

December 2017

/\ FINAL (!)

Incubator PMC report for December 2017

Timeline

Wed December 06	Podling reports due by end of day
Sun December 10	Shepherd reviews due by end of day
Sun December 10	Summary due by end of day
Tue December 12	Mentor signoff due by end of day
Wed December 13	Report submitted to Board
Wed December 20	Board meeting

Shepherd Assignments

Dave Fisher		Griffin	
Dave Fisher		SAMOA	
Drew Farris		Daffodil	
Drew Farris		Spot	
John Ament		Gearpump	
John Ament		Quickstep	
Justin Mclean		Wave	
P. Taylor Goetz		OpenWhisk	
P. Taylor Goetz		PageSpeed	
P. Taylor Goetz		Superset	
Timothy Chen		Amaterasu	
Timothy Chen		Myriad	
Timothy Chen		Omid	
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="1395e425-6793-473f-ad0a-c411760cd0f5"><ac:plain-text-body><![CDATA[[none]	AriaTosca]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="04b2ebbe-30fd-4829-be21-8379243f8eff"><ac:plain-text-body><![CDATA[[none]	Crail]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="233aab7c-1541-4575-b5af-b2f163025278"><ac:plain-text-body><![CDATA[[none]	HTrace]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="fc416481-533f-44d0-96e8-17fee18070b6"><ac:plain-text-body><![CDATA[[none]	Hivemall]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="d743d84e-14a1-4fa3-8234-fbc0518b0045"><ac:plain-text-body><![CDATA[[none]	Pony Mail]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="8235672d-a359-413e-baf1-8a450a58cb6"><ac:plain-text-body><![CDATA[[none]	Pulsar]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="28b66d34-ca59-4a9d-b492-5627ed201181"><ac:plain-text-body><![CDATA[[none]	SDAP]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="b234ae53-476c-4f6c-8207-ebde8a6f0ace"><ac:plain-text-body><![CDATA[[none]	SINGA]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="7bd453f7-33fa-4d1f-9ec6-06694ac75390"><ac:plain-text-body><![CDATA[[none]	Taverna]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e3a1bfc8-716d-4dae-9fa0-c878e6bbca54"><ac:plain-text-body><![CDATA[[none]	Tephra]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="09b0d8e2-25a8-48d0-ba80-87860e90a587"><ac:plain-text-body><![CDATA[[none]	Trafodion]]></ac:plain-text-body></ac:structured-macro>

Report content

Incubator PMC report for December 2017

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

There are currently 53 podlings incubating. In the month of November, we executed six releases, added one IPMC member and received the resignation of another (which is in flight until their sole podling retires). We added two podlings to our roster, have a few more in the pipeline and have one podling planning to graduate this month.

* Community

New IPMC members:

- Michael Semb Wever

People who left the IPMC:

- Upayavira (not processed in LDAP yet)

* New Podlings

- Crail
- Service Comb

* Podlings that failed to report, expected next month

- Amaterasu - Activity stopped late November
- HTrace - 0 on list activity
- Pony Mail - Low activity, suspect just missed
- Wave - Retiring
- Myriad - Received report, no mentor sign off
- Spot - Received report, no mentor sign off

* Graduations

The board has motions for the following:

- Trafodion

* Releases

The following releases entered distribution during the month of November:

- 2017-11-01 Apache Freemarker 2.3.27
- 2017-11-08 Apache Netbeans HTML4J 1.5
- 2017-11-15 Apache Mnemonic 0.10.0
- 2017-11-15 Apache MXNet 0.12.1
- 2017-11-18 Apache Netbeans HTML4J 1.5.1
- 2017-11-28 Apache Griffin 0.1.6

Table of Contents

AriaTosca
Crail
Daffodil
Gearpump
Griffin
Hivemall
Omid
OpenWhisk
PageSpeed
Pulsar
Quickstep
SAMOA
SDAP

SINGA
Superset
Taverna
Tephra
Trafodion

AriaTosca

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

AriaTosca has been incubating since 2016-08-27.

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

1. Grow the community and enroll new committers.
2. Have frequent Release cycles to be compliant with the Apache Way.
3. Grow the project's contributions from existing contributors

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

The project is ready for its third release as far as the code is concerned. Considerable functionality has been added since the last release in addition to a growth in diversity of contributions.

How has the community developed since the last report?

- * All project communication takes now place on either the mailing list or ASF's Slack, and is open for everybody. We have worked hard to make this the normal operating mode of the community. There have been some hiccups, but in general the majority of discourse has taken place on the list. Slack has been used as a secondary method.
- * Increased chatter on the mailing list including new subscribers who ask questions, make suggestions and are likely to become contributors themselves in the near future.
- * Weekly grooming meetings are now attended by the majority of the community. This is not only a forum to review the weekly work items list, but also as a place for open discussion about issues, questions as well as for the community to get to better know each other.
- * Over the last 3 months we have seen an over %50 increase in contributors and the existing group of committers voted to elect a new Committer to the project.

How has the project developed since the last report?

- ARIA now has a robust parser testing framework (ARIA-1)
- Huge performance improvement (ARIA-1)
- Significant website and documentation improvements for usability and to enhance new user Usability in terms of "getting started" instructions
- Clearwater IMS Example
- Package names fixed/made consistent
- Configuration loading mechanism
- Removed support for Python 2.6
- 42 JIRA issues have been resolved since the last report
- Moved project to gitbox

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:

2017-07-13 0.1.1

When were the last committers or PPMC members elected?

Tom Nadeau was added to the PMC/committers on 2017-11-28

Signed-off-by:

(ariatosca) Suneel Marthi

Comments:

(ariatosca) John D. Ament

Comments: Activity on list is picking up over the slack conversations. A community member is planning to write a report, however it indicates a corporate hierarchy issue within the project that needs to get solved. In addition, the podling needs to figure out how to operate more on list (without grooming meetings) and do things more asynchronously.

(ariatosca) Jakob Homan

Comments:

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 issues to address in the move towards graduation:

1. Finish project setup
2. Community building
3. Create a first release

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

None.

How has the community developed since the last report?

N/A

How has the project developed since the last report?

- * This is the first report.
- * Apache ICLA signed by all contributors from IBM Zurich Research.
- * Software Grant and Corporate Contributor License Agreement signed by IBM ZUrich Research representative.
- * Infrastructure setup underway (missing JIRA due to infrastructure issues)
- * Available Crail code to be transferred soon.

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?

2017-11-01 (entering incubation)

Signed-off-by:

(crail) Julian Hyde

Comments:

Project is still boot-strapping. No activity on dev list yet (except for preparation of this report).

(crail) Luciano Resende

Comments:

[] (craill) Raphael Bircher

Comments:

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 issues to address in the move towards graduation:

1. Obtain remaining SGA from NCSA
2. First release as an Apache Incubator
3. Broadening base of contributors and building the 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?

- No additional committers or community growth, however activity on the dev@daffodil list has increased, with multiple design discussions, code reviews, and project discussions taking place.

How has the project developed since the last report?

- All tickets from the old JIRA system were successfully transferred to ASF JIRA
- IBM has submitted their SGA. We are now only waiting on an SGA from NCSA, who we believe is very close to submission. Once that is in, we believe all necessary SGA's and CLA's will have been submitted and we can begin the process of relicensing the Daffodil codebase to Apache v2 and switching the package namespace to org.apache.daffodil
- Development pace has increased, with multiple issues, commits, and pull requests having gone through the new Daffodil Code Contributor Workflow
- Continuous integration has been enabled via TravisCI GitHub Integration
- Migrated repository to GitBox
- A website has been developed based on the Apache Incubator Website template. Still need to enable automatic building/publishing--manually publishing for now.

How would you assess the podling's maturity?

Please feel free to add your own commentary.

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

Date of last release:

- None

When were the last committers or PPMC members elected?

- None, same as project incubator inception

Signed-off-by:

[X](daffodil) John D. Ament

Comments: Solid on list communication. Waiting on SGAs to be processed to relicense the code. Podling executed a non-ASF release in the interim.

[X](daffodil) David Fisher

Comments:

IPMC/Shepherd notes:

Drew Farris (shepherd): Both mentors active on the mailing list. Project is still getting set up yet showing healthy mailing list activity.

Gearpump

Gearpump is a reactive real-time streaming engine based on the micro-service Actor model.

Gearpump has been incubating since 2016-03-08.

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

1. Have frequent release cycles.
2. Continue to evolve community interest and support.
3. Keep up with Apache Beam's evolving API.

- 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 was a company managed to deploy Gearpump on their secured Hadoop cluster. They reported some bugs occurred in their environment and contributed the fixes back to Gearpump community.

- How has the project developed since the last report?

15 issues created and 9 issues resolved.

An official Gearpump docker image was published.

Continuous improvements for Gearpump Runner in Apache Beam were made.

Date of last release:

2017-07-17

When were the last committers or PMC members elected?

No new committers or PMC members elected yet.

Signed-off-by:

(gearpump) Jean-Baptiste Onofré

Comments:

(gearpump) Andrew Purtell

Comments:

(gearpump) Jarek Jarcec Cecho

Comments:

(gearpump) Reynold Xin

Comments:

(gearpump) Todd Lipcon

Comments:

(gearpump) Xuefu Zhang

Comments:

Griffin

Griffin is a open source Data Quality solution for distributed data systems at any scale in both streaming or batch data context

Griffin has been incubating since 2016-12-05.

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

1. Continuous enhancing to provide better accuracy and profiling data quality dimensions.
2. Release more dq dimensions like timeliness and completeness based on feedbacks from community.
3. Onboard more user cases and grow the community.

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

- Apache Griffin 0.1.6-incubating released.

How has the community developed since the last report?

- Four new contributors on-board, some more users had contacted us for use cases.
- Discussed and Voting to on-board some committers.

How has the project developed since the last report?

- Active development is moving on well, 67 commits in last three months.
- Profiling measure for streaming mode has been developed.
- Scheduler has been optimized, support flexible job scheduling.
- UI module for profiling measure has been developed.
- Work toward for next version to support enhanced scheduler end to end, and support more data quality dimensions.

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:

2017-11-07

When were the last committers or PPMC members elected?

- Initial set of committers / PPMC.
- Discussed and Voting to on-board some committers.

Signed-off-by:

- (griffin) Kasper Sørensen
Comments:
- (griffin) Uma Maheswara Rao Gangumalla
Comments:
- (griffin) Luciano Resende
Comments:
- (griffin) Henry Saputra
Comments:

IPMC/Shepherd notes:

Dave Fisher: Interesting project. Seems to be going well. No concerns.

Hivemall

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

Hivemall has been incubating since 2016-09-13.

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

1. Create the first Apache release
2. Community growth (committers and users)
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?

- * Got contributions from new external contributors
<https://github.com/apache/incubator-hivemall/pull/116>
- * Github watchers/stars are gradually increasing:
108 stars as of Dec 6 (was 99 on Sept 6)
- * Twitter account @ApacheHivemall followers are gradually increasing:
101 followers as of Dec 6 (was 87 on Sept 6)

How has the project developed since the last report?

- * Approaching the first Apache release while the released delayed from Q3 to Q4.
 - Blocking issues are decreased
<https://issues.apache.org/jira/browse/HIVEMALL-21>
 - Remaining tasks are just updating Maven release scheme and documents
<https://issues.apache.org/jira/browse/INFRA-15465>

Since the last report, we have

- * In the last 3 months, we opened 19 JIRA issues and closed 16 JIRA issues as seen in <https://goo.gl/QFQEF5> (as of Dec 1)

	Created	Resolved
Sept 2017	3	8
Oct 2017	12	5
Nov 2017	4	3

- * Created 14 Pull Requests and merged 21 Pull Requests between June 01 and Aug 31.
<https://github.com/apache/incubator-hivemall/pulls?q=is%3Apr%20created%3A2017-09-01..2017-11-30>
<https://github.com/apache/incubator-hivemall/pulls?q=is%3Apr%20closed%3A2017-09-01..2017-11-30>

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
(planning the first Apache release in Q4, 2017)

When were the last committers or PPMC members elected?

- * Oct 10: Invited Kento Nozawa as a committer

Signed-off-by:

- (hivemall) Reynold Xin
Comments:
- (hivemall) Markus Weimer
Comments:
- (hivemall) Xiangrui Meng
Comments:
- (hivemall) Daniel Dai
Comments:

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 issues to address in the move towards graduation:

1. Boost project dissemination in the Apache community.
2. Collaborate/integrate Omid with other Apache projects.
3. Get positive feedback from other projects currently in Apache.

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?

Continue integration with Apache Phoenix

How has the project developed since the last report?

Critical fixes on security and HBase 0.x & 1.x integration.
HBase 0.x version being tested in prod environment at Yahoo.
Started testing HBase version 1.x in local environments.

Quarter Stats (from: 2017-09-01 to: 2017-11-30):

Metric	counts
# of msgs in dev list	53
Active Contributors (incl mentors)	8
Jira New Issues	5
Resolved Issues	1
Pull Requests merged	1
Pull Requests proposed	2

How would you assess the podling's maturity?

Please feel free to add your own commentary.

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

Date of last release:

2016-06-24

When were the last committers or PPMC members elected?

Signed-off-by:

- (omid) Alan Gates
Comments:
- (omid) Lars Hofhansl
Comments:
- (omid) Flavio P. Junqueira
Comments:
- (omid) Thejas Nair
Comments:
- (omid) James Taylor
Comments:

OpenWhisk

OpenWhisk is an open source, distributed Serverless computing platform able to execute application logic (Actions) in response to events (Triggers) from external sources (Feeds) or HTTP requests governed by conditional logic (Rules). It provides a programming environment supported by a REST API-based Command Line

Interface (CLI) along with tooling to support packaging and catalog services. Additionally, it now provides options to host the platform components as Docker containers on various Container Frameworks such as Mesos, Kubernetes, and Compose.

OpenWhisk has been incubating since 2016-11-23.

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

1. Est. release process for ZIP of source code
2. Increase additional company and individual Contributors to maintain all project repos.

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

- It had been our intent to have Apache infra. to host our staging (and move it out of disparate private builds within various companies). However, INFRA indicates they only typically allocate 1 VM per project for use with demos and minor things. They indicate that perhaps a corporate donation is in order (ala. Spark).

- Goal: standup testing resources at Apache and utilize for public CI and performance testing of

OpenWhisk on Kubernetes

- Needed resources:

Minimum 5 worker node kubernetes cluster. Each worker nodes can be fairly modest (2-4 virtual cores; 4-8GB of memory)

- Formal hand-off of OpenWhisk trademark/logo from IBM needs to be executed; need to identify process for this.

Discussion started w/ Apache legal via "legal-discuss" mailing list with subject "Trademark handoff for "OpenWhisk" name and logo".

How has the community developed since the last report?

- incubator-openwhisk Github stars: 2181 (+253 since last report)
- incubator-openwhisk GitHub forks: 469 (+52 since last report)
- dev mailing list activity continues to be active
- The bi-weekly Zoom "Technical Interchange" continues to be well received and attended.
- Complete videos posted to OW YouTube channel and detailed notes to our CWIKI.
- YouTube Channel: Apache Meetings Playlist: https://www.youtube.com/watch?v=n419L8wyQlk&list=PLxVTI8yc_iX4qqr_45gduBFQxhSTfJ68g

- CWiki Meeting Notes: OpenWhisk Technical Interchange Meeting Notes: <https://cwiki.apache.org/confluence/display/OPENWHISK/OpenWhisk+Technical+Interchange+Meeting+Notes>

New Contributors of note (not comprehensive):

- Brendan McAddams: Red Hat: Akka expert working on creating separate Akka dispatchers for CouchDB and Kafka Clients

- Kavitha Vallari Devara: Samsung FaaS team in "Open Service Lab of Convergence Team (SW R&D Center): incubator-openwhisk organizing project directory structure.

- Gabriel Nicolas Avellaneda: independent. CLI contributions

- Manjiri Tapaswi: N. Caroline State Univ. (NCSU): Fixed all "good first issues" for "incubator-wskdeploy"

- Stig Bakken, Zedge (mobile phone customization), first review of PR in Kube repo.

How has the project developed since the last report?

OpenWhisk general developments:

- Reorg. of "core" openwhisk directories to better match architecture
 - Removal of openwhisk cli and cli related test cases
- Community agreed to Scala formatting guidelines; re-enforced via github pre-commit hooks
- Focus on improved per-invocation (of Actions) performance via logging and other improvements.
 - Changed default logging backend from log4j to Logback
- Redis is now optional on deployment
- Support for metrics streaming (via plug-in, default Kamon).
- Continued focus in HA/Clustering enablement while continuing to reduce reliance on Ansible for configuration. Work included:
 - Enabling Controller HA support (defaults to active cluster of 2).
 - Enabling CouchDB clustering
 - Using ZooKeeper as persistent store for InvokerId assignment
 - Adaptive timeouts for controller and nginx
 - Programmatic lazy creation of completedN and invokerN topics
- Certificate checking for secure mode
- Now generating maven snapshot artifacts for some sub-projects
- Kubernetes deployment improvements:
 - Moved Travis testing to using Minikube & more current Kubernetes version (1.7.4)
 - Significant progress on supporting deploy of complete OpenWhisk system on Kubernetes, eliminating custom containers, and streamlining deploy process

- Separation of Language runtimes into their own top-level repos. including: docker (SDK) runtime, pythin, nodejs, java, php, swift
- Updates to Akka and focus on improving messaging
- Updates to improve API Gateway resiliency
- Nginx performance improved
- NodeJS 6 upgraded to 6.11.4
- NodeJS 8 support
- Swift 3.1.1 now default, deprecated Swift 3

openwhisk-package-deploy:

- New repo created with proper license, docs, tests, and Travis CI
- New action wskdeploy allows to deploy wskdeploy git repos
- New "managed" sync support

openwhisk-packages-cloudant/couchdb:

- Update to feed action to update trigger config (dbName)
- Update to feed action to return config and status

openwhisk-packages-kafka:

- New ability for feed action to return config and status
- Update to feed action to update trigger config (topic)

openwhisk-packages-alarms:

- New ability for feed action to return config and status
- New alarm trigger type interval
- New alarm trigger type fire once
- New StartDate and stopDate parameters for cron

openwhisk-cli:

- New canonical source after duplication removal from main repo
- New open issues moved from main repo to CLI repo
- New trigger feed lifecycles on `trigger get` and `trigger update`

Significant. discussions (dev list):

- Warm containers
- Are Java actions first class citizens?
- Git commit (accreditation)
- Moving out runtime images and tests
- Enablement of controller clustering
- Invoker activation queueing proposal
- proposal to remove trigger activations when no rules are matched
- [Scala] code style guideline
- NodeJS 8 support
- Use an explicit property to define type of Whisk entity

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:

Comments:

Need greater variety of contributors and contributing companies

Date of last release:

N/A

When were the last committers or PPMC members elected?

None during this period

Signed-off-by:

- (openwhisk) Felix Meschberger
Comments:
- (openwhisk) Isabel Drost-Fromm
Comments:

[X](openwhisk) Sergio Fernández
Comments:

PageSpeed

PageSpeed represents a series of open source technologies to help make the web faster by rewriting web pages to reduce latency and bandwidth.

PageSpeed has been incubating since 2017-09-30.

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

1. Finish project setup
2. Community building
3. Create a first release

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

None

How has the community developed since the last report?

- A new potential committer indicated interest in joining the PageSpeed initiative by maintaining the FreeBSD port he contributed and updating it to the latest, and upstreaming required changes to support the new platform to make future updates easier.

How has the project developed since the last report?

- Some technical hurdles were encountered while trying to run the RAT tools on the codebase. The plan is to pick this up again after the initial code drop has been performed to the ASF repo.
- We are working on getting the source code into the ASF repository.
- Progress was made on new features and getting the module distributed more widely.

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?
2017-09-30 (entering incubation)

Signed-off-by:

- (pagespeed) Jukka Zitting
Comments:
- (pagespeed) Leif Hedstrom
Comments:
- (pagespeed) Nick Kew
Comments:
- (pagespeed) Phil Sorber
Comments:

Pulsar

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

subscribers, and cross-datacenter replication.

Pulsar has been incubating since 2017-06-01.

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

1. Grow the community with new Committers/PPMC members.
2. Set up a test cluster to be able to run system tests
3. Improve release process by having more committers taking part in 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?

The community added 5 contributors. There is a healthy discuss on issues related to development, tools and processes among the community members. There has been a good number of questions asked by users on the Slack channel where there are around 33 weekly active users.

How has the project developed since the last report?

18 authors have pushed 150 commits to master in the last 3 months.

The project has made the its second release in the Apache Incubator on October 13th. The next release (1.21-incubating) is being voted right now in the dev mailing list. Several design documents with improvement proposals have been submitted to the wiki and discussed in the mailing list.

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

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

Date of last release:

2017-10-13, 1.20.0-incubating

When were the last committers or PPMC members elected?

None

Signed-off-by:

[X](pulsar) Dave Fisher

Comments: Seems like the podling is engaging new contributors mostly through slack channel. We will need to make sure that enough information is archived by either discussion on dev@ or in some other way. They are on their fourth RC in the current release and are doing things correctly.

[X](pulsar) Jim Jagielski

Comments: Agreed. Need to push more actual conversation to the dev@ list

[X](pulsar) P. Taylor Goetz

Comments: I share the sentiment of other mentors. Would like to more dev@ list traffic vs. slack.

[](pulsar) Francis Liu

Comments:

Quickstep

Quickstep is a high-performance database engine.

Quickstep has been incubating since 2016-03-29.

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

1. Work on adoption of the Quickstep technology.
2. Continue building a developer community.
3. Work towards a second release.

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

None.

How has the community developed since the last report?

There are a number of active developers working with the code base, who may become committers in the near future.

How has the project developed since the last report?

A number of bug reports have been filed and features added to move the single-node version of the system forward. The newer features added to the system are accompanied by a JIRA issue. There is more activity on the dev mailing list, and the newer members of the community are involved in the discussions.

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

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

Date of last release:

2017-03-25

When were the last committers or PPMC members elected?

2017-10-03 Robert Claus (committer)

Signed-off-by:

- (quickstep) Konstantin Boudnik
Comments:
- (quickstep) Julian Hyde
Comments:
Community building is this project's biggest challenge.
Robert Claus is, I believe, the first committer outside
of the UW team, so his appointment is good progress.
- (quickstep) Roman Shaposhnik
Comments:

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 Flink, and Apache Samza.

SAMOA has been incubating since 2014-12-15.

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

1. Elect new PPMC members
2. Enlarge the community

3.

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

None

How has the community developed since the last report?

Mailing list activity (September 2017 - November 2017):

* @dev: 49 messages

Jira issues backlog (September 2017 - November 2017):

* Created: 5

* Resolved: 1

How has the project developed since the last report?

* Planning a new release in the next month.

* Presentation at the Flink Forward Conference from Orange's team

* Integration of a new Boosting algorithm

* Improved Apache Kafka support and support for JSON and Avro.

* Improved SAMOA instances.

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?

None

Signed-off-by:

(samoa) Alan Gates

Comments: In response to shepherd's comments (see below) community started a vote on a couple of contributors becoming committers.

(samoa) Ashutosh Chauhan

Comments:

(samoa) Enis Soztutar

Comments:

(samoa) Ted Dunning

Comments:

IPMC/Shepherd notes:

Concerned about the lack of community growth. Very few commits this year yet there are pull requests by contributors.

Sent an email to private@samo

SDAP

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

SDAP has been incubating since 2017-10-22.

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

1. Grow community
2. Make formal SDAP (Incubating) releases
3. Evangelize SDAP as an integrated data analytic center for Big Science problems

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 have not formally reported before. We hope the community to grow once we make our first Incubating release.

How has the project developed since the last report?

All infrastructure and codebase(s) have successfully been transitioned over to the ASF. We are working to comply with branding, trademarks and source code. This work is ongoing.

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?

No new committers or PPMC have been added other than the original SDAP committers.

Signed-off-by:

- (sdap) Jörn Rottmann
Comments:
- (sdap) Raphael Bircher
Comments:
- (sdap) Suneel Marthi
Comments:
- (sdap) Lewis John McGibbney
Comments: SDAP is doing pretty well. Branding, trademarks and source code compliance are the primary issues right now. Once these are addressed SDAP can focus on the first Incubating release.

SINGA

SINGA is a distributed deep learning platform.

SINGA has been incubating since 2015-03-17.

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

1. Finish the existing jira tickets, especially the Keras API and Open Neural Network Exchange format.
2. Prepare the graduation documents and start discussion in dev@ list.

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 new developer (Wengtong Li) contributed to the repository.
1 new committer (Moaz) joined officially.

How has the project developed since the last report?

1. We fixed several bugs related to Tensor operations and unit-testing.
2. Some new issues/features are created, mainly about supporting including Keras API and Open Neural Network Exchange format.

How would you assess the podling's maturity?

We are discussing the graduation process and want to start it with one more release.

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

Date of last release:

2017-06-29

When were the last committers or PPMC members elected?

2017-09-01

Signed-off-by:

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

Superset

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

Superset has been incubating since 2017-05-21.

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

1. Grow the community and enroll new committers.
2. Plan and execute our first Apache release
3. Make progress

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

* We're working on finding a way to link Jira Issues to GitHub Issues as the team has a preference for using GitHub Issues for tracking.

How has the community developed since the last report?

* Organic growth of our Github contributors (184), forks (2,675), watchers (888) and stars (16,969)

* Group of active committers meeting bi-weekly from Airbnb, Lyft and Twitter (Apple joining soon)

* Slack channel for Apache Superset created at <https://apache-superset.slack.com/>

How has the project developed since the last report?

* Committers have created initial UI/UX mocks and engineering plan for redesign of Explore and Dashboarding views

* A healthy, constant flow of bug fixes, quality improvements and new features, take a look at the project's Pulse on Github for more details

* The team has completed all of the items on the Incubation Checklist

How does the podling rate their own maturity.

Initial setup

Working towards first release

Community building

Nearing graduation

Other:

Date of last release: No official release yet since being voted into Apache Incubation. (Planning for the first Apache release in Q2, 2018)

When were the last committers or PPMC members elected?

* Grace Guo - PPMC & Committer (2017-10-29)

* Chris Williams - PPMC & Committer (2017-10-20)

* John Bodley - PPMC & Committer (2017-10-20)

Signed-off-by:

(superset) Ashutosh Chauhan

Comments:

(superset) Luke Han

Comments:

(superset) Jim Jagielski

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 issues to address in the move towards graduation:

1. Re-engage PMC to encourage regular activity
2. Retiring irrelevant git repositories
3. Graduate!

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

Community slow. Considering retiring, but will try one last push.

How has the community developed since the last report?

Quite a few pull requests, but need follow ups.

Non-PPMC committers have taken more charge in release preparation.

Email stats since 2017-09-01:

dev@taverna: 29

users@taverna: 2

How has the project developed since the last report?

Very slow progress towards graduation.

Some progress towards next releases - need implicit Release Manager delegation.

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:

2016-07-01

When were the last committers or PPMC members elected?

2016-10-20

Signed-off-by:

[x](taverna) Andy Seaborne
Comments:

In Dec 2016, the report said "Preparing to graduate." In 2017, the dev@ list has been quiet and the last release was in 2016.

While the PPMC is large, not everyone has been active in incubation. I am worried that pushing through graduation will only result in the same situation as a TLP.

The community needs to consider what is the best environment for Taverna.

- (taverna) Daniel J Debrunner
Comments:
- (taverna) Marlon Pierce
Comments:
- (taverna) Stian Soiland-Reyes
Comments: I have admitted lack of leadership. I think we should seriously consider retiring, but only after a 6 month push towards getting graduation done.
- (taverna) Suresh Marru
Comments:
- (taverna) Suresh Srinivas
Comments:

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.

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

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?

- 1 new subscriber in dev mailing list since the last report
- 18 new JIRAs filed since the last report
- 1 external contributors submitted patches since the last report
- 1 external contributor created a ticket in JIRA since the last report

How has the project developed since the last report?

- Tephra 0.13.0-incubating released
- Working on 0.14.0-incubating 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 the last release:

2017-09-21

When were the last committers or PMC members elected?

- None since coming to incubation

Signed-off-by:

- (tephra) Alan Gates
Comments:
- (tephra) Andrew Purtell
Comments:
- (tephra) James Taylor
Comments:
- (tephra) Lars Hofhansl
Comments:

Trafodion

Trafodion is a webscale SQL-on-Hadoop solution enabling transactional or operational workloads on Hadoop.

Trafodion has been incubating since 2015-05-24.

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

- 1.
- 2.
- 3.

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

How has the community developed since the last report?

- New mentor added: Jacques Le Roux
- Slowly attracting more adopters and subscribers.
- ML postings: 37 user, 396 dev, 528 coderev, 668 issues, 533 commits
- Twitter: 228 followers, 73 tweets, 66 likes

How has the project developed since the last report?

- * 115 commits from 17 contributors.
- * 95 tickets filed and 78 resolved.
- * We are working towards our next release, 2.2. Ming Liu is the Release Manager.

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:

Request to graduate submitted to board after discussions.

Date of last release:

2017-05-01 2.1

When were the last committers or PPMC members elected?

2017-03-10

Signed-off-by:

[](trafodion) Michael Stack

Comments:

[X](trafodion) Jacques Le Roux

Comments: I'm new to this poddling, but I have already seen that the activity is steady, the CI pipeline is working good and over all I like the responsiveness and kindness of the community. Ready for graduation!