

April2014

/!\ *FINAL* (!)

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

Incubator PMC report for April 2014

Timeline

Wed April 02	Podling reports due by end of day
Sun April 06	Shepherd reviews due by end of day
Sun April 06	Summary due by end of day
Tue April 08	Mentor signoff due by end of day
Wed April 09	Report submitted to Board
Wed April 16	Board meeting

Shepherd Assignments

Alan D. Cabrera	Usergrid
Alan D. Cabrera	VXQuery
Andrei Savu	ODF Toolkit
Andrei Savu	Ripple
Dave Fisher	NPanday
John Ament	Phoenix
Justin Mclean	DataFu
Matthew Franklin	Samza
Raphael Bircher	Stratos
Roman Shaposhnik	MetaModel
Ross Gardler	DeviceMap
Suresh Marru	Aurora
Suresh Marru	Celix

Report content

Incubator PMC report for April 2014

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 32 podlings currently under incubation.

* Community

New IPMC members:

Rob Weir
Sean Owen

People who left the IPMC:

(None)

* New Podlings

(None)

* Graduations

The board has motions for the following:

(None)

The following graduations have been approved by the previous board meeting but haven't been recorded in the incubator report:

Tajo
Olingo
Allura

* Releases

The following releases were made since the last Incubator report:

March 13 Phoenix 2.2.3-incubating

It took 2-10 days for the third IPMC vote to arrive.

Release	VOTE start	Third IMPC +1	Days
-----	-----	-----	-----
Apache Phoenix 2.2.3-incubating	Mar 7	Mar 13	6

* Miscellaneous

- * A proposal for Stratosphere: an open source system for parallel data analysis has been submitted
- * A proposal for Slider has superseded the former proposal for Hoya
- * A POC with Black Duck software and a few incubator projects has been submitted and actively discussed

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

* Not yet ready to graduate

No release:

NPanday
ODF Toolkit
Samza

Community growth:

NPanday
ODF Toolkit

* Did not report, expected next month

Celix
DeviceMap
Ripple

Table of Contents

Aurora
DataFu
MetaModel
NPanday
ODF Toolkit
Phoenix
Samza
Stratos
Usergrid

VXQuery

Aurora

Aurora is a service scheduler used to schedule jobs onto Apache Mesos.

Aurora has been incubating since 2013-10-01.

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

1. First release, progress being tracked in ticket AURORA-147
2. Establish a consistent development process and release rhythm.
3. Expanding the community and continuing to add new committers.

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?

Latest Additions:

- * PMC addition: Mark Chu-Carroll, 2014-02-03
Jake Farrell, 2014-01-23
- * Contributor addition: Mark Chu-Carroll, 2014-01-14

Issue backlog status since last report:

- * Created: 296 (includes initial import)
- * Resolved: 141

Mailing list activity since last report:

- * @dev 624 messages
- * issues 503 messages
- * @reviews 1499 messages

How has the project developed since the last report?

- Jira issues have been imported to issues.apache.org
- Documented contributing guide and committers workflow
- Automating release workflows and preparing for our first release candidate

Date of last release:

- No releases as of yet. Targeting first release within the next couple weeks.

Signed-off-by:

- [X](aurora) Jake Farrell
- [](aurora) Benjamin Hindman
- [](aurora) Chris Mattmann
- [X](aurora) Henry Saputra

DataFu

DataFu provides a collection of Hadoop MapReduce jobs and functions in higher level languages based on it to perform data analysis. It provides functions for common statistics tasks (e.g. quantiles, sampling), PageRank, stream sessionization, and set and bag operations. DataFu also provides Hadoop jobs for incremental data processing in MapReduce.

DataFu has been incubating since 2014-01-05.

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

1. Building ASF community
2. Release
3. Remaining incubator paperwork

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

None.

How has the community developed since the last report?

A talk was given at an Apache Pig meetup held on March 14th. A talk is scheduled to be given at ApacheCon in Denver on April 7th. Jian Wang accepted the invitation to become a committer.

How has the project developed since the last report?

Two new Jiras have been filed and received patches.

Date of last release:

None. Third month of incubation.

When were the last committers or PMC members elected?

2014-02-22

Signed-off-by:

[](datafu) Ashutosh Chauhan
[X](datafu) Roman Shaposhnik
[](datafu) Ted Dunning

Shepherd/Mentor notes:

Justin Mclean (jmclean):

Relative new podling yet to make a release. One mentor is active on public mailing list no obvious issues that need attention.

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. Finalize 4.1 release to demonstrate new functionality and evolution within Apache
2. More contributors to increase diversity of the community
3. Determine approach to technical documentation across different software versions

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 finished our first release and are seeing some initial interest from users. Mostly there are interactions to add new features or connectivity to new types of datastores. For instance, it has been proposed to look into building connectivity to semantic web technologies, iRods, Apache Phoenix - and within the contributors group there has been some focus on

adding HBase and Cassandra support.

Discussions about community by-laws are not very active at this point. It does not seem to be top-of-mind for most community members, and also shepherd/mentor notes from January 2014 indicate that this is not important right now.

We're improving on the point of registering JIRA issues and running patches through the review board. By now this is the approach taken every time, unless a commit is found to be very trivial.

Still the core group of developers seems a bit small, but by focusing on broader connectivity and keeping the releases coming, we think we can attract a larger community to the project.

How has the project developed since the last report?

Apache MetaModel 4.0.0-incubating was released on Jan 21.

We are working on new features for a 4.1 version. The 4.1 version will represent the point in time where Apache MetaModel starts to show significant new functionality when compared to the old eobjects.org MetaModel project.

A new module for providing easy Spring framework integration has been implemented.

Finishing touches are being made to a new module for connecting to Apache HBase.

Various API changes are being planned and discussed.

Many bugfixes and performance improvements has been implemented.

Date of last release:

Jan 21, 2014

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
[X](metamodel) Noah Slater

NPanday

NPanday allows projects using the .NET framework to be built with Apache Maven.

NPanday has been incubating since 2010-08-13.

In previous reports, we indicated that there was a regression that blocked releasing the next version. This has now been resolved, and the only thing that stands in the way is sufficient cycles to push it out. There have been a small number of users (not yet committers) who've stepped forward to help test when that happens, which might help build some much needed momentum. However, as before, it's still a concern whether sufficient release votes could be found as many of the existing PPMC are inactive. The last release was made on 16 May 2011.

The steps that have been identified are:

- apply patches that have been contributed recently
- convert the website to svnpubsub so updated documentation can be published
- perform the release steps

The NPanday automated builds have been impacted by changes to the Windows build infrastructure again, but this is probably something that needs to be worked around on our end.

Signed-off-by:

[](npanday) Raphael Bircher

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 in terms of committers
2. Have regular releases
3. Decide if we would like to be a top level project or join an existing project

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

We've been struggling to get 3 +1 IPMC votes for our latest release candidate. We're stuck at two +1 votes.

How has the community developed since the last report?

We issued a "call for volunteers" to our users list and got a good response: <http://s.apache.org/4Vx>

Of course, it will be easier to grow the community if we can reliably conduct a release vote.

How has the project developed since the last report?

We've been trying to get a release out since February 11th.

Date of last release:

2013-06-22

When were the last committers or PMC members elected?

2012-10-29

Signed-off-by:

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

Phoenix

Phoenix is an open source SQL query engine for Apache HBase, a NoSQL data store. It is accessed as a JDBC driver and enables querying and managing HBase tables using SQL.

Phoenix has been incubating since 12/11/2013.

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

- 1) Make our customers successful and prove value of Phoenix+HBase over other alternatives
- 2) Attract new committers and continue to get contributions from existing committers to the project

- 3) Figure out a bundling strategy with HBase so that Phoenix is pre-installed in HBase distributions.

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?

Continued growth of user/dev traffic on email list. James Taylor is giving a talk on Apache Phoenix at ApacheCon.

How has the project developed since the last report?

- * First release has been completed.
- * Vote is in progress on general list for our second and third releases.
- * Work in parallel on next point releases has already begun.

Date of last release:

March 13, 2014

When were the last committers or PMC members elected?

No new committers or PMC members, but a discussion has begun to add a new committer

Signed-off-by:

[](phoenix) Lars Hofhansl
[X](phoenix) Andrew Purtell
[](phoenix) Devaraj Das
[](phoenix) Enis Soztutar
[](phoenix) Steven Noels

Shepherd/Mentor notes:

John Ament (johndament):

Phoenix is a newer podling, however they are growing quite fast. While there were some release issues, those were resolved quite quickly on the general list and new releases cut quickly. Based on this I would strongly recommend that the podling consider graduation sooner than later. Most of their status report is complete as well.

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. Release a version of Samza.
2. Pull in at least one more committer.
3. Building expertise in the community.

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?

Community development has been strong. Our mailing list saw an all-time high in March of 275 emails. Collectively, March and February had more emails than all prior months.

How has the project developed since the last report?

The project is developing well. We're trying to stabilize the codebase, and have our first release. JIRA information:

Period	Created	Resolved
January 2014	24	18
February 2014	27	26
March 2014	51	38

Date of last release:

Currently in the process of beginning first release.

When were the last committers or PMC members elected?

March 17th, 2014

Signed-off-by:

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

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. Resolve name search issue
<https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-45>
2. Community diversity
3. Start Graduation vote

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

The Stratos PPMC would be glad if the name search issue could be resolved soon. <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-45>

How has the community developed since the last report?

Six new committers & PMC members have been added to the project. (In Stratos, committers are also invited to join the PMC). Community members have written several blog posts. Lakmal Warusawithana was unanimously elected as the potential Stratos PMC chair.

How has the project developed since the last report?

Stratos has developed a very active community. Community diversity was a concern during the last report as well, and it is being addressed by bringing in a number of non-WSO2 committers & PPMC members. There is lot of interest expressed from & participation from many people. A single Stratos distribution with multiple profiles has been created. The team is also working on creating a simple VM which will allow people to easily try Stratos out. This will bring down the barrier to entry most potential contributors are facing right now, since setting up Stratos is a bit tedious right now. The project is actively working towards further diversification by attracting more non-WSO2 contributors. The domain stratoslive.com has been donated by WSO2 to the Stratos PMC. The project seems ready for graduation. A draft resolution towards graduation has been prepared & is being reviewed. The community is thinking about starting a graduation vote in the near future.

Date of last release:

2014-02-27 - Stratos (incubating) 4.0.0-Alpha (Developer preview only)

When were the last committers or PMC members elected?

2013-03-31

Signed-off-by:

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

Usergrid

Usergrid is Backend-as-a-Service (BaaS) composed of an integrated database (Cassandra), application layer and client tier with SDKs for developers.

Usergrid has been incubating since 2013-10-03.

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

1. Grow the Usergrid community
2. Further embrace the Apache methodology including using Apache infrastructure to move Usergrid forwards
3. Make our first incubating release

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

There are currently some issues with regards to licensing of existing Usergrid code. The community is however aware of this and are working towards suitable replacement(s).

How has the community developed since the last report?

The dev@ mailing list has seen increased amounts of traffic in recent months. As of close of March there are 55 subscribers to dev@. Individual posts to dev@ are shown below.

Jan --> 113
Fed --> 106
Mar --> 179

How has the project developed since the last report?

We have addressed a number of release dependent issues. We continue to work towards these issues with the aim of having our first incubating release soon.

Date of last release:

Usergrid has not yet made an incubating release.

When were the last committers or PMC members elected?

The Usergrid PPMC elected Shawn Feldman (sfeldman) to join the PPMC on Mon, 17 Mar 2014.

Signed-off-by:

[](usergrid) Dave Johnson
[X](usergrid) Jake Farrell
[](usergrid) Jim Jagielski
[X](usergrid) Lewis John Mcgibbney
[](usergrid) Luciano Resende

VXQuery

A standards compliant parallel XML Query processor.

VXQuery has been incubating since 2009-07-06.

As stated in the last report, there are no formal issues blocking graduation. While not a formal requirement, the current plan is (still) to do a second incubator release before graduation (the reasons to do this have not changed).

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?

No change.

How has the project developed since the last report?

There has been a lot of progress

- improving the performance of the system, and on
- running VXQuery on a cluster.

For most tests that we are running now, we see near-linear speed-up and scale-up.

Source code has been moved from svn to git.

The second (and hopefully last) incubator release is being finalized.

Date of last release:

2013-12-06

When were the last committers or PMC members elected?

Marvin Humphrey was added to the PPMC as a mentor in November 2013.
Steven Jacobs was elected as a committer in September 2013.

Signed-off-by:

[x](vxquery) Anthony Elder
[x](vxquery) Jochen Wiedmann
[x](vxquery) Marvin Humphrey
[](vxquery) Sanjiva Weerawarana
[](vxquery) Radu Preotiuc-Pietro