

October2013

⚠ FINAL ⚠

This report has been submitted to the Board. Please do not edit further.

Incubator PMC report for October 2013

Timeline

Wed October 2	Podling reports due by end of day
Sun October 6	Shepherd reviews due by end of day
Sun October 6	Summary due by end of day
Tue October 8	Mentor signoff due by end of day
Wed October 9	Report submitted to board
Wed October 16	Board meeting

Shepherd Assignments

Raphael Bircher	Olingo
Alan Cabrera	Celix
Alan Cabrera	VXQuery
Dave Fisher	DeviceMap
Dave Fisher	Spark
Matt Franklin	ODF Toolkit
Matt Franklin	Sentry
Ross Gardler	Ripple
Suresh Marru	Samza
Suresh Marru	MetaModel
Andrei Savu	Marmotta
Roman Shaposhnik	Helix
Roman Shaposhnik	Stratos
	Chukwa (graduating)
	jclouds (graduating)
	Storm (just starting out)

Report content

Incubator PMC report for October 2013

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

The Incubator grew to 39 podlings this month.

Mailing list activity on `general@incubator` went back up after light months in July and August.

* Community

New IPMC members:

Tammo van Lessen
Christian Müller

People who left the IPMC:

(None)

* New Podlings

The Incubator has voted in four new podlings.

* Aurora is a scheduler that provides all of the primitives necessary to quickly deploy and scale stateless and fault tolerant services in a datacenter.

* Storm is a distributed, fault-tolerant, and high-performance realtime computation system that provides strong guarantees on the processing of data.

* BatchEE will be an ALv2-licensed implementation of the JBatch Specification which is defined as JSR-352 (for version 1.0).

* Usergrid is a multi-tenant Backend-as-a-Service stack for web & mobile applications, based on RESTful APIs.

* Graduations

The board has motions for the following:

Chukwa
jclouds

* Releases

The following releases were made since the last Incubator report:

Sep 11 Apache Ambari 1.2.5-incubating
Sep 27 Apache Drill 1.0.0-m1-incubating
Sep 24 Apache Spark 0.8.0-incubating
Sep 25 Apache Sentry 1.2.0-incubating
Sep 27 Apache Blur 0.2.0-incubating
Oct 03 Apache Marmotta 3.1.0-incubating

It took 1-8 days for the third IPMC vote to arrive.

Release	RC VOTE start	Third IMPC +1	Days
Apache Ambari 1.2.5-incubating	Sep 04	Sep 10	6
Apache Drill 1.0.0-m1-incubating	Sep 11	Sep 18	7
Apache Spark 0.8.0-incubating	Sep 17	Sep 23	6
Apache Sentry 1.2.0-incubating	Sep 19	Sep 24	5
Apache Blur 0.2.0-incubating	Sep 18	Sep 19	1
Apache Marmotta 3.1.0-incubating	Sep 24	Oct 02	8

* Miscellaneous

* Allura's release candidate, which spent weeks waiting for IPMC feedback, was ultimately withdrawn after flaws were discovered.

* Commons-monitoring seems likely to enter the Incubator. The name is still under discussion.

* A question about old Chukwa binary artifacts was resolved by reaffirming that the ASF makes no guarantees about convenience binaries.

* The Incubator paid elevated attention to trademarks this month.

* Old controversies about what criteria to apply when adding personnel to

a proposed podling finally bit an unlucky live podling candidate (Usergrid). Hopefully it will never happen again.

----- Summary of podling reports -----

* Still getting started at the Incubator

Storm

* Not yet ready to graduate

No release:

MetaModel
Olingo
Samza
Stratos
VXQuery

Community growth:

Celix
ODF Toolkit
Sentry

* Ready to graduate

Chukwa
Helix
jclouds
Marmotta

* Did not report, expected next month

DeviceMap
Ripple

Table of Contents

Celix
Chukwa
DeviceMap
Helix
jclouds
Marmotta
MetaModel
ODF Toolkit
Olingo
Ripple
Samza
Sentry
Spark
Storm
Stratos
VXQuery

Celix

Celix is an OSGi like implementation in C with a distinct focus on interoperability between Java and C.

Celix has been incubating since 2010-11-02.

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

1. Grow the community
2. Find more committers
3. Publish a new 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?

Several new people are working on additional code for the project. One person is working on a GSoC project, someone else is working on a new remote services protocol and is actively communicating over the list and using jira to post patches. Overall this is an increase in activity.

The Native-OSGi project is still progressing, but rather slow. So there isn't anything new on that end. Celix is still the reference implementation, so I still expect this to be a great drive for Celix, for the code as well as the committers.

How has the project developed since the last report?

As mentioned above, a student is working on a GSoC project and has implemented an Event Admin which will be included in Celix. Also a additional Remote Services implementation based on shared memory is being made.

Work for a new release is still ongoing, the focus is slightly changing, one of the main users has requested a release. Since they don't use Win32, Win32 support can be stripped from the release plan and postponed to a later release.

Date of last release:

16-12-2012

When were the last committers or PMC members elected?

2013-01-07 Roman Shaposhnik joins as a mentor.
2012-03-20 Pepijn elected as committer.

Signed-off-by:

[x](celix) Marcel Offermans
[](celix) Karl Pauls

Shepherd notes:

Chukwa

Chukwa is a log collection and analysis framework for Apache Hadoop clusters.

Chukwa has been incubating since 2010-07-14.

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

1. Update DNS entry for chukwa.apache.org
2. Update mailing list location
3. Update chukwa source and website to the new url.

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

Chukwa has been voted to graduate to TLP in September in general@incubator. This will be the last incubator report.

How has the community developed since the last report?

Community has been stable and adding new features to Chukwa.

How has the project developed since the last report?

New heatmap feature and updating to work with being developed. New demux parser improvement for configuring reducer parser has been developed.

Date of last release:

2012-01-26

When were the last committers or PMC members elected?

2012-09-06

Signed-off-by:

[X](chukwa) Alan Cabrera
[](chukwa) Anthony Elder
[](chukwa) Bernd Fondermann
[X](chukwa) Jukka Zitting

Shepherd notes:

DeviceMap

Apache DeviceMap is a data repository containing device information, images and other relevant information for all sorts of mobile devices, e.g. smartphones and tablets. While the focus is initially on that data, APIs will also be created to use and manage it.

Shepherd notes:

Dave Fisher (wave):

I see good activity recently. It looks like it took DeviceMap quite awhile to get started, but the community has grown and it looks like they are working towards a release. Mentors are active. I'm not sure why they have not reported. (Oct 6)

Marvin Humphrey (marvin):

Podling did not report, expected next month.

Helix

Apache Helix is a generic cluster management framework used to build distributed systems and provides automatic partition management, fault tolerance and elasticity.

Helix has been incubating since 2012-10-14.

We feel that we are ready for graduation and have started the discussion internally. We need to update the status page with new committers and also resolve trade mark. Our processes are continuing to improve and adoption is increasing both inside and outside of LinkedIn.

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?

- Added Kanak Biscuitwala as a committer
- Receiving external contribution in developing JavaScript bindings for

Helix (see <https://github.com/sudowork/node-helix-participant>)

- Steady activity on the mailing list

How has the project developed since the last report?

- Making progress on 0.7.0, a release that examines and rethinks every interaction that Helix exposes to increase ease of use. We hope to release this month.
- Introduced new coding style guidelines, and simplified the code review process with new scripts
- Increased documentation of the development process on JIRA pages
- Actively populating the developer wiki and increased the amount of documentation on the project website
- Received a patch integrating Helix with YARN, to automate management of YARN containers.

Date of last release:

2013-05-30 (0.6.1-incubating)

When were the last committers or PMC members elected?

- Kanak Biscuitwala (committer) - 2013-09-20

Signed-off-by:

[X](helix) Patrick Hunt
[X](helix) Olivier Lamy
[](helix) Mahadev Konar
[](helix) Owen O'Malley

Shepherd notes:

Roman Shaposhnik (rvs):

A very solid project on the clear path to graduation.

jclouds

A cloud agnostic library that enables developers to access a variety of supported cloud providers using one API

jclouds has been incubating since 2013-04-29.

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

The jclouds IPMC graduation vote passed on Sep 23rd, 2013.

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

The jclouds community is looking forward to the next ASF Board meeting and hopes to be confirmed as a TLP.

How has the community developed since the last report?

jclouds held its first committer vote as an Apache project and is proud to have Zack Shoylev on board!

How has the project developed since the last report?

Our second incubating release was successful, with no issues related to ASF process or guidelines raised. The remaining trademark and IP issues have been resolved, and changes to the project website to bring it in line with ASF guidelines successfully made. The project also held a successful vote to adopt a set of Bylaws.

With those items completed and a new committer on board, the jclouds PPMC proceeded to a graduation vote, which was passed and confirmed by the

IMPC. Andrew Bayer was elected as the first project PMC chair and will thus be proposed as Vice President, Apache jclouds.

Date of last release:

2013-08-28 (1.6.2-incubating)

When were the last committers or PMC members elected?

2013-08-08 (Zack Shoylev as committer)

Signed-off-by:

[](jclouds) Brian McCallister
[](jclouds) Tom White
[](jclouds) Henning Schmiedehausen
[X](jclouds) David Nalley
[](jclouds) Jean-Baptiste Onofré
[](jclouds) Mohammad Nour El-Din
[X](jclouds) Olivier Lamy
[](jclouds) Tomaz Muraus
[](jclouds) Suresh Maru
[](jclouds) Carlos Sanchez

Shepherd notes:

Marmotta

An Open Platform for Linked Data.

Marmotta has been incubating since 2012-12-03.

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

1. Check whether Marmotta is a suitable name
(<https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-39>)
2. PPMC & IPMC discussion and vote about graduation
3. Graduate tasks

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

No such issues.

How has the community developed since the last report?

- * The NOTICE&LICENSE issues from last time are now properly solved.
- * The 3.1.0-incubating release was rescheduled to "after summer vacation time" and prepared in late September (released October 3rd).
- * Traffic on dev@marmotta.i.a.o was decreasing over the summer months, but is rising again in September during the preparation of the last release.
- * We are starting to get bug reports from non-committers in JIRA while traffic on users@marmotta.i.a.o is slowly evolving.

How has the project developed since the last report?

- * The 3.1.0-incubating release has been accepted on October 3rd.
- * The project was explicitly recommended for graduation after the last report in July (<<http://s.apache.org/LIF8>>) but we wanted to have another release before graduation to stabilize the release process and have another check of the N&L by the incubator community.
- * Consequently the discussion graduation is being resumed after the release, trying to do it by the end of October if no new issues appear.

Date of last release:

2013-10-03 (3.1.0-incubating)

When were the last committers or PMC members elected?

2013-05-21 (Raffaele Palmieri, committer&PPMC)

2013-06-24 (Peter Ansell, committer&PPMC)

So no new committers and/or PPMC member since the last report.

Signed-off-by:

[X](marmotta) Fabian Christ
[](marmotta) Nandana Mihindukulasooriya
[X](marmotta) Andy Seaborne

Shepherd notes:

Andrei Savu (asavu):

Project is on track for graduation. Nothing jumped out at me as requiring any special attention. Good luck guys!

MetaModel

MetaModel is a data access framework, providing a common interface for exploration and querying of different types of datastores.

MetaModel has been incubating since 2013-06-12.

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

1. More contributors to increase diversity of the community
2. Make first release under ASF umbrella
3. Looking forward to setup CI builds under ASF infra

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?

Ankit Kumar has taken the role of being release engineer for our first ASF release.

A draft for bylaws for the project has been added to the wiki, we are discussing and finalizing the details.

We are still not seeing a lot of involvement from new contributors, but expect that a release will give us more users and thus involvement from more newcomers.

How has the project developed since the last report?

The primary focus has been on making our first release under ASF umbrella. The work is coming along well and we have a practical TODO list of less than 5 points now.

Date of last release:

None

When were the last committers or PMC members elected?

June 2013

Signed-off-by:

[X](metamodel) Henry Saputra

[X](metamodel) Arvind Prabhakar
[X](metamodel) Matt Franklin
[](metamodel) Noah Slater

Shepherd notes:

ODF Toolkit

Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents.

ODF Toolkit has been incubating since 2011-08-01.

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

1. Grow the community, especially attracting new developers.
2. Address issues of too much code development happening outside of Apache
3. Hold an informal graduation readiness vote

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

It has recently been pointed out that the project status page on the Incubator project site is out of date, we hope to update that shortly.

How has the community developed since the last report?

Higher feed-back from the users mailing list, but no new developers.

Discussion has started, but without much conclusion, of whether we should aim for TLP status, or merge into an existing TLP with similar alignment, such as POI or OpenOffice.

How has the project developed since the last report?

Clear unnecessary dependency from the Clarizza project.

Date of last release:

2013-06-22

When were the last committers or PMC members elected?

Not for a long while...

Signed-off-by:

[](odftoolkit) Sam Ruby
[X](odftoolkit) Nick Burch
[](odftoolkit) Yegor Kozlov

Shepherd notes:

Marvin Humphrey (marvin):

Report not filed in time for shepherd review.

Olingo

Apache Olingo provides libraries which enable developers to implement OData producers and OData consumers. While starting with an initial code base implementing OData version 2.0 it is also a clear goal to start implementing a library for OData 4.0 once the OData standard is published at OASIS. The focus within the community is currently on the Java technology but it is up to the community to discuss if other environments find interest.

Olingo has been incubating since 2013-07-08.

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

1. Make the first release
2. Work on documentation
3. Grow community

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?

- Committers actively working on making the release. Effort is coordinated by Jira and mailing list.
- Committers are in mailing list contact with mentors and infrastructure to extend infrastructure setup (e.g. Jenkins, KEYS files ...).
- Some committers have started with KEY signing. This is in progress to get all committers involved.

How has the project developed since the last report?

- There is currently no feature development
- Instead committers do concentrate on making the first release
- A release candidate is available and discussed by community on mailing list
- Release voting has started
- Documentation on web site is growing and will be published as soon as the first release is available

Date of last release:

- not released so far

When were the last committers or PMC members elected?

- n/a

Signed-off-by:

[X](olingo) Alan Cabrera
[X](olingo) Dave Fisher
[X](olingo) Florian Müller

Shepherd notes:

Raphael Bircher (rbircher):

Apache Olingo goes direction first release. They already use mailing list in a good way. It looks like they already know how to work with lazy consensus. The Mentors are active in a good way. Well, the documentation is poor at the moment. But the project knows this problem and work on it. Looking forward for a release.

Ripple

Ripple is a browser based mobile phone emulator designed to aid in the development of HTML5 based mobile applications. Ripple is a cross platform and cross runtime testing/debugging tool. It currently supports such runtimes as Cordova, WebWorks and the Mobile Web.

Shepherd notes:

Marvin Humphrey (marvin):

Podling did not report, expected next month.

Samza

Samza is a stream processing system for running continuous computation on infinite streams of data.

Samza has been incubating since 2013-07-30.

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

1. Bring in new community members
2. Make a release
3. Cement ASF way to 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?

Community development continues. Chris Riccomini presented Samza to YARN Dev User Group with great interest. Similar talks planned for Bay Area and London Hadoop User Groups this month. Interest (and patches) continue to begin to appear.

How has the project developed since the last report?

Several JIRAs opened and committed. Ready to start working on first release.

Date of last release:

None.

When were the last committers or PMC members elected?

None.

Signed-off-by:

[X](samza) Chris Douglas
[](samza) Arun Murthy
[X](samza) Roman Shaposhnik

Shepherd notes:

Sentry

Sentry is a highly modular system for providing fine grained role based authorization to both data and metadata stored on an Apache Hadoop cluster.

Sentry has been incubating since 2013-08-08.

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

1. Build community and add new committers
2. Continue to release at regular intervals

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?

Sentry User Meetup group has nearly 50 members. The project's first meetup is planned for the end of Oct. We have 23 sign ups on the dev list.

How has the project developed since the last report?

We have migrated all of the infrastructure to ASF. The website is up and running. Podling name search has been successfully completed. Project had its first release since starting incubation. On the development front, a number of new JIRAs have been filed and are being actively worked on.

Date of last release:

2013-09-25 (1.2.0-incubating)

When were the last committers or PMC members elected?

Last committers and PMC members were elected when the project entered incubation.

Signed-off-by:

[X](sentry) Arvind Prabhakar
[X](sentry) Joe Brockmeier
[X](sentry) David Nalley
[](sentry) Olivier Lamy
[X](sentry) Patrick Hunt
[](sentry) Thomas White

Shepherd notes:

Matt Franklin (mfranklin):

No issues. Looks like the project is off to a great start.

Spark

Spark is an open source system for fast and flexible large-scale data analysis. Spark provides a general purpose runtime that supports low-latency execution in several forms.

Spark has been incubating since 2013-06-19.

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

1. Move JIRA over to Apache (still haven't gotten success from INFRA on this: <https://issues.apache.org/jira/browse/INFRA-6419>)
2. Add more committers under Apache process
3. Make further Apache releases

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

We still need some help importing our JIRA -- see INFRA-6419. For some reason we've had a lot of trouble with this. It should be easier now because Apache's JIRA was updated and now matches our version.

How has the community developed since the last report?

We made the Spark 0.8.0 release, which was the biggest so far, with 67 developers from 24 organizations contributing. The release shows how far our community has grown -- our 0.6 release last October had only 17 contributors, and our 0.7 release in February had 31. Most of the contributors are now external to the original UC Berkeley team.

How has the project developed since the last report?

We made the Spark 0.8.0 release mentioned above, which so far seems to be doing well. It brings a number of deployability features, improved Python support, and a new standard library for machine learning; see <http://spark.incubator.apache.org/releases/spark-release-0-8-0.html>

for what's new in the release.

Date of last release:

2013-09-25

When were the last committers or PMC members elected?

June 2013

Signed-off-by:

[X](spark) Chris Mattmann
[](spark) Paul Ramirez
[](spark) Andrew Hart
[](spark) Thomas Dudziak
[](spark) Suresh Marru
[X](spark) Henry Saputra
[X](spark) Roman Shaposhnik

Shepherd notes:

Dave Fisher (wave):

Very active community on a fast track. Good report. Get your JIRA over and you are getting close. (Oct. 7)

Marvin Humphrey (marvin):

Report not filed in time for shepherd review.

Storm

Storm is a distributed, fault-tolerant, and high-performance realtime computation system that provides strong guarantees on the processing of data.

Storm has been incubating since 2013-09-18.

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

1. Transition community to new infrastructure and mailing lists
2. Test and mature the Netty message passing implementation so all future releases do not include Omg transport (which is LGPL)
3. Make a release under Apache

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

We are still waiting for ICLA's from Flip Kromer and David Lao.

We are waiting for infrastructure to be set up so we can switch mailing lists and code hosting over to Apache.

How has the community developed since the last report?

A few more users have been added to the Powered By page. Otherwise, we've received a steady amount of new patches and messages on the mailing list.

How has the project developed since the last report?

A software grant has been filed by Nathan Marz for the project. Also, all contributor agreements Nathan has collected pre-Apache have been submitted to Apache. All committers except for two have filed ICLA's and CCLA's.

We are making one last release via the old route during the transition into the incubator (0.9.0). We have switched development to an Apache-like consensus-driven approach and have made lots of progress on closing pull requests (as well as getting committers active in

discussions, merging, and managing release process).

Date of last release:

None

When were the last committers or PMC members elected?

N/A, this is the first status report.

Signed-off-by:

[X](storm) Ted Dunning
[X](storm) Arvind Prabhakar
[X](storm) Devaraj Das
[X](storm) Matt Franklin
[X](storm) Benjamin Hindman

Shepherd notes:

Stratos

Stratos will be a polyglot PaaS framework, providing developers a cloud-based environment for developing, testing, and running scalable applications, and IT providers high utilization rates, automated resource management, and platform- wide insight including monitoring and billing.

Stratos has been incubating since 2013-06-20.

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

1. Moving all Stratos related trademarks from WS02 to Apache. Certain WS02 technologies such as StratosLive still bear the Stratos name. These have to be resolved in the near future. We are in the process of sorting these out.
2. Get the first Apache Stratos (incubating) release out, and in the process ensure all ASF release requirements are met.
3. Build 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 committers & PMC members have been added to the project. (In Stratos, committers are also invited to join the PMC). Six Google Hangouts were held in order to educate the community about the project & the technical details/architecture in the project. A Cloud BarCamp is being organized for end of October to enable the Stratos community members as well as other Apache cloud project community members to meet up & exchange ideas & build synergy. Community members have also written several blog posts. The project has also seen a number of new contributors who have started actively participating in the project.

How has the project developed since the last report?

Code has been refactored to get rid of non ASL 2.0 compatible third party library dependencies. Package names have been changed. License files added. Several bugs have been fixed. Builds have been made available for review & testing. In anticipation of the first release, the project has made available several RC packs to get community feedback. Documentation has been added to the wiki. The project site has been properly setup. The roadmap discussions are currently taking place on the dev list.

Date of last release:

- none -

When were the last committers or PMC members elected?

2013-09-23

Signed-off-by:

[x](stratos) Afkham Azeez
[](stratos) Anthony Elder
[x](stratos) Chip Childers
[](stratos) Marlon Pierce
[](stratos) Mohammad Nour
[](stratos) Noah Slater

Shepherd notes:

Roman Shaposhnik (rvs):

Stratos appears in a solid shape and working on its first incubator release. That said, their report is late which makes it difficult for the shepherding work to feel complete. I will try to ping some of the mentors.

Marvin Humphrey (marvin):

Report not filed in time for shepherd review.

VXQuery

A standards compliant parallel XML Query processor.

VXQuery has been incubating since 2009-07-06.

The most important issue to address in the move towards graduation is still the first release. The third release candidate is currently up for vote on vxquery-dev. So far there at 4 positive votes.

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

We are still a little low on mentors, but since Ant Elder has agreed to help out, the situation has improved.

How has the community developed since the last report?

Steven Jacobs has been voted in as a committer.

How has the project developed since the last report?

Benchmarking data and queries have been added to the project. A number of issues discovered when running the benchmark have been fixed.

Steven contributed an integration with Lucene.

A new release candidate has been created and is currently being voted on.

Date of last release:

never

When were the last committers or PMC members elected?

2013-09-19

Signed-off-by:

[](vxquery) Sanjiva Weerawarana
[](vxquery) Radu Preotiuc-Pietro
[x](vxquery) Jochen Wiedmann

[](vxquery) Ant Elder

Shepherd notes: