

# September2017

/\ FINAL (!)

## Incubator PMC report for September 2017

### Timeline

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

### Shepherd Assignments

Dave Fisher	Heron		
Dave Fisher	SINGA		
Drew Farris	Spot		
Drew Farris	Wave		
John Ament	Gobbli n		
John Ament	HTrace		
Justin Mclean	Mnemo nic		
P. Taylor Goetz	Griffin		
P. Taylor Goetz	Hivem all		
P. Taylor Goetz	Omid		
Timothy Chen	AriaTo sca		
Timothy Chen	MRQL		
Timothy Chen	Supers et		
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="5da66ed0-8c64-41d6-a786-db4f1ba9e12d"><ac:plain-text-body><![CDATA[	[none]	Gearpu mp	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e058e952-fc42-415b-89de-e02a474fb220"><ac:plain-text-body><![CDATA[	[none]	Livy	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="1ecc6929-98c6-483e-806a-d9f3e5d9247b"><ac:plain-text-body><![CDATA[	[none]	Myriad	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="487e17e0-93e7-4e52-80ae-8c5f4eb64791"><ac:plain-text-body><![CDATA[	[none]	OpenW hisk	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="07a3ff7e-1030-4761-9249-77b62400ddd3"><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="524addfe-ef0e-4e36-ade8-f132c517e472"><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="23b85c3f-59be-4f0a-a788-ad4cc26f1926"><ac:plain-text-body><![CDATA[	[none]	Quickst ep	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="e95418cf-0759-4af3-b9ff-d7d5c928b574"><ac:plain-text-body><![CDATA[	[none]	SAMOA	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="da70118c-a41f-4171-b45e-b1ab85bd4cf4"><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="1e02a3fe-9e57-417a-95f7-a352fe75d227"><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="9a211cac-0a57-49b9-953d-d88a3ee087fc"><ac:plain-text-body><![CDATA[	[none]	Trafodio n	]]></ac:plain-text-body></ac:structured-macro>
<ac:structured-macro ac:name="unmigrated-wiki-markup" ac:schema-version="1" ac:macro-id="ef68afd4-bd08-4b3e-9875-95cf56ef0e16"><ac:plain-text-body><![CDATA[	[none]	Daffodil	]]></ac:plain-text-body></ac:structured-macro>

## Report content

Incubator PMC report for September 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 54 podlings incubating. We executed nine podling releases and have one podling planning to graduate this month. No changes to the PMC structure this month.

### \* Community

New IPMC members:

- None

People who left the IPMC:

- None

### \* New Podlings

- Amaterasu
- Daffodil

### \* Podlings that Retired

- MRQL

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

- Myriad
- Spot

### \* Graduations

The board has motions for the following:

- RocketMQ
- Mnemonic (may be next month)

### \* Releases

The following releases entered distribution during the month of August:

- 2017-08-01 Apache Juneau 6.3.1
- 2017-08-01 Apache Tamaya 0.3.0
- 2017-08-05 Apache Tamaya Extensions 0.3.0
- 2017-08-08 Apache Pulsar 1.19.0
- 2017-08-16 Apache HTrace 4.3.0
- 2017-08-16 Apache Spot 1.0
- 2017-08-17 Apache Fluo 1.0.0
- 2017-08-23 Apache S2Graph 0.2.0
- 2017-08-29 Apache Livy 0.4.0

### \* Legal / Trademarks

- Thanks to hard work on both the Oracle and ASF side, the Netbeans team was able to get the first three repositories populated from a software grant.

---

### Table of Contents

AriaTosca  
Daffodil  
Gearpump  
Gobblin  
Griffin  
Heron

Hivemall  
HTrace  
Livy  
Mnemonic  
MRQL  
Myriad  
Omid  
OpenWhisk  
Pony Mail  
Pulsar  
Quickstep  
SAMOA  
SINGA  
Spot  
Superset  
Taverna  
Tephra  
Trafodion  
Wave

-----  
-----  
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. Rotate Release Managers for future releases.
- 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?

Seeing some contributions now from independent developers.  
Active discussions on mailing lists with potential users and contributors.

How has the project developed since the last report?

The project has had its first two releases in July 2017 - 0.1 and 0.1.1

Important new features have been added:

- Resumable workflows
- Scaling nodes during service creation
- AWS plugin support
- Additional work towards being more TOSCA-compliant (e.g. interface inputs support)
- Additional usage examples have been added, including ones that use IaaS plugins (Openstack, AWS), as well as one covering a real-life NFV use-case (Clearwater example)
- More comprehensive tests, including for documentation generation, and parallelization of test suites
- Standardized the release process with automatic shell scripts, allowing anyone to easily create additional releases for the project
- 64 JIRA issues have been resolved 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
- [X] Community building
- [ ] Nearing graduation

☐ Other:

Date of last release:

July 17, 2017

When were the last committers or PPMC members elected?

Project is still functioning with the initial set of committers.

Signed-off-by:

☒ (ariatosca) Suneel Marthi

Comments:

☒ (ariatosca) John D. Ament

Comments: The project has started to see more community interest. Need to figure out how to drive development more on list to get more outside contributions.

☐ (ariatosca) Jakob Homan

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. Initial setup (finalize infrastructure setup, get SGAs in place, move code/documentation/bugs/etc. over to Apache infrastructure, modify licenses)
2. Establish guidelines regarding release cycles, code reviews, bug management, etc.
3. Work to grow the community to establish a diversity of backgrounds and expertise

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?

N/A

How would you assess the podling's maturity?

Please feel free to add your own commentary.

☒ Initial setup

☐ Working towards first release

☐ Community building

☐ Nearing graduation

☐ Other:

Date of last release:

N/A

When were the last committers or PPMC members elected?

N/A

Signed-off-by:

☒ (daffodil) John D. Ament

Comments:

☒ (daffodil) David Fisher

Comments: Just starting off.

-----

## 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. Increase user adoption.

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?

- Gearpump took part in Google Summer of Code 2017 project and it has been successfully accomplished by student, Buddhi Ayesha, who is interested in making more contributions beyond the project.

How has the project developed since the last report?

- 32 issues created and 20 issues resolved.  
- Gearpump Runner has been merged into Apache Beam trunk.  
- An experimental SQL support has been added and state management is under development.

How would you assess the podling's maturity?

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

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:

[X](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:

-----

## Gobblin

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

Gobblin has been incubating since 2017-02-23.

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

1. Cut our 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?

- \* We are working towards our first release.
- \* Email stats for last month:  
user@gobblin.incubator.apache.org : 14  
dev@gobblin.incubator.apache.org : 259
- \* There have been 54 Commits in last month:  
git log --format='%ci' | grep -cE '2017-0(8)'
- \* 30 of those commits were by non-committers:  
git log --format='%ae %ci' | grep -E '2017-0(8)' | cut -d ' ' -f 1 | sort | uniq -c | sort -n
- \* A video conference based meetup happened last month.

How has the project developed since the last report?

- \* Site has been setup.
- \* Apache wiki has been populated with relevant content.
- \* Code development is actively being tracked via JIRA / Sprint dashboard.

How would you assess the podling's maturity?

The podling is working towards its first release. Like last time, continued progress and activity on all fronts.

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

Date of last release:

N/A

When were the last committers or PPMC members elected?

N/A

Signed-off-by:

- ☐ (gobblin) Jean-Baptiste Onofré  
Comments: Waiting the board report. I will help/ping.
- ☒ (gobblin) Olivier Lamy  
Comments:
- ☐ (gobblin) Jim Jagielski  
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 development of more dq dimensions for batch and streaming mode.
2. Make the next release.
3. Grow 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?

- Three new contributors on-board, some users had contacted us for use cases.

How has the project developed since the last report?

- Active development is moving on well, 83 commits in last three months.
- The first release of Griffin has been published.
- Accuracy measure for streaming mode has been developed.
- Griffin DSL has been optimized, to support accuracy and profiling cases well.
- Service module refined API, UI module upgrade to Angular 2.
- Work toward for next release version (general enough) to support streaming dq measure end to end.

How would you assess the podling's maturity?

Please feel free to add your own commentary.

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

Date of last release:

2017-07-17

When were the last committers or PPMC members elected?

- Initial set of committers / PPMC.

Signed-off-by:

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

-----  
Heron

A real-time, distributed, fault-tolerant stream processing engine.

Heron has been incubating since 2017-06-23.

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

1. Finish bootstrapping project, IP clearance, initial website
2. Expanding the community and adding new committers
3. 1st ASF release

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?

- \* Some issues encountered with new mailing list signups have been resolved.
- \* All initial committers have submitted ICLAs and have gotten their ASF accounts setup.

How has the project developed since the last report?

\* The project has made 3 non-Apache releases.

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 Apache releases as of yet. Working on code import.

When were the last committers or PPMC members elected?

N/A, still bootstrapping the project.

Signed-off-by:

☒ (heron) Jake Farrell  
Comments:  
☐ (heron) Jacques Nadeau  
Comments:  
☐ (heron) Julien Le Dem  
Comments:  
☒ (heron) P. Taylor Goetz  
Comments: Currently waiting on SGA (pinged for status, but no response yet). Little activity on mailing lists so far. Decision making seems to be taking place off-list (3 non-Apache releases). I concur with Dave Fisher that monthly reporting should be considered (although this is report #2 in their initial monthly reporting cycle).

IPMC/Shepherd notes:

Dave Fisher: It seems that with 3 non-apache releases while there is not even any use of Apache resources that Heron ought to focus on migration. I think that they should be monthly until the code is moved.

-----  
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/91>  
<https://github.com/apache/incubator-hivemall/pull/93>  
<https://github.com/apache/incubator-hivemall/pull/111>  
<https://issues.apache.org/jira/browse/HIVEMALL-124>



- \* Udemy course featuring Hivemall was opened by someone outside PPMC  
<https://www.udemy.com/new-in-big-data-apache-hivemall-machine-learning-with-sql/>
- \* Github watchers/stars are gradually increasing:  
 99 stars as of Sept 6 (was 79 on May 31)
- \* Twitter account @ApacheHivemall followers are gradually increasing:  
 87 followers as of Sept 6 (was 70 on May 31)

How has the project developed since the last report?

- \* Approaching the first Apache release while the released delayed from Q2 to Q3 due to delay in merging large patches.
  - Waiting for incoming pull requests to review and merge  
<https://github.com/apache/incubator-hivemall/pulls>
  - Merging HIVEMALL-24 takes time more than expected  
<https://github.com/apache/incubator-hivemall/pull/105>
  - Blocking issues are decreased  
<https://issues.apache.org/jira/browse/HIVEMALL-21>

Since the last report, we have

- \* In the last 90 days, we opened 29 JIRA issues and closed 21 JIRA issues as seen in <https://goo.gl/88Qr11> (as of Sept 6)

		Created	Resolved
June	2017	7	8
July	2017	15	12
Aug	2017	5	0
Sept	2017	2	1

- \* Created 30 Pull Requests and merged 24 Pull Requests between June 01 and Aug 31.  
<https://github.com/apache/incubator-hivemall/pulls?q=is%3Apr%20created%3A2017-06-01..2017-08-31>  
<https://github.com/apache/incubator-hivemall/pulls?q=is%3Apr%20closed%3A2017-06-01..2017-08-31>

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

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

Date of last release:

No release yet  
 (planning the first Apache release in Q3, 2017)

When were the last committers or PPMC members elected?

- \* June 5: Invited Takuya Kitazawa to PPMC

Signed-off-by:

```
[ ](hivemall) Reynold Xin
  Comments:
[X](hivemall) Markus Weimer
  Comments:
[ ](hivemall) Xiangrui Meng
  Comments:
[ ](hivemall) Daniel Dai
  Comments:
```

-----  
 HTrace

HTrace is a tracing framework intended for use with distributed systems written in java.

HTrace has been incubating since 2014-11-11.

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

1. Continue to grow the Apache HTrace community
2. Continue to explore projects integrating Apache HTrace

3. Continue to develop and release stable Apache HTrace incubating artifacts

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

In a DISCUSSION thread on our dev list [1], one of us suggested moving the podling htrace to the attic because of inactivity.

A healthy back and forth suggested that we are not dead yet, that we cut back our scope radically, and that we give it another try before calling it quits.

1. <https://s.apache.org/9ybr>

How has the community developed since the last report?

Mailing lists continue quiet (bar the above noted thread).

How has the project developed since the last report?

No development.

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

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

Date of last release:

2016-10-06

When were the last committers or PPMC members elected?

2016-10-03

Signed-off-by:

[ ](htrace) Jake Farrell  
Comments:  
[ ](htrace) Todd Lipcon  
Comments:  
[x](htrace) Lewis John McGibbney  
Comments: I think the following comment resonates on the mailing list  
"...A healthy back and forth suggested that we are not dead yet,  
that we cut back our scope radically, and that we give it another  
try before calling it quits." I think HTrace may have a good opportunity  
to come back to life if the above cut backs are implemented.  
[ ](htrace) Andrew Purtell  
Comments:  
[x](htrace) Billie Rinaldi  
Comments:  
[X](htrace) Michael Stack  
Comments:

-----  
Livy

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

Livy has been incubating since 2017-06-05.

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

1. Grow more committers to maintain the community

2. Expand the community, be more open to the newcomers
3. Have a regular release cadence

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?

We are seeing nascent activity on the new ASF mailing lists for users.

How has the project developed since the last report?

We released the first Apache release of Livy, and migrated all the assets including JIRA, website to ASF.

One big feature about support shared session across languages was merged.

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:

2017-09-01

When were the last committers or PPMC members elected?

N/A

Signed-off-by:

- ☒ (livy) Bikas Saha  
Comments:
- ☐ (livy) Brock Noland  
Comments:
- ☐ (livy) Luciano Resende  
Comments:
- ☒ (livy) Jean-Baptiste Onofré  
Comments:

-----  
Mnemonic

Apache Mnemonic is a transparent NVM hybrid memory oriented library that has created Durable Object Model (DOM), Durable Computing Model (DCM) and Durable Query Model (DQM) that bring many advantages to significantly improve the performance of massive real-time data processing/analytics. developers are able to use this library to design their cache-less, SerDe-less and direct-native high-performance applications or frameworks.

Mnemonic has been incubating since 2016-03-03.

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

1. Have elected new PPMC, Release Manager and committers
2. Have proposed Durable Query Model (DQM) based on DOM and DCM
3. Have prepared Apache maturity assessment checklist for review

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

Not yet

How has the community developed since the last report?

Our last report was in Jun. Since then

- \* fulfilled the gaps from graduation
- \* elected new release manager
- \* elected new committers
- \* introduce owner of CQEngine to our community
- \* some contributors have started to vote for graduation

How has the project developed since the last report?

- \* improved Apache mnemonic official website
- \* implemented durable query in memory service level
- \* implemented pure Java volatile memory service
- \* improved the documentation of project specific release process
- \* released mnemonic-0.8.0-incubating
- \* started to vote mnemonic-0.9.0-incubating

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

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

Date of last release:

2017-06-23

When were the last committers or PPMC members elected?

Johnu George (PPMC on Aug. 04 2017)  
PeiLi Shen (Committer on Aug. 13 2017)

Signed-off-by:

- ☒ (mnemonic) Patrick Hunt  
Comments: afaict ready to graduate (there is an active DISCUSS thread going on now)
- ☐ (mnemonic) Andrew Purtell  
Comments:
- ☒ (mnemonic) James Taylor  
Comments:
- ☐ (mnemonic) Henry Saputra  
Comments:

-----  
Myriad

IPMC/Shepherd notes:

johnndament: No report from podling. Discussion back in July indicated the podling did not want to retire and had enough active members to stay alive. No on list activity since then.

-----  
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?

Started integration with Apache Phoenix

How has the project developed since the last report?

Several bugs discovered and fixed.  
Continue testing prod environment at Yahoo.  
Branch 0.9.0.0 started for next release

Quarter Stats (from: 2017-03-01 to: 2017-05-31):

+-----+		
Metric	counts	
+-----+		
# of msgs in dev list	108	
Active Contributors (incl mentors)	8	
Jira New Issues	6	
Resolved Issues	6	
Pull Requests merged	3	
Pull Requests proposed	3	
+-----+		

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

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

Date of last release:

2016-06-24

When were the last committers or PPMC members elected?

Signed-off-by:

[X](omid) Alan Gates

Comments: Activity continues to be low but there has been a noticeable uptick since the last report, and there is now a consistent if low level of activity on the project. The big challenge now for Omid is to attract new contributors to the project.

[ ](omid) Lars Hofhansl

Comments:

[ ](omid) Flavio P. Junqueira

Comments:

[ ](omid) Thejas Nair

Comments:

[X](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. Est. release process for Docker and other platform binaries
3. 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 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 indicated that typically a donation is in order (ala. Spark).

- Note: the desire for this staging capability was expressed in the Incubator submission where we anticipated needing ~12 VMs to host each platform component and be able to load balance/schedule Functional code as part of the OpenWhisk Serverless (across multiple supported language runtimes). This would also be needed to effectively test auto-scaling/load balancing and container pools and to get accurate performance test measures.

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

How has the community developed since the last report?

- First substantive design changes allowing for Service Provider Interfaces (SPIs) developed within Apache community (dev list, Wiki, etc.).

- Established bi-weekly Zoom "Technical Interchange" calls that have been well received.

- Complete videos posted to OW YouTube channel and detailed notes to our CWIKI.

- Kubernetes on OpenWhisk work hosted a few ad-hoc to continue design discussions (CWIKI has designs and comments)

- for how best to move away from Ansible deployments to self-configuring, autoscalable containers (for use in Kube, Mesos, Compose, etc.).

- Increased interest/list/Slack traffic in general due to experiment Kube work, as it is applicable to Mesos, etc.

- Increased community discussions (design topics, features) of substance moved to dev list.

How has the project developed since the last report?

- incubator-openwhisk Github stars: 1928

- incubator-openwhisk GitHub forks: 417

- Movement to become more "pluggable" via Service Provider Interfaces (SPI), several PRs that are ongoing.

- So far there are now interfaces for MessagingProvider.scala, ArtifactStoreProvider.scala and proposals for others such as LogStore, ContainerFactory, LoadBalancer,

Sig. PRs:

- initial work on no root ansible #2669

- This allows to deploy with no sudo/root access. Significant to enable Kube/Mesos deployments.

- Simplify Spi implementations. #2663

- Service Provider Interfaces (SPI) become more mature as patterns emerge.

- Healthy discussion/growth area for project to allow most components to become (eventually) pluggable (e.g., Log, Trace, Container build, Load Balancing, Scheduler, Routing, Message Queue, etc.)

- Apply standard scala formatting #2650

- Coding/linting standards debated/agreed to by community (and CI automated enforcement encoded)

- Allow multiple controllers and invokers in local environments #2633

- implement a ContainerPool which is completely independent of the underlying container driver to facilitate testing and extensibility.

- More improvements to allow different pooling of invoker containers based upon different application use case

- (new ideas generated in community) including requirements for diff. scheduling/co-tenanting,/user responsiveness (http), etc. needs.

Sig. discussions (dev list):

- Release for OpenWhisk

- Releases and SPI plans

- OpenWhisk shell tool

- authentication models for openwhisk

- Performance tests for OpenWhisk

- Adaption of hosts-files

- Default handling for HTTP web actions

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?

- Tyson Norris, Committer, 08/31/2017
- James Dubee, PPMC, Committer, 09/05/2017

Signed-off-by:

- ☐ (openwhisk) Felix Meschberger  
Comments:
- ☐ (openwhisk) Isabel Drost-Fromm  
Comments:
- ☒ (openwhisk) Sergio Fernández  
Comments:

-----  
Pony Mail

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 issues to address in the move towards graduation:

1. Work on the next release (0.10)
2. Work on expanding the community, be more open to newcomers
3. Resolve design strategies

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

None that I am aware of.

How has the community developed since the last report?

No major development. A few more bug reports from outside contributors have been filed.

How has the project developed since the last report?

Work continues towards the next release, 0.10, which should be ready for a vote soon.

How would you assess the podling's maturity?

The podling is somewhat mature, but it is my assessment that more people are needed for it to be a fully viable community.

Please feel free to add your own commentary.

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

Date of last release:

2016-08-02

When were the last committers or PPMC members elected?

sebb as PMC, 2017-04-17.

No new committers as of late.

Signed-off-by:

[ ](ponymail) Andrew Bayer

Comments:

[X](ponymail) John D. Ament

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 2 contributors. There is a healthy discuss on issues related to development, tools and processes among the community members.

How has the project developed since the last report?

8 authors have pushed 69 commits to master in the last 30 days.  
The project has made the first release from the Apache Incubator on August 9th.  
The next release is planned for mid-September. 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

[X] Community building

[ ] Nearing graduation

[ ] Other:

Date of last release:

2017-08-09, 1.19.0-incubating

When were the last committers or PPMC members elected?

<none>

Signed-off-by:



[X](pulsar) Dave Fisher  
Comments: Project is moving towards a second release  
[ ](pulsar) Jim Jagielski  
Comments:  
[x](pulsar) P. Taylor Goetz  
Comments: Pulsar seems to be doing well.  
[ ](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. Building a Quickstep community
2. More adoption of the Quickstep technology
3. More releases

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

None

How has the community developed since the last report?

In the summer of 2017, we took efforts to spread awareness about the Quickstep project and technology. We had many talks, particularly in the San Francisco Bay Area. The talks/venues are listed here:

1. SRI International in Menlo Park
2. A meetup hosted by H2O.ai in Mountain View
3. Hortonworks in Santa Clara (Thanks to our mentor Julian Hyde for hosting)

We had also applied for a talk at BigDataWisconsin meetup but the proposal was not selected.

How has the project developed since the last report?

During the summer, the development activity was slow. There are ongoing efforts to build a distributed version of Quickstep.

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

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

Date of last release:

2017-03-25

When were the last committers or PPMC members elected?

2017-05-12 Tianrun Lee (committer)

Signed-off-by:

- [ ](quickstep) Konstantin Boudnik  
Comments:
- [x](quickstep) Julian Hyde  
Comments:

Community building remains the project's biggest challenge.  
The project has been putting considerable effort into talks  
and I hope this effort starts to pay off.

[ ](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 (June 2017 - August 2017):  
\* @dev: 120 messages

Jira issues backlog (June 2017 - August 2017):  
\* Created: 4  
\* Resolved: 5

We saw the engagement of new contributors, who are plausible candidates for the PPMC.

How has the project developed since the last report?

\* Planning a new release in the next month. Work on the release was delayed due to summer holidays for most of the contributors.

\* Integrated Apache Kafka support and support for JSON and Avro.

\* Developed a new Boosting algorithm that will be integrated shortly.

\* Integrated new version of the SAMOA instances.

\* Integrated ability to store predictions on disk.

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

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

Date of last release:

2016-09-30

When were the last committers or PPMC members elected?

None

Signed-off-by:

[X](samoa) Alan Gates  
Comments:  
[ ](samoa) Ashutosh Chauhan  
Comments:  
[ ](samoa) Enis Soztutar  
Comments:  
[ ](samoa) Ted Dunning  
Comments:

-----  
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. Prepare all documentations and go through the graduation voting process.
- 2.
- 3.

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?

A new committer (moaz) has been added into SINGA.

How has the project developed since the last report?

1. We extended SINGA Python support from Py2 to Py3.
2. We added CPack to generate debian packages for SINGA
3. Several example models (VGG and InceptionV3/V4) were added.
4. A couple of bugs were fixed and the documentation was updated.

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

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

Date of last release:

2017-12-02

When were the last committers or PPMC members elected?

2017-29-08

Signed-off-by:

[ ](singa) Daniel Dai  
Comments:  
[X](singa) Alan Gates  
Comments:  
[ ](singa) Ted Dunning  
Comments:  
[ ](singa) Thejas Nair  
Comments:

IPMC/Shepherd notes:

Dave Fisher - Graduation looks close.

-----  
Spot

IPMC/Shepherd notes:

- Drew Farris (shepherd): No mentor activity apparent from the mailing lists, ongoing project activity observed.
- John D. Ament: Podling needed some reminding about the board report, I suspect next report will be in and of high quality.

-----  
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. Address the challenge that React is no longer allowed in Apache projects
2. Grow the community and enroll new committers.
3. Plan and execute our first Apache release

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

\* Superset is heavily reliant upon React, but React is no longer allowed in ASF projects (<https://code.facebook.com/posts/112130496157735/explaining-react-s-license/>). We are currently evaluating Preact as an alternative. (<https://github.com/apache/incubator-superset/pull/3355>)

How has the community developed since the last report?

\* Grace Guo and Riccardo Magliocchetti have been elected as Superset Committers  
\* Organic growth of our Github contributors (160), forks (2,338), watchers (829) and stars (15,807)

How has the project developed since the last report?

\* We have completed all of the items on the Apache Incubation Checklist  
\* We have a healthy, constant flow of bug fixes, quality improvements and new features. Take a look at the project's Pulse on Github for more details

How would you assess the podling's maturity?

Please feel free to add your own commentary.

[ ] Initial setup  
[X] 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 Q4, 2017)

When were the last committers or PPMC members elected?

Grace Guo and Riccardo Magliocchetti were both elected on August 17, 2017

Signed-off-by:

[X](superset) Ashutosh Chauhan  
Comments:  
[X](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. Complete the remaining IP review (by retiring irrelevant git repositories)
2. Re-engage PMC to encourage regular activity
3. Graduate!

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?

Fairly quiet over the summer. Lack of progress towards graduation. Need more active drive.

Email stats since 2017-07-01:

dev@taverna: 74 emails  
users@taverna: 4 emails

How has the project developed since the last report?

Progress towards release of the Taverna Mobile Android App.  
Been in contact with legal/infra to figure out how to do code signing for Google Play store.

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:

☐ (taverna) Andy Seaborne  
Comments:  
☐ (taverna) Daniel J Debrunner  
Comments:  
☐ (taverna) Marlon Pierce  
Comments:  
☒ (taverna) Stian Soiland-Reyes  
Comments: Mailing list has gone very quiet  
☐ (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
- 16 new JIRAs filed since the last report
- 2 external contributors submitted patches since the last report
- 4 external contributor created a ticket in JIRA since the last report

How has the project developed since the last report?

- Working on 0.13.0-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:

2017-05-30

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

Apache Trafodion is a webscale SQL-on-Hadoop solution enabling transactional or operational workloads on Hadoop. Trafodion builds on the scalability, elasticity, and flexibility of Hadoop. Trafodion extends Hadoop to provide guaranteed transactional integrity, enabling new kinds of big data applications to run on Hadoop.

Trafodion has been incubating since 2015-05-24.

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

1. Develop the community and continue to grow the diversity.
2. Continue the momentum and become more integrated with the rest of the Apache community.
3. Continue to create software releases.

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

of?

- \* Two of our three remaining mentors, Devaraj Das and Enis Söztutar, resigned on July 21.

How has the community developed since the last report?

- \* Participation in the project remains healthy, with the usual small dip due to the summer vacation season:
  - ML postings: 106 user, 273 dev, 389 coderev, 734 issues, 341 commits
  - Twitter: 221 followers, 72 tweets, 67 likes
- \* We are working towards our next release, 2.2. Ming Liu is the Release Manager.
- \* Discussion regarding readiness to graduate continues. Diversity continues to grow.

How has the project developed since the last report?

- \* 162 commits from 18 contributors.
- \* 100 JIRAs filed and 63 resolved.

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

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

Date of last release:

2017-05-01 2.1

When were the last committers or PMC members elected?

2017-03-10

Signed-off-by:

[x](trafodion) Michael Stack  
Comments:

-----  
Wave

A wave is a hosted, live, concurrent data structure for rich communication. It can be used like email, chat, or a document.

Wave has been incubating since 2010-12-04.

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

1. Appoint members of the PMC
2. Update release procedure
- 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?

There were no relevant events regarding the community in this period.

How has the project developed since the last report?

Some members of community work on side projects - like UI improvements, but overall we had a low level of activity.

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:

March 2016

When were the last committers or PPMC members elected?

March 2017

Signed-off-by:

☒ (wave) John D. Ament

Comments: Reached out to get the dates fixed up. In order to graduate, we have to figure out how to get more activity.

☐ (wave) Upayavira

Comments:

IPMC/Shepherd notes:

Drew Farris (shepherd): Two mentors active. Light activity on the project in general, some discussion towards graduation a couple months back.