

April2013

- REPORT HAS BEEN SUBMITTED - DO NOT EDIT ANYMORE*

Shepherds

Podling	Shepherd
Celix	
Chukwa	
Curator	Suresh Marru
DeviceMap	
Falcon	Arun C. Murthy
Helix	Suresh Marru
JSPWiki	
Knox	
Marmotta	Dave Fisher
Mesos	
MRQL	
ODF Toolkit	
Open Climate Workbench	
Provisionr	Tom White
Ripple	Dave Fisher
Tajo	Hyunsik Choi
Tashi	
Tez	Arun C. Murthy
VXQuery	Dave Fisher

- REPORT HAS BEEN SUBMITTED - DO NOT EDIT ANYMORE*

Incubator PMC report for Apr 2013

April was a month of much email at the IPMC. Two different non-Foundation-member candidates for IPMC membership led to controversy (each in their own way), and exposed the difficulties of achieving consensus in a large, diverse, group. Eventually, the two cases were resolved, and the community proceeded to a public conversation about decision-making process, agreeing to use a 3/4-majority voting rule when considering non-Foundation-member candidates.

The voluminous and frustrating email on this subject, combined with the ongoing questions around supervision (are enough mentors paying attention?) led to a renewed debate on the structure, and indeed, existence, of the IPMC as currently constituted. This debate continues.

o Community

New IPMC members:

Henry Saputra
David Nalley

People who left the IPMC:

Felix Meschberger
Nicolas Lalevée
Stefan Bodewig
Jeremias Maerki

o New Podlings

Curator

o Graduations

The board has motions for the following:

Apache Onami

o Releases

o Legal / Trademarks

o Infrastructure

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

Still getting started at the Incubator

These projects are still getting started, so no immediate progress towards graduation is yet expected.

Not yet ready to graduate

Marmotta

Not yet ready to graduate, needs attention

VXQuery

Ready to graduate

Table of Contents

Celix
Chukwa
Curator
DeviceMap
Falcon
Hadoop Development Tools
Helix
JSPWiki
Knox
Marmotta
Mesos
MRQL
ODF Toolkit
Open Climate Workbench
Provisionr
Ripple
Tajo
Tashi
Tez
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 a community
2. Find more committers
3. ..

Celix is maintained by 2 committers and a small group of users. Questions to the mailing list are (most of the

time) answered quickly. We are also trying to give talks on Celix or related stuff on conferences. We see and recognise the problem of the small community, but are also confident that Celix and its community will grow.

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 is no active growth, but there is currently one student doing his internship with Apache Celix and we are trying hoping to get some projects in the Google Summer of Code 2013. We hope that this will lead in more activity on the mailing-list and hopefully to extra committers.

How has the project developed since the last report?

Technically progress has been slow the last months. There was a hackaton organized for remote services in Java [1], which although not directly related to Apache Celix, should eventually work together with the Apache Celix and as result the committers of Apache Celix were present. We also expect an increase in contributions the coming months because of the student internship / GSoC activities.

[1] <http://lists.apache.org/pipermail/dev/2013-March/000020.html>

Please check this [] when you have filled in the report for Celix.

Signed-off-by:

Marcel Offermans: [](celix)

Karl Pauls: [](celix)

Luciano Resende: [](celix)

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. Chukwa wiki is currently hosted in Hadoop subproject, it may be good to move it fully independent of Hadoop wiki.

2. As of late, the overall activity has declined. The podling needs to grow its active community.

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

* The podling is still working on increasing community activity. It's hoped that the addition of three new committers will help.

How has the community developed since the last report?

* 3 Committers have been added.

How has the project developed since the last report?

* The project has a few patches and discussion on improving Chukwa architecture design. The main focus is removing Chukwa collectors to reduce deployment complexity.

Please check this [] when you have filled in the report for Chukwa.

Signed-off-by:

Jukka Zitting: [x] (chukwa)

ant: [x] (chukwa)

Alan Cabrera [X](chukwa)

Bernd Fondermann: [](chukwa)

Shepherd notes:

Curator

Curator - ZooKeeper client wrapper and rich ZooKeeper framework

Curator has been incubating since 2013-03-11.

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

1. build community
2. create a first release
3. finalize name search/trademarks

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?

This is the first report.

There are currently 12 subscribers on the dev list. There are 13 subscribers to the user list.

How has the project developed since the last report?

Initial resources have been created (Mailing lists, SCCS, Jira, Jenkins, etc...). All mentors are signed up for the mailing lists. All committers have accounts and are also registered for the ML.

Mentors have been educating the initial committers/community on the workings of Apache.

A first release candidate for a 2.0.0 release has been created and feedback provided by mentors. The community is iterating on new RCs.

Please check this [X] when you have filled in the report for Curator.

Signed-off-by:

Enis Söztutar: [](curator)

Luciano Resende: [x](curator)

Mahadev Konar: [](curator)

Patrick Hunt: [X](curator)

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.

DeviceMap has been incubating since 2012-01-03.

It looks like DeviceMap needs to be rebooted - there's almost zero activity from the initial committers, but a few people have proposed contributions that could help restart with a mostly new set of committers (without kicking existing ones out of course - they're welcome to be more active if they're willing and able).

Two committer/PPMC elections are pending, and with two active mentors this can provide the embryo of a new community.

The podling could be considered a failure at this point, but giving it another chance does no harm, so we recommend trying that.

Please check this [X] when you have filled in the report for DeviceMap.

Signed-off-by:

Bertrand Delacretaz: [X](devicemap)

Kevan Miller: [](devicemap)

Andrew Savory: [X](devicemap)

Shepherd notes:

Falcon

Falcon is a data processing and management solution for Hadoop designed for data motion, coordination of data pipelines, lifecycle management, and data discovery. Falcon enables end consumers to quickly onboard their data and its associated processing and management tasks on Hadoop clusters.

Falcon has been incubating since 2013-03-27.

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

1. Bootstrap the project (source code, documentation, issue tracking)
2. Add new and diverse committers
3. Build and grow community
4. Releases at frequent and regular intervals

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?

This is the first report for Falcon since its acceptance into Incubator. We have requested for infrastructure for bootstrapping the podling (awaiting source control, mailing list and issue tracking) - (<https://issues.apache.org/jira/browse/INFRA-6115>). We will look to move the source code and make documentation and easy startup-guide available for users to try it out and collaborate.

How has the project developed since the last report?

N/A. This is the first report being filed for the project.

Please check this [] when you have filled in the report for Falcon.

Signed-off-by:

Arun Murthy: [X](falcon)

Chris Douglas: [X](falcon)

Owen O'Malley: [](falcon)

Devaraj Das: [](falcon)

Alan Gates: [](falcon)

Shepherd notes:

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.

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

1. Add new committers.
2. Grow community.
3. Better documentation to increase 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?

There has been significant increase in the activity since last report resulting in good discussions on the mailing list.

How has the project developed since the last report?

We had our first Apache Incubator release in January and we are in the process of making another release shortly.

Lots of new features and bug fixes added.

- Issues +58 issues created, 35 resolved.
- Additional features to be part of 0.6.2-incubating release
- Standalone Helix agent
- Task execution recipe
- Task scheduling and support for custom re-balancer.
- S4 integration with Helix in progress.

Please check this [X] when you have filled in the report for Helix.

Signed-off-by:

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

Shepherd notes:

JSPWiki

Java-based wiki engine

JSPWiki has been incubating since 2007-09-17.

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

1. resolution acceptance
2. community vote
3. IPMC graduation vote

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

If the release vote of 2.9.1 is succesful, we're thinking on
discussing leaving the Incubator + graduate as TLP

How has the community developed since the last report?

Compared to the last period, the traffic on jspwiki-dev has been halved,
whereas the traffic at jspwiki-user has increased slightly. 1 patch has
been incorporated into trunk

How has the project developed since the last report?

15 JIRAs resolved, ongoing vote to release Apache JSPWiki (incubating) 2.9.1

Please check this [] when you have filled in the report for JSPWiki.

Signed-off-by:

Dave Johnson: [](jspwiki)
Craig Russell: [](jspwiki)
Henning Schmiedehausen: [](jspwiki)
Sam Ruby: [](jspwiki)
Siegfried Goeschl: [X](jspwiki)

Shepherd notes:

Knox

Knox Gateway is a system that provides a single point of secure access
for Apache Hadoop clusters.

Knox has been incubating since 2013-02-22.

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

1. Complete first release and refine/document release processes.
2. Expand community to include more diverse committers.
3. Clear the project name with legal and pick a new name if required.

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

1. No.

How has the community developed since the last report?

1. No new members of the community have shown up or been added.

How has the project developed since the last report?

1. Now have all infrastructure in in place.
2. Significant increase in traffic on dev mailing list.
3. Significant improvement to website presence.

4. Internal audit of dependencies and creation of LICENSE and NOTICE files.

5. Preparing for first release for 0.2.0-incubating.

Please check this [] when you have filled in the report for Knox.

Signed-off-by:

Owen O'Malley: [] (knox)

Chris Douglas: [X] (knox)

Mahadev Konar: [] (knox)

Alan Gates: [X] (knox)

Devaraj Das: [X] (knox)

Chris Mattmann: [X] (knox)

Tom White: [X] (knox)

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. Ingest code and clear IP.
2. A release
3. Build dev and PMC.

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 project has been completely setup in the ASF infrastructure, and community starts to evolve healthily (i.e., the monthly evolution of traffic at the dev mailing list is 77, 274, 292, 490 messages, starting in December).

* After intensively discussing and reviewing 5 internal release candidates, now, besides little legal details, the first release will be probably submitted to the IPMC in the upcoming days.

How has the project developed since the last report?

After the publication of LMF 2.6.0, the base code had been imported on February 19th. Since that, the team has been working hard to implement all necessary refactoring and, in parallel, some bug fixing; which in total means almost 200 issues already registered at Jira (70% resolved, the rest mostly planned tasks/features).

Please check this [X] when you have filled in the report for Marmotta.

Signed-off-by:

Fabian Christ: [X](marmotta)

Nandana Mihindukulasooriya: [X](marmotta)

Andy Seaborne: [X](marmotta)

Shepherd notes:

Dave Fisher (wave): This project is doing well and looks really cool.

Mesos

Mesos is a cluster manager that provides resource sharing and isolation across cluster applications.

Mesos has been incubating since 2010-12-23.

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?

How has the project developed since the last report?

Please check this [] when you have filled in the report for Mesos.

Signed-off-by:

Tom White: [](mesos)

Shepherd notes:

MRQL

MRQL is a query processing and optimization system for large-scale, distributed data analysis, built on top of Apache Hadoop and Hama.

MRQL has been incubating since 2013-03-13.

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

1. Expand the development community to include more diverse committers
2. Complete the first release
3. Create a wiki to document both the software use and the software development process

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?

This is the first report for MRQL. We have decided to adopt the review-then-commit policy that requires at least two +1 from committers and no -1 from committer. In addition, we have started the process of creating a new infrastructure for the MRQL podling. We have already created three email lists, mrql-private, mrql-dev, and mrql-user, and we have requested the creation of a new JIRA project and a new wiki site for MRQL.

How has the project developed since the last report?

We have restructured the original source code and we have modified the in-file copyright notices of the source code to be compliant with the Apache License 2.0, to make it ready for the first release.

Please check this [] when you have filled in the report for MRQL.

Signed-off-by:

Alex Karasulu: [X](mrql)

Anthony Elder: [](mrql)

Alan Cabrera: [](mrql)

Mohammad Nour: [X](mrql)

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.
- 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?

How has the project developed since the last report?

Please check this ☐ when you have filled in the report for ODF Toolkit.

Signed-off-by:

Sam Ruby: ☐ (odftoolkit)

Nick Burch: ☐ (odftoolkit)

Yegor Kozlov: ☐ (odftoolkit)

Shepherd notes:

Open Climate Workbench

Apache Open Climate Workbench (Incubating) is an effort to develop software that performs climate model evaluation using model outputs from a variety of different sources (the Earth System Grid Federation, the Coordinated Regional Downscaling Experiment, the U.S. National Climate Assessment and the North American Regional Climate Change Assessment Program) and temporal/spatial scales with remote sensing data from NASA, NOAA and other agencies. The toolkit includes capabilities for regridding, metrics computation and visualization.

Open Climate Workbench has been incubating since 2013-02-15.

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

1. Develop an Apache community for Open Climate Workbench and connect to other relevant Apache efforts (Tika, Hadoop, SIS, OODT)
2. Make an initial release.
3. Add new contributors to the project.

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?

We're up to 79% (22) of the initial project members having ICLAs on file and accounts. We've got 21% more (6) people still who have yet to submit their ICLA.

Chris Douglas and Chris Mattmann have reached out to Chris Jack and Bruce Hewitson and are working on getting those 2 ICLAs so we can process their accounts.

How has the project developed since the last report?

INFRA-5874 is pretty much done, the big thing we're waiting on now is the SVN import of the RCMES codebase. The SGA is on file as of Thursday March 14, 2013, so we're ready to go there. Infra needed a checksum file and evidence that we have an SGA to do the import, so we provided that information on INFRA-5966, so should be loaded into SVN any day now.

There was some limited discussion about others with the Climate name in their software -- after discussion with Shane, this was deemed not a problem.

Once we get the code base into SVN, all of the RCMES folks who have been working on the project at JPL will conduct all development at Apache.

Please check this ☒ when you have filled in the report for Open Climate Workbench.

Signed-off-by:

Chris Mattmann: ☒ (openclimateworkbench)

Suresh Maru: ☒ (openclimateworkbench)

Chris Douglas: ☐ (openclimateworkbench)

Nick Kew: ☐ (openclimateworkbench)

Shepherd notes:

Provisionr

Provisionr provides a service to manage pools of virtual machines on multiple clouds.

Provisionr has been incubating since 2013-03-07.

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

1. release 0.4.0-incubating compliant with the ASF policies
2. develop a community
3. improve support for golden images & cloudstack

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?

This is the first report. Andrei Savu presented the project at the Bucharest JUG meetup (<http://www.slideshare.net/savu.andrei/apache-provisionr-incubating-bucharest-jug-10>) and we've got our first contribution from Ioannis Canellos.

How has the project developed since the last report?

We are still working on ASF compliance - mandatory for our first release.

Please check this [] when you have filled in the report for Provisionr.

Signed-off-by:

Roman Shaposhnik: [](provisionr)

Tom White: [X](provisionr)

Mohammad Nour: [X](provisionr)

Shepherd notes: Provisionr has been incubating for about a month and has made good progress on the initial set up tasks.

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.

Ripple has been incubating since 2012-10-16.

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

1. Check if some source modules are not compatible with ASF guidelines and need to be extracted to f. e. GitHub
2. Check license of images and - if necessary - replace them
3. Prepare 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?

No new contributors/committers.

Due to changes in the working life of some of the original developers there are no more full-time committers working at the Ripple project. The committers expressed there interest in contributing further in their spare time.

How has the project developed since the last report?

Less activity

Please check this [X] when you have filled in the report for Ripple.

Signed-off-by:

Jukka Zitting: [](ripple)

Christian Grobmeier: [X](ripple)

Andrew Savory: [](ripple)

Shepherd notes:

Tajo

Tajo is a distributed data warehouse system for Hadoop.

Tajo has been incubating since 2013-03-07.

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

1. Make a release
2. Attract users and contributors
3. Foster more and diverse committers

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?

* Project infrastructures (i.e., site, jira, wiki, mailing list and git repository) setup has been completed, and the site and the source code are moved to the Apache infrastructure.

* All members have submitted necessary CLAs.

* There have been no new committers or PPMC members elected.

How has the project developed since the last report?

* We've applied ALv2 header to all files and adopted maven-rat-plugin to audit them.

* We've made and improved developer/user documents, including GettingStarted, Build Instruction, and HowToContribute.

* We've resolved 11 jira issues, including bug fixes and improvements.

* Discussion is underway to move current site-based managed documentation to the wiki.

Please check this [] when you have filled in the report for Tajo.

Signed-off-by:

Chris Mattmann: [X](tajo)

Owen O'Malley: [](tajo)

Alex Karasulu: [X](tajo)

Shepherd notes:

Tashi

An infrastructure for cloud computing on big data.

Tashi has been incubating since 2008-09-04.

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?

How has the project developed since the last report?

Please check this [] when you have filled in the report for Tashi.

Signed-off-by:

Matthieu Riou: [](tashi)

Craig Russell: [](tashi)

Shepherd notes:

Tez

Tez is a framework for processing arbitrarily complex data-processing tasks.

Tez has been incubating since 2013-02-24.

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

1. Develop collaborations with other Apache projects, including Hadoop, YARN
2. Make an initial Tez release.
3. Grow the Apache Tez community.

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?

The community is still bootstrapping. Chris started a thread on dev@tez to request clarification on who needs ICLAs filed, and who needs accounts set up. Chris can help with this as can the other mentors. Jira is ready, we are beginning to bootstrap work.

How has the project developed since the last report?

A few jiras have been created and some patches have been filed.

Mailing lists for dev, private, and commits are up and running.

Please check this [X] when you have filled in the report for Tez.

Signed-off-by:

Alan Gates: [](tez)

Arun C Murthy: [X](tez)

Chris Douglas: [](tez)

Chris Mattmann: [X] (tez)

Jakob Homan: [] (tez)

Owen O'Malley: [] (tez)

Shepherd notes:

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. Two release candidates have been created and rejected so far. Work on the third release candidate is ongoing.

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

For RC1 only 2 mentors (Jochen and Paul) voted, for RC2 only Jochen voted.

Getting a 3 votes from the IPMC might be difficult.

How has the community developed since the last report?

no changes

How has the project developed since the last report?

There has been progress on the rewriting of queries. A larger class of queries is now correctly translated into plans that can be executed in parallel on a

Hyracks cluster.

Please check this [x] when you have filled in the report for VXQuery.

Signed-off-by:

Sanjiva Weerawarana: [](vxquery)

Radu Preotiuc-Pietro: [](vxquery)

Jochen Wiedmann: [x](vxquery)

Shepherd notes:

Dave Fisher (wave): There is a current attempt to pass a VOTE on a release. This podling could use more mentors. BUT I am rather unsure if there is a viable community.