

September2019

FINAL DO NOT EDIT - REPORT SUBMITTED

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

It 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] (X and 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 September 04	Podling reports due by end of day
Sun September 08	Shepherd reviews due by end of day
Sun September 08	Summary due by end of day
Tue September 10	Mentor signoff due by end of day
Wed September 11	Report submitted to Board
Wed September 18	Board meeting

Shepherd Assignments

Dave Fisher	SAMOA
Dave Fisher	Warble
Drew Farris	Pony Mail
Drew Farris	Tephra
Justin Mclean	BRPC
P. Taylor Goetz	Daffodil
P. Taylor Goetz	SDAP
P. Taylor Goetz	Spot
Timothy Chen	Crail
Timothy Chen	DLab
Timothy Chen	IoTDB
	Amaterasu
	Druid
	Hivemall
	Iceberg
	Marvin-AI
	Myriad

	Nemo
	Omid
	SINGA
	ShardingSphere
	Superset
	Tamaya
	Taverna
	Toree

Incubator PMC report for September 2019

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

This monthly report is in markdown so that it's easier to read. If you are not viewing this in that format, it can be seen here: <https://cwiki.apache.org/confluence/display/INCUBATOR/September2019>

There are presently 48 podlings incubating. In August, podlings executed 9 distinct releases. We added three new IPMC members and no IPMC members retired.

We have two new podlings this month DolphinScheduler and MesaTEE. There have been a couple of ongoing enquires about new projects, and some proposals are being worked on. A couple of podlings are heading towards graduation in the next few months. Druid was up for graduation last month but withdrew as there's a trademark issue with its name.

Several podlings Amaterasu, Tamaya and Tephra didn't submit a report and have been asked to do so next month. Tamaya has now failed to report twice in a row.

Amaterasu has failed to report three times in a row. It looks like the PPMC is missing and there's little or no activity on the list. One final chance will be given to the project to explain what its status is and how it hopes to graduate before it is asked to retire.

Omid and ShardingSphere submitted insufficient reports and will be asked to report next month.

A few mentors have been added and removed in various projects.

There were 3 IP clearances.

The incubator web site menu was reorganized to group similar content together, and further work was done on simplifying the language used on the incubator policy page. A few other pages were also improved.

The new disclaimer policy accepted by the board last month has received good feedback, and we've had a podling already make a release with it. The board asked for some updated documentation on the new disclaimer, and that was undertaken.

Some podlings are still struggling to get 3 +1 votes on their release by mentors and other IPMC votes are needed to make releases. However, those votes seem to be happening quite quickly with several IPMC members checking releases.

A new mentor guidance document is in the process of being created, so far not many mentors have contributed to this. Given the many years of experience over many mentors, it would be nice to see mentors contribute to this. This document would help future mentors and podling deal with everyday situations that come up that they may be unsure of how to deal with.

The IPMC consider dropping that ASF member can join just back asking. Looking at mentor it more than that mentors join this way go missing, other than those who are voted in. Some IPMC member didn't think this was a good idea, so the policy stands.

There was a long discussion on the requirement of IPMC vote on podling release which included several suggested changes to policy and new ideas. Consensus was not able to be reached on any changes, and the discussion has stalled. The main issue being that PMC votes are needed on releases and some IPMC members want to do away with IPMC votes on releases altogether. A new project was suggested as an experiment to try this out (exact details still to be worked out), but so far that has not happened. It was even suggested by a couple of IPMC members that a new incubator project be made.

It may be best to wait until we see the impact of the work in progress disclaimer before revisiting the IPMC vote situation.

There will be a Podling's Shark Tank at the upcoming ApacheCON North America and server talks on incubating subjects and projects.

ECharts is showing slow progress and thought they were ready to graduate when they were not. A face to face meeting with several experienced IPMC members has been organized at ApacheCon to get them back on the right track.

There was a discussion on the private on some podlings that probably should be retired.

Community

New IPMC members:

- Gian Merlino
- Zhijie Shen
- Jamie Mark Goodyear

People who left the IPMC:

None

New Podlings

DolphinScheduler MesaTEE

Podlings that failed to report, expected next month

- Amaterasu
- Tamaya
- Tephra

Graduations

- None

The board has motions for the following:

- Apache Rya
- Apache SINGA

Releases

The following releases entered distribution during the month of August:

- Apache Flagon UserALE.js v2.0.2
- DataSketches memory 1.1.0
- Druid 0.15.1
- IoTDB 0.8.0
- Milagro crypto 1.0.0
- Myriad 0.4.0
- ShardingSphere 4.0.0
- Superset 0.34.0
- Tuweni 0.8.1

IP Clearance

- Apache Milagro - Decentralized Trust Authority
- Apache Celix - Websocket pubsub admin
- Apache Beam ZetaSQL Translator

Legal / Trademarks

None.

Infrastructure

None.

Miscellaneous

Nothing of interest.

Table of Contents

[BRPC](#)
[Crail](#)
[Daffodil](#)
[DLab](#)
[Druid](#)
[Hivemall](#)
[Iceberg](#)
[IoTDB](#)
[Marvin-AI](#)
[Myriad](#)
[Nemo](#)
[Omid](#)
[PonyMail](#)
[SAMOA](#)
[SDAP](#)
[ShardingSphere](#)
[SINGA](#)
[Spot](#)
[Superset](#)
[Taverna](#)
[Toree](#)
[Warble](#)

BRPC

brpc is an industrial-grade RPC framework for building reliable and high- performance services.

brpc has been incubating since 2018-11-13.

Three most important unfinished issues to address before graduating:

1. finish Apache release regularly
2. attract more users and contributors
3. build website

Are there any issues that the IPMC or ASF Board need to be aware of?

None.

How has the community developed since the last report?

brpc team had organized a forum on Aug 3rd in Shanghai China, over 100 engineers attended it, Contributors from Baidu/bilibili/iQiyi shared their experience on using brpc in their projects. we hope to attract more engineers to use and contribute to brpc project.

A new team from baidu shanghai office have began to working on brpc, including liuzhiwei, shizhan, tangjie, xingyongxu and yuhao.

How has the project developed since the last report?

we proposed the first apache release, the vote of rc1 was failed due to one 3rd party's license missed from LICENSE. the vote of rc2 was still on, waiting for JB's comments.

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:

Still working on the first release

When were the last committers or PPMC members elected?

No new ones.

Have your mentors been helpful and responsive?

JB promised to check our release candidate 2. Von gave us many suggestions about how to discuss in maillist.

Signed-off-by:

- (brpc) Kevin A. McGrail
Comments: None. Good work on the rc2 work. Thanks to Von and JB for their input.
- (brpc) Jean-Baptiste Onofré
Comments:
- (brpc) Von Gosling
Comments: There is some issues about rc2, begin to prepare rc3. What's more, I'm encouraging them to interact more in the community, and I'm sure their first major release will be back soon after enter the incubator

IPMC/Shepherd notes:

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.

Are there any issues that the IPMC or ASF Board need to be aware of?

none

How has the community developed since the last report?

- We presented an Apache Crail based, unified ephemeral data storage architecture at USENIX ATC'19 in July.
- The USENIX ATC paper, lightening talk, and talk are online.

How has the project developed since the last report?

- Incorporated pull requests to fix several resource management issues.
- Working towards new 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:

2018-11-26

When were the last committers or PPMC members elected?

December 4th, 2018

Have your mentors been helpful and responsive?

Our mentors are very helpful.

Signed-off-by:

- (crail) Julian Hyde
Comments: Activity is slow but steady. A paper & talk at a major conference should help build community. It's also about time for a release; I will make a suggestion on the dev list.
- (crail) Luciano Resende
Comments:
- (crail) Felix Cheung
Comments: Agreed with report. Let's try to get more traffic on dev list also.

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 from outside Tresys (main priority)
2. Work with other Apache projects where Daffodil could provide extra functionality
3. Establish a frequent release schedule

Are there any issues that the IPMC or ASF Board need to be aware of?

- None

How has the community developed since the last report?

- Voted to add two new PPMC members. One accepted and has been added. One declined due to CLA concerns
- Accepted for talk at ApacheCon NA 2019
- Accepted for Podling SharkTank at ApacheCon NA 2019
- Gained a new involved contributor with multiple improvements and bug fixes
- Outside contributor released tutorials related to DFDL and Daffodil

How has the project developed since the last report?

- Released Apache Daffodil (incubating) 2.4.0
- 42 commits merged from 5 different contributors
- 54 issues created, 57 issues resolved--improvements made at keeping up with new bugs
- Greatly improved compatibility with IBM DFDL implementation
- Added new features to continue to meet DFDL specification
- Added new extensions to be proposed as additions to future DFDL versions

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:

Daffodil PPMC members are attendeing ApacheCon NA and are scheduled for a talk and are taking place in the podling shark tank. Plan to search out other opportunities at the conference to grow our community.

Date of last release:

- 2019-07-12

When were the last committers or PPMC members elected?

- 2019-06-20 - Brandon Sloane (PPMC and comitter)

Have your mentors been helpful and responsive?

Yes, mentors have been helpful and responsive.

Signed-off-by:

- (daffodil) Dave Fisher
Comments: Once some more PPMC from other entities are brought onboard the podling will be ready to graduate.
- (daffodil) Christofer Dutz
Comments: I also hope the acna19 participation will help attract new people.

IPMC/Shepherd notes:

DLab

DLab is a platform for creating self-service, exploratory data science environments in the cloud using best-of-breed data science tools.

DLab has been incubating since 2018-08-20.

Three most important unfinished issues to address before graduating:

1. Come up with a new name for DLab
2. Extend the amount of committers. The plan here is to participate in a number of conferences, starting with ApacheConf in Berlin on October 23, EPAM SEC conference in Gansk on SolutionsHub area
3. After release 2.2 is made, the team will have to do major refactoring of the code, updating of product documentation, which might potentially slow down the process of releasing version 2.3

Are there any issues that the IPMC or ASF Board need to be aware of?

No issues.

How has the community developed since the last report?

A new committer is added:

- Adams Disturber (Andrii Dumych) Also awaiting for committer Mykola Bodnar, pending account creation. ICLA files are signed off and submitted, waiting for Apache ID to be created.

How has the project developed since the last report?

Team is actively working on release 2.2 features. There will be many new features and conceptual changes introduced in upcoming release. We are introducing the concept of projects within DLab, which will allow to control permissions for groups of DLab users, improve collaboration capabilities. Major improvements will be made in terms of GCP platform deployment. GCP will support the majority of the features, currently available in DLab for AWS and Azure Clouds. One of the conceptual features implemented is the ability to provision DLab's self-service noe in k8s and being able to integrate with many identity providers, supporting OAuth2 and SAML2.0 via Keycloak.

How would you assess the podling's maturity?

Please feel free to add your own commentary.

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

Date of last release:

2019-15-04

When were the last committers or PPMC members elected?

2019-20-08

Have your mentors been helpful and responsive?

Yes, mentors have been helpful and responsive.

Signed-off-by:

- (dlab) P. Taylor Goetz
Comments:

- (dlab) Henry Saputra
Comments:
- (dlab) Konstantin I Boudnik
Comments:
- (dlab) 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. Resolve trademark issues

Are there any issues that the IPMC or ASF Board need to be aware of?

We will need to resolve the trademark issues before graduation, which the board has been made aware of privately. We are grateful for the board's support, and we will work with trademarks@ to resolve the issues.

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>.
- Two Druid community meetups have been scheduled on 9/17 and 9/19.

How has the project developed since the last report?

- Since the last report, we have had a total of 293 commits from 30 contributors.
- The project website has been migrated to Apache infrastructure.
- We have released 2 versions, 0.15.0 and 0.15.1.
- We have code frozen the upcoming 0.16.0 release branch and are preparing the first release candidate.

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 0.15.1-incubating was released on Aug 15, 2019.

When were the last committers or PPMC members elected?

The Druid PPMC elected 2 new committers to the project on Sep 2, 2019.

Have your mentors been helpful and responsive?

Julian has continued to be very helpful.

Signed-off-by:

- (druid) Julian Hyde
Comments: We got very close to graduation this month; the issues that prevented graduation were frustrating to everyone but people are working hard to resolve them. Well done.
- (druid) P. Taylor Goetz
Comments:
- (druid) Jun Rao
Comments:

IPMC/Shepherd notes:

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

Are there any issues that the IPMC or ASF Board need to be aware of?

None

How has the community developed since the last report?

- Takuya (PPMC) will give a talk at ApacheCon, LAS VEGAS.
- Got and merged a pull request from new external contributor.
- Github watchers/stars are gradually increasing: 234 stars as of Mar 1 (was 199 on Mar 1)
- Twitter account @ApacheHivemall followers are gradually increasing: 208 followers as of Mar 1 (was 189 on Mar 1)

How has the project developed since the last report?

- Working toward the 3rd Apache release. After HIVEMALL-245 is fixed, it will be ready to the vote.

Since the last report, we have

- In the last 3 months, we opened 10 JIRA issues and closed 7 JIRA issues as seen in <https://goo.gl/QFQEF5> (as of Sept 5)
- Created 7 Pull Requests and closed 5 Pull Requests between May 1st and Aug 31th

How would you assess the podling's maturity?

The pace of community growth is slow. No much active developers.

Considering a pivot of the project scope to get more attention:
moving from Hive/Spark to a standalone library focusing on Python/Jupyter notebook integration.

- 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 and responsive?

- Koji is active at mentoring.

Signed-off-by:

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

IPMC/Shepherd notes:

Justin Mclean: The included stats don't really mean much to anyone outside of your project, please drop them from future reports.

Iceberg

Iceberg is a table format for large, slow-moving tabular data.

Iceberg has been incubating since 2018-11-16.

Three most important unfinished issues to address before graduating:

1. Make the first Apache release. (<https://github.com/apache/incubator-iceberg/milestone/1>)
2. Grow the Iceberg community
3. Add more committers and PPMC members

Are there any issues that the IPMC or ASF Board need to be aware of?

No issues.

How has the community developed since the last report?

The community continues to grow steadily. In the last month:

- 59 pull requests have been merged
- 17 people contributed the merged PRs
- 18 issues have been closed, 22 issues were opened

For comparison, the last report had 74 pull requests merged over 3 months.

How has the project developed since the last report?

- License documentation has been completed for the Java project, unblocking the first release
- Added more documentation to iceberg.apache.org
- Started vectorized read branch with significantly better performance
- Added metadata tables
- Added configuration to control statistics and truncate long values
- Improved Hive Metastore integration
- A working python read path has been submitted in PRs

How would you assess the podling's maturity?

- 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?

- Anton Okolnychyi was added 30 August 2019

Have your mentors been helpful and responsive?

Yes

Signed-off-by:

- (iceberg) Ryan Blue
Comments:
- (iceberg) Julien Le Dem
Comments:
- (iceberg) Owen O'Malley
Comments:
The project also gave two presentations: * Berlin Buzzwords (June 2019) * ApacheCon NA (Sep 2019) Iceberg is being used in production at Netflix on huge tables, up to 25 petabytes.

- (iceberg) James Taylor
Comments:
- (iceberg) Carl Steinbach
Comments: Approval added by Ryan Blue, Carl had trouble editing the new report location

IPMC/Shepherd notes:

Justin Mclean: The included stats don't really mean much to anyone outside of your project, please drop them from future reports. The community growth section might as well be blank. I find it surprising that this project thinks that it is near graduation. Please discuss this with your mentors.

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. Make the community (both the contributors and committers/PPMCs) more diverse and learn to discuss things on the mailing list.
2. More committers should know how to release a version.
3. Write more documentation to help new contributors.

Are there any issues that the IPMC or ASF Board need to be aware of?

No.

How has the community developed since the last report?

1. Some new contributors joined the community. E.g., PR#339 comes from a new contributor and it is really an interesting feature.
2. Some contributors (who are not committers yet) are more active than before. E.g., Boris Zhu and Jack Tsai.
3. More users began to report their issues. E.g., issue #171, #184 and #186.

As all the above contributions and contributors are not from the institutions/companies where initial committers are, we think it is a good signal that indicates more guys noticing the community.

How has the project developed since the last report?

1. IoTDB releases its first Apache version: v0.8.0 !
2. IoTDB is more stable and the performance is far better than before. However, the version is incompatible with previous version (which are released before IoTDB is denoted to ASF).
3. In v0.8.0, many issues are resolved (see release notes: <http://iotdb.apache.org/#/Materials/Release%20Notes>)

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:

2019-08-25

When were the last committers or PPMC members elected?

2019-07-15

Have your mentors been helpful and responsive?

They are helpful, especially in the process of releasing the first version.

Signed-off-by:

- (iotdb) Justin Mclean
Comments: Good to see your first release out, but more active involvement by the PPMC is needed in vetting releases. While it improved, still more discussion needs to happen on the mailing list.

- (iotdb) Christofer Dutz
Comments: I really would like to see more mailing list usage, cause the number of commit messages is in strong contrast to the discussions. I as a mentor some time don't have a clue what the project is working on.
- (iotdb) Willem Ning Jiang
Comments:
- (iotdb) Kevin A. McGrail
Comments: Congrats on 0.8.0!!!!

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 issues to address in the move towards graduation:

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

Are there any issues that the IPMC or ASF Board need to be aware of?

- None

How has the community developed since the last report?

- Participation in the local PyData event with a presentation of the project by Rafael Novello.
- Workshop participation in University of São Paulo with a presentation of the project by Zhang Yifei.
- Participation in the B2W Summit event with a presentation of the project by Zhang Yifei.

How has the project developed since the last report?

- New public engine(post-purchase service prediction) in construction by contributors from University of São Carlos.

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?

- Wei Chen added as new PPMC member (2019-06-07)

Have your mentors been helpful and responsive?

- At least two of the mentors are sometimes active.

Signed-off-by:

- (marvin-ai) Luciano Resende
Comments:
Removed Jim as mentor per his resignation
The report mention that there are some collaboration with University of São Carlos, but the dev list seems very quiet since July. Where are the communications and technical discussions happening ?
 - (marvin-ai) William Colen
Comments:
-

Myriad

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

Myriad has been incubating since 2015-03-01.

Three most important unfinished issues to address before graduating:

1. Builds Apache Myriad releases regularly
2. Attract more users and contributors. Community growing.
3. Attract new mentors for helping the project towards the Apache Way.

Are there any issues that the IPMC or ASF Board need to be aware of?

We need new mentors, we have only one active mentor right now. We have to increase the community and we have to increase the number of mentors too in order to help us to address the Apache way with solid foundations.

How has the community developed since the last report?

We have four active committers (PPMCs) right now. A new committer since the last report. This is a small increment from the community growing point of view, but it helps regarding the minimal amount of PPMC's for achieving a release cycle with regularity.

How has the project developed since the last report?

The most important thing is the activity of creating new releases in a regular way. The Apache resources (project page, documentation page, ...) are maintained and enhanced with new content. Taking into account the effort of rebooting the project (more than two years orphaned), the project is slowly addressing the right way, and we consider the project is successfully rebooted right now. This state allow us working towards attract new users and contributors in our humble opinion.

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: New important features focuses in production environments for attracting new users.

Date of last release:

The 0.4.0 release was issued on Aug 25, 2019

When were the last committers or PPMC members elected?

2019-08-02 New committer/PMC member Oscar Fernandez

Have your mentors been helpful and responsive?

We have only two mentors right now. One of the is for ever missing, and the other is ready for help when the project needs. So we have only one active mentor. This is an issue to address with the help of IPMC.

Signed-off-by:

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

IPMC/Shepherd notes:

Justin Mclean: Please ask for another mentor on the incubator general list.

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. Create more releases

Are there any issues that the IPMC or ASF Board need to be aware of?

None.

How has the community developed since the last report?

- Attracted new contributors through the GSoC program
- Two GSoC students finished up their projects successfully.
- Presented Nemo at Berlin Beam Summit 2019.
- Discussion with Google Beam, Shopify, Seznam.cz developers At the Beam Summit Europe 2019, which occurred in Berlin from June 19th to the 20th, there had been many discussions regarding Nemo from engineers working in the industry, as well as with engineers working with Apache Beam. There had been people from various companies, as Shopify and Seznam.cz, facing problems as large-shuffle and data skew, and had shown a lot of interest in using the Apache Nemo runner. Also, in the following ApacheCon summit happening in Vegas in September, the Apache Beam team will demonstrate running their workloads using the Apache Nemo runner.

How has the project developed since the last report?

- Integrate Apache Crail with Nemo
- Off-heap support to improve shuffle speed
 - DirectByteBufferInput/OutputStream for Off-heap SerializedMemoryStore
 - Off-heap SerializedMemoryStore
 - Off-heap memory management (reuse ByteBuffer)
 - Separation of JVM heap region and off-heap memory region
- Enable Nemo single-stage programs to run on a serverless framework, AWS Lambda
 - Refactor Executor to support different types of execution environment
 - Add a Lambda Executor
 - Add a compiler pass for Lambda
- Progress on a ML-based smart compiler

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. Recent two contributors are good committer candidates.

Have your mentors been helpful and responsive?

Yes.

Signed-off-by:

- (nemo) Davor Bonaci
Comments: Lots of good work on community building, awaiting results. In-project activity lower than before.
- (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. Do podling name search.

Are there any issues that the IPMC or ASF Board need to be aware of?

N/A

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: New important features focuses in production environments

Date of last release:

2019-05-21

When were the last committers or PPMC members elected?

2018-09-25

Have your mentors been helpful and responsive?

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.

Comments

Going forward, we need to document the additional features added for the integration with Apache Phoenix, as well as, the low latency algorithm.

Signed-off-by:

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

IPMC/Shepherd notes:

Justin Mclean: This report need more detail and tell us very little about the podlings progress. Pleas include more detail in you next report. I there any reason you have not added any committers/PPMC member for a year?

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. Growing contributor base
2. Getting release processes nailed to a tee
3. Expanding unit tests

Are there any issues that the IPMC or ASF Board need to be aware of?

How has the community developed since the last report?

Not much to report, progress is slow-paced as usual. This isn't to say that the project is dormant or in danger of retirement, just that the pace is to be considered more casual. There is oversight, and important issues are being addressed in a timely manner.

How has the project developed since the last report?

No noteworthy change, we are operating at a pretty stable level.

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?

Sebb, 2017-06-01. Perhaps we should be looking for more candidates :)

Have your mentors been helpful and responsive?

No answer

Signed-off-by:

- (ponymail) John D. Ament
Comments:
- (ponymail) Sharan Foga
Comments: At least one request from a user has not yet been responded to. Think it would be good to work on more engagement to ensure timely responses to user queries.

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

Are there any issues that the IPMC or ASF Board need to be aware of?

None

How has the community developed since the last report?

- Mailing list activity:
- @dev: 12 messages

How has the project developed since the last report?

- Retirement has been suggested due to its very low activity
- To restart engagement, Corey Sterling has been elected as committer in late August. With his help, we are preparing a new release of Apache SAMOA in early September. Corey Sterling has experience on open source software, as an example, he did the last release of MOA, an open source software non-distributed for data streams very related to SAMOA. He works at University of Waikato.

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?

August 2019

Have your mentors been helpful and responsive?

Yes, the mentors have been helpful and responsive.

Signed-off-by:

- (samoa) Alan Gates
Comments: Happy to see a new committer, hopefully this will help kickstart the project.
- (samoa) Ashutosh Chauhan
Comments:

IPMC/Shepherd notes:

Dave Fisher: Let's see if bringing on a new committer revitalizes the podling.

SDAP

SDAP is an integrated data analytic framework for Big Science problems.

SDAP has been incubating since 2017-10-22.

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

1. Make official SDAP (Incubating) Release
2. Improve/create user guide documentation
3. Improve committer participation

Are there any issues that the IPMC or ASF Board need to be aware of?

No.

How has the community developed since the last report?

We are continuing to attract new committers.

How has the project developed since the last report?

SDAP is still working towards its first release. Work towards a repeatable build pipeline has been started.

New functionality has been added to the Distributed Oceanographic Matchup Service (DOMS) and Extensible Data Gateway Environment (EDGE) components that allow for filtering in situ measurements by a standardized quality flag.

How would you assess the podling's maturity?

There are several deployments of SDAP actively being used and interest is high. However, active participation from project members is low. A push is needed to cross the t's and dot the i's in regards to meeting the requirements Apache has for releasing software so that we can claim a 1.0 version of SDAP.

Please feel free to add your own commentary.

- Initial setup
- Working towards first release
- Community building
- Nearing graduation

Date of last release:

XXXX-XX-XX

When were the last committers or PPMC members elected?

Maya Debellis was elected as a committer on 2019-02-08

Have your mentors been helpful and responsive?

Yes, mentors have been helpful and responsive.

Signed-off-by:

- (sdap) Jörn Kottmann Comments:
- (sdap) Trevor Grant Comments:
- (sdap) Lewis John McGibbney Comments:

IPMC/Shepherd notes:

Justin Mclean: Please use the format provided rather than an old format.

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 candidate releases
3. Make a GA release to make user API stable

Are there any issues that the IPMC or ASF Board need to be aware of?

No

How has the community developed since the last report?

The community has been growing. We have more contributors during those 3 months. A UI project has already donate to ShardingSphere. We have finished the second apache release on 24 August. The count of GitHub issues are 1800+, the count GitHub pull requests are 1000+.

How has the project developed since the last report?

We are merging new UI project to dev branch. The new features data mask and distribute transaction have finished.

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:

2019-08-24

When were the last committers or PPMC members elected?

2019-01-21, but there are some active potential committers may prompt soon.

Have your mentors been helpful and responsive?

Mentors are responsive and helpful. Things tend to be on the right way.

Signed-off-by:

- (shardingsphere) Craig L Russell
Comments: The project is reviewing the "committer bar" which may be set too high. There is plenty of activity, mainly focused on Github dialog.
- (shardingsphere) Benjamin Hindman
Comments:
- (shardingsphere) Willem Ning Jiang
Comments: It's good to see that shardingsphere PPMC put lot efforts on building community, and making a good release.
- (shardingsphere) Von Gosling
Comments: well in community activities, especially in Github Issues and PRs

IPMC/Shepherd notes:

Justin Mclean: The stats show don't mean much to people outside your project, please drop them from future reports and include commentary on how you project is growing instead. Given there lots of activity I would expect this report to have a lot more detail. A quick look at that stats, without explanation, suggests to me that this project is in trouble,. You have a large number of contributors, and outstanding pull request, but are not voting committers in. Please discuss with your mentors what the committer bar should be and what it means to be a committer.

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. Fix the code grant issue being discussed in the general@ list.
2. Update the documentation and packages (e.g., docker images and conda packages)

Are there any issues that the IPMC or ASF Board need to be aware of?

N.A

How has the community developed since the last report?

- The community has been discussing about the graduation
- Currently, we have 45 contributors, 1803 stars and 479 forks on github.

How has the project developed since the last report?

- Since the last report, we have been focusing on enhancing the ONNX feature and distributed training. More than 20 ONNX operators are added. Distributed training using NCCL and MPI has been tested.
- 180 new commits are merged since last report.
- There are 57, 142 and 472 emails on dev@ list for June, July and August 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 and responsive?

Yes. The mentors raised very important issues like the code grant in the graduation discussion.

Signed-off-by:

- (singa) Alan Gates
Comments: This podling is ready to graduate as soon as the IP clearance is in place. I'm happy to see them ready to take their place as a TLP.
- (singa) Ted Dunning
Comments:
- (singa) Thejas Nair
Comments:

IPMC/Shepherd notes:

Justin Mclean: It might be better to say that the issue is that a SGA was never submitted.

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. Create a new roadmap and release plan that will improve adoption.
3. Make it easier to contribute to the project (e.g. documentation, framework).

Are there any issues that the IPMC or ASF Board 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 are a series of PRs that have been evaluated for the SPOT-181 epic, and Spot website branch (asf-site) that are in the process of being merged. We have enough votes to merge, and have put out a request to committers to review so they can be auto-merged if possible, otherwise we will just manually merge them via cli.

Once those are merged, we will work on the plan to merge SPOT-181 into master so that the Open Data Model is consolidated into the main branch of our code.

We're also working on re-focusing the project release roadmap. Our first step will be to re-engage with the community to better understand how/where Spot is being used today, what features/enhancements would benefit the community best, and which committers (existing and potential) are interested in contributing to the roadmap and releases moving forward.

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 and responsive?

Yes, but we would welcome more mentors and further guidance on positioning the project for future graduation.

Signed-off-by:

- (spot) Uma Maheswara Rao G
- Comments:

IPMC/Shepherd notes:

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

Superset

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

Superset has been incubating since 2017-05-21.

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

- Build up towards a steady stream of Apache 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: N/A

How has the community developed since the last report?

- Organic growth of our Github contributors (343->403), forks (4300->5076), watchers (1170->)1212 and stars (23,326->26,021)

How has the project developed since the last report?

- First official Apache release is out!!! "0.34.0"
- Started an effort around a community design-focussed interest group with a set of community contributors from different organizations
- New committer Kim Truong, Erik Ritter

New features * Dashboard interactive filters improvements * See commit log / release notes

Bug fixes, code quality and housekeeping * See the commit log, 169 PRs merged since July 1st

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:

2019-08-27

When were the last committers or PPMC members elected?

- Kim Truong (2019-07-30)
- Erik Ritter (2019-08-29)

Signed-off-by:

- (superset) Ashutosh Chauhan Comments:
- (superset) Luke Han Comments:
- (superset) Alan Gates Comments: Congratulations on the first release!

IPMC/Shepherd notes:

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!

Are there any issues that the IPMC or ASF Board need to be aware of?

Considering retiring.

How has the community developed since the last report?

Very quiet through the summer. Podling considering retirement from incubator.

How has the project developed since the last report?

No development since last report.

Release Candidates planned during June has not progressed further to a vote.

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 and responsive?

No answer

Signed-off-by:

- (taverna) Stian Soiland-Reyes
Comments: While the project has previously been active, done several releases, and recruited PPMC and committers, it is now very quiet. For graduation it only needs a couple of maintenance releases of the remaining code-bases (or not include those in graduation). I suggest the project use the next reporting period to discuss retiring vs. graduation.

IPMC/Shepherd notes:

Justin Mclean: I'm a little confused if you want to graduate or retire.

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

Are there any issues that the IPMC or ASF Board need to be aware of?

None

How has the community developed since the last report?

Toree will have a talk at ApacheCon Las Vegas where we will talk about its overall status and current functionality as well as try to attract more users and in particular contributors.

After the release of Toree 0.3.0 the community slowed down, mostly due to Toree being somewhat stable and running on top of Spark, where the public APIs in use are also stable.

The community needs to work on a release that supports Apache Spark 2.4.x and Scala 2.12.

Toree might be one of those low activity projects which have enough active PPMCs (see last ppmc voting thread), and we might start graduation talks based on these facts.

How has the project developed since the last report?

Very little activity on the last few months after the Toree 0.3.0 release, mostly around some compatibility issues with Spark 2.4.x and Scala 2.12.

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?

Kevin Bates was added to the PPMC on 2019-08-14

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:

Justin Mclean: Please use the format provided rather than an old format.

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. Develop documentation
2. Develop community
3. Develop community

Are there any issues that the IPMC or ASF Board need to be aware of?

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:

None.

When were the last committers or PPMC members elected?

Have your mentors been helpful and responsive?

No answer.

Signed-off-by:

- (warble) Daniel Takamori
Comments:
- (warble) Chris Lambertus
Comments: Warble continues to be stagnant. If no progress is made at ACNA we will likely shutter the effort.

IPMC/Shepherd notes:

Dave Fisher: A similar note was made last year. No trouble if this retires and then returns as an Infra tool.