3. Build a community of users around the project.

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

aware of? Currently there are no issues that require the involvement of either the Incubator PMC or the ASF board.

How has the community developed since the last report?

This is the second month of the project incubation. The community is still low, even though a couple patches were given by external contributors.

4 initial committers have not joined the PPMC yet by subscribing to the private list.

How has the project developed since the last report?

The project has attempted to make a release. The second release candidate went to a vote with the IPMC and was rejected. The committers are now working with the IPMC to make sure the NOTICE and LICENSE files are correct.

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:

No release yet.

When were the last committers or PPMC members elected?

No vote has taken place yet.

Have your mentors been helpful?

No issue to report.

Signed-off-by:

- (tuweni) Jim Jagielski Comments:
- (tuweni) Dave Fisher Comments:
- (tuweni) Kenneth Knowles Comments:
- (tuweni) Jean-Baptiste Onofré Comments:
- (tuweni) Michael Wall Comments:
- (tuweni) Furkan Kamaci Comments:

IPMC/Shepherd notes:

TVM

TVM is a full stack open deep learning compiler stack for CPUs, GPUs, and specialized accelerators. It aims to close the gap between the productivity-focused deep learning frameworks, and the performance- or efficiency-oriented hardware backends.

TVM has been incubating since 2019-03-06.

Three most important unfinished issues to address before graduating:

1. Source code and website migration to ASF infra
2. Make the first Apache release
3. Continue to grow the community

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

aware of?

No answer.

How has the community developed since the last report?

TVM community has welcomed one PPMC member in the past month. There is also ongoing votes about adding new committers.

How has the project developed since the last report?

Summary statistics: in the past month, 44 contributors have pushed 98 commits to master and 98 commits to all branches. On master, 366 files have changed and there have been 21,156 additions and 2,810 deletions. The contributions covers areas including documentations, bugfixes, user interface and backend hardware support.

Some highlights of recent developments:

- More robust frontend support to support various machine learning models.
- Cycle accurate simulation to make it easy to add new architecture backends.
- Quantized models to enable deployment to embedded devices.

For detailed information about the project development, please refer to the monthly TVM community's monthly summary:

- April: <https://discuss.tvm.ai/t/tvm-monthly-april-2019/2426>
- May: <https://discuss.tvm.ai/t/tvm-monthly-may-2019/2793>

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:

no release yet

When were the last committers or PPMC members elected?

May 23rd

Have your mentors been helpful?

Mentors are very helpful in helping to set up the podling and provide helpful guidance.

Signed-off-by:

- (tvm) Markus Weimer Comments:
- (tvm) Byung-Gon Chun Comments:
- (tvm) Sebastian Schelter Comments:
- (tvm) Henry Saputra Comments:
- (tvm) Timothy Chen Comments:
- (tvm) Furkan Kamaci Comments:

IPMC/Shepherd notes:

Dave Fisher: This already busy project and community is just getting started in the Incubator. They look to be on the right track. They have some resources which may require a VM. I suggested that they reach out to Infra on Slack for some answers. It's good to see mentor engagement.

Warble

a distributed endpoint monitoring solution where the agent is hosted on your own hardware.

Warble has been incubating since 2018-06-11.

Three most important unfinished issues to address before graduating:

1. Build community
2. Expand documentation
3. Develop a basic test harness

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

aware of?

Warble is hibernating pending further work on developing the community (perhaps via ACNA19)

How has the community developed since the last report?

N/A

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

Have your mentors been helpful?

No answer.

Signed-off-by:

- (warble) Daniel Takamori Comments:

- (warble) Chris Lambertus Comments:

IPMC/Shepherd notes:

We still want the code within the ASF, but have yet to develop critical mass necessary for community development.

Zipkin

Zipkin is a distributed tracing system. It helps gather timing data needed to troubleshoot latency problems in microservice architectures.

Zipkin has been incubating since 2018-08-29.

Three most important unfinished issues to address before graduating:

1. Release all migrated components under Apache Incubator umbrella
2. Migrate the inter-component dependencies to the ones released under Apache Incubator
3. Finalize and polish all repositories to have no outstanding (minor) issues with respect to release policies and processes

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

aware of?

NONE

How has the community developed since the last report?

There have been 2 community events which has happend in April and May respectively:

- UX workshop and Lens GA planning at LINE Fukuoka
- Zipkin Pow-wow

Both events have been quite successful and contributed new interesting ideas regarding the project's current and future developments.

How has the project developed since the last report?

All designated repositories have been migrated under Apache organization.

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:

2019-06-02

When were the last committers or PPMC members elected?

2019-03-15

Have your mentors been helpful?

No answer.

Signed-off-by:

- (zipkin) Michael Semb Wever Comments:
- (zipkin) John D. Ament Comments:
- (zipkin) Willem Ning Jiang Comments:
- (zipkin) Andriy Redko Comments: Very good pace with releases, really glad to see things moving.
- (zipkin) Sheng Wu Comments: Zipkin has released most of components in incubator, by different release managers. New PPMC members are selected. Near graduation.

IPMC/Shepherd notes:

Dave Fisher: Podling suitable name search is needed. Also a download page on the website.