

January 2012

January 2012 Board reports (see [ReportingSchedule](#)).

These reports were due here by the **first Wednesday of the month** (i.e. two weeks before the meeting) so that the Incubator PMC can consider them and be ready for this month's board meeting. This report is **CLOSED**.

Your project might need to report even if it is not listed below, please check your own reporting schedule or exceptions.

Please remember to include:

- The "incubating since" info.
- The project's top 2 or 3 things to resolve prior to graduation.
- The short [description](#) of what your project's software does.
- The **Signed off by mentor**: is for Mentor(s) to show that the Report has been reviewed.

Remember that this is a "text" report, so no markup please.

```
-----
Any23

Any23 is defined as a Java library, a Web service and a set of command line tools to extract and validate
structured data in RDF format from a variety of Web documents and markup formats. Any23 is what it is
informally named an RDF Distiller.

A list of the three most important issues to address in the move towards graduation

    1. Port Any23 code to ASF infrastructure and update license headers
    2. Develop a strong community with organizational diversity and with strong connections to other relevant
ASF communities.
    3. At least one Any23 incubating release

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

Yes, all of the code has been ported from Google Code to the ASF. Thanks to Daniel Shahaf and Michele Mostarda
for
leading the charge here.

How has the community developed since the last report?

All ASF karma has been granted on the repository, and we've received a few JIRA issues, but not from outside
the core set of PPMC members as of yet. The team needs to respond to Paolo Castagna's points regarding RDF
frameworks and collaboration, and will hopefully do so this month.

How has the project developed since the last report?

Any23 was voted into the Incubator by the IPMC on October 1, 2011.

We have Jenkins CI builds going thanks to Lewis John McGibbney, code up and running at the ASF thanks to Michele
and to Daniel, so we're all set to really get kicking!

Signed off by mentor: mattmann

-----
Bloodhound

(project add text here)

Signed off by mentor:

-----
Celix

Celix is an implementation of the OSGi Specification in C.

Celix entered incubation on November 2, 2010.

Over the last few months lots of work has been put into integrating APR and updating the Celix code base to the
proposed code style. This code style has partially been documented on [1]. Also some effort has been put into
updating the source for Visual Studio, even though not yet finished some interesting and helpful patches where
submitted and applied.
```

In Oktober an event was held to create more awareness for services, OSGi and Celix in the Dutch embedded community. The attendance was great, and a follow-up will be planned. Also, in November a talk was given at the ApacheCon.

As listed below, one of the most important issues is the slow growth of the Celix community. With the current discussion about poddlings and how long a poddling is in the incubator, we are discussing a plan how we can move to a more diverse community and be able to graduate. This plan will be included in the next board report.

Most important issues are:

- * Improve robustness (APR, error handling etc), resulting in a first release
- * Generate awareness and grow a community!

Signed off by mentor: marrs

Chukwa

Chukwa is an open source data collection system for monitoring large distributed systems. Chukwa is built on top of the Hadoop Distributed File System (HDFS), HBase and Map/Reduce framework and inherits Hadoop's scalability and robustness. Chukwa also includes a exible and powerful toolkit for displaying, monitoring and analyzing results to make the best use of the collected data.

- * Updated LICENSE and NOTICE files to comply with IPMC requests
- * Updated Hadoop dependency to Hadoop 1.0.0
- * Vote in progress for Chukwa 0.5.0 Release Candidate 3

When Chukwa 0.5.0 is officially released, Chukwa will complete all required goals to graduate incubation.

Signed off by mentor: wrowe, cdouglas

Cordova

December was a busy month. Lots of updates to the code base and our working toolchain.

We would particularly like to thank infra for all the work getting us setup with git. We really appreciate the effort and valuable discussion surrounding it. It is this sort of attention to community that makes us grateful to be members of the Apache community.

- * shipped 1.3 with full windows phone support (NOTE: this was not an ASF release; this is next for us)
- * BlackBerry Playbook support
- * started migration to Cordova name
- * git migration complete
- * battery events api (http://docs.phonegap.com/en/1.3.0/phonegap_events_events.md.html#Events)
- * first version of the project web site up

Graduation concerns:

- * Complete the migration to Apache infrastructure
- * Make an official Apache release
- * Grant committership to new contributors

Signed off by mentor: jukka, gianugo

Deft

Deft is a non-blocking, asynchronous, event driven high performance web framework running on the JVM. The project entered incubation on 2011-07-08.

The most important issues to address in moving to graduation:

1. Prepare an initial incubation release.
2. Attract more committers, and other community members.
3. Secure the previously used URL, "deftserver.org".

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

- Not at this time.

How has the community developed since the last report?

- There have been no significant developments.

How has the project developed since the last report?

- Outstanding tasks are still being worked on.

Signed off by mentor: mnour

DeltaSpike

(incubating since December 2011)

DeltaSpike will consist of a number of portable CDI extensions that provide useful features for Java application developers. The goal of DeltaSpike is to create a de-facto standard of CDI-Extensions that is developed and maintained by the community.

There are currently no issues requiring IPMC or Board attention.

Since our last report on December 12, 2011, we have accomplished the following:

- Finished the discussion about coding conventions and added the corresponding checkstyle rules
- Finished the discussion on the first features
- Finished infrastructure topics together with the infrastructure team (GIT repository, nightly builds with Jenkins, Sonar setup)
 - (Only the JIRA plug-in for GIT is missing because there are issues with using it side by side with the plug-in for SVN)
- Most initial committers signed their ICLA and got their account
- Started discussions on the next features
- Started with the build structure
- Started with the test-infrastructure
- Started to migrate core feature contributions from Seam3 and CODI
- Started to add information to the Wiki to help new folks getting involved
- Started discussions about tools for our documentation and website
- Started with the status page (based on <http://incubator.apache.org/guides/website.html>)

Upcoming major goals:

- Finish test-infrastructure to allow in-container tests
- Finish the setup for the documentation and the website
- Finish the discussions for all features which will be part of the first release

Top 2 or 3 things to resolve before graduation:

1. Build community
2. Create at least one release
3. Create Documentation

Signed off by mentors: Gerhard Petracek, Mark Struberg, jim

DirectMemory

(incubating since October 2011)

Apache DirectMemory is a multi layered cache implementation featuring off-heap memory management (a-la BigMemory) to enable efficient handling of a large number of java objects without affecting jvm garbage collection performance.

A list of the three most important issues to address in the move towards graduation.

- 1) closed - Branding requirements have been met and modified logos and artwork have been published
- 2) Understanding process/decision making guidelines (new committer process is undergoing testing, release process still yet to be worked out)

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

- None - we are going well even having lost one mentor (but always hope he'll keep in touch ;)

How has the community developed since the last report

- Tasks and proposals contributed by non-committers
- The team is voting a new committer

How has the project developed since the last report.

- Some examples are being added to the code base including another interesting integration with Apache Solr

Signed off by mentor: Olivier Lamy, Tim Williams

EasyAnt

Easyant is a build system based on Apache Ant and Apache Ivy. Incubating since 2011-01-31.

Towards graduation, we need to:

- * Create a release
- * Build a community

Since the last report:

- * no particular progress on the diversity of the community
- * very little activity on the developer mailing list

Signed off by mentor: bodewig

Empire-db

Apache Empire-db is a relational database abstraction layer that allows developers to take a more SQL-centric approach in application development than traditional ORM frameworks. Its focus is to allow highly efficient database operations in combination with a maximum of compile-time-safety and DBMS independence. Empire-db has entered Incubation in August 2008.

Activity since last report:

- * Last month we voted for graduation, submitted our resolution which was revoked because of a copy-paste issue
- * A fixed resolution will be submitted for this month's board meeting

Signed off by mentor: dashorst

Flex

Flex is an application framework for building Flash-based applications.

Entered incubation December 30th, 2011.

The podling is just getting started, accounts and infrastructure are being created.

Signed off by mentor: bdelacretaz, greddin

Giraph

Giraph is a large-scale, fault-tolerant, Bulk Synchronous Parallel (BSP)-based graph processing framework that runs on Hadoop. Giraph entered the incubator in August 2011.

Project developments:

- * Sebastian Schelter, and Claudio Martella were added to the PPMC
- * Review Board in frequent use and high traffic on dev and commit lists
- * ICLA on file from all committers

Next steps:

- * Adding new committers.
- * Making a release.

Signed off by mentor: mattmann, omalley

HISE

(project add text here)

Signed off by mentor:

JSPWiki

(project add text here)

Signed off by mentor:

Kafka

(introduced to Apache incubator on Jul 4, 2011)

Kafka provides an extremely high throughput distributed publish/subscribe messaging system. Additionally, it supports relatively long term persistence of messages to support a wide variety of consumers, partitioning of the message stream across servers and consumers, and functionality for loading data into Apache Hadoop for offline, batch processing.

A list of the most important issues to address in the move towards graduation

1. Invite diverse new active committers

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

We incurred significant overhead during our first release; see <http://apache.markmail.org/message/rsxjgrrufc6khlqy?q=list:org%2Eapache%2Eincubator%2Egeneral+from:%22junrao%22> for reference. Firstly, there is no good example to follow for a new podling. Secondly, the licensing rules are complex and the documentation is not enough for a new podling to understand. This gets more complicated since different IPMC members have different views and apply those rules slightly differently. This delays releases and causes overhead to the team. We hope that some of these things can be addressed to make the process easier to other projects.

How has the community developed since the last report?

Activities in the mailing list increased 4x since the last report. Several patches from non-committers have been submitted, reviewed and committed.

How has the project developed since the last report?

We have successfully completed the first release of Kafka after joining the incubator. This releases adds the end-to-end compression feature and mirroring support in Kafka. In addition to this, it also fixes more than 40 issues.

Signed off by mentor: cdouglas

Alan Cabrera

Kato was accepted into the Incubator on 6 November 2008.

Kato is a project to develop the Specification, Reference Implementation, and TCK for JSR 326: the JVM Post-mortem Diagnostics API.

Recent Activity:

* Discussions have been going on with Oracle regarding their participation in JSR 326 and Kato. Initially when the project was started Oracle expressed enthusiasm at participating, however during the acquisition of Sun Oracle was understandably quiet. In more recent times Oracle expressed interest in the project, but this interest hasn't amounted to anything concrete and we've been pushing for Oracle to make a definitive statement regarding it's position on the JSR.

The following is planned for next reporting period:

- * Without any response from Oracle we will assume that Oracle has no interest in the project, and decide what to do with the JSR and Kato unilaterally.
- * Decide in what form the podling should continue, if at all.

Before this project can graduate we need to encourage more participation in the project and grow the community.

Signed off by mentor:

Kitty

Incubating since 2010-10-03

Description

Kitty is a command line JMX Client written in Groovy.

It is a lightweight, production focused, application server performance diagnostic and management utility.

The three most important issues to address in the move towards graduation:

- 1) The community around this project is absent, looking at the amount of mails sent to the mailing list in the past months
- 2) The lack of a "Product Owner" who can prioritise stories/tasks, leaving the developers without a clue about the functionalities to implement
- 3) The code is a port from another language (Jython), and the history of the past project seems lost, so there isn't any reference about how the actual code should work. Also, there is no test (acceptance criteria) in support of that nor for the new features to implement.

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

What seems relevant to me at the moment is the lack of communication between the committers;

it would be fine if there was a clear statement about who is still interested and who is not in participating in the project, and how to reorganise the now shrunken team, if still there is one.

How has the community developed since the last report?

Sadly, the community hasn't develop at all; even worse, we faced the loss of some members who were actively involved in the project development/management.

How has the project developed since the last report?

Jira is not updated, leaving a shade of uncertainty about what has been fixed and what is still in progress since the last report.

It seems that the last update was done in October, the 4th, with a fix regarding the correct navigation through the server tree in a file-system fashion.

Signed off by mentor: kevan, jim

Lucy

Lucy is a loose port of the Apache Lucene search engine library, written in C and targeted at dynamic language users.

Lucy was voted into the Incubator on July 22, 2010.

Progress since the last report:

- * Added new Committers/PPMC-members Nick Wellnhofer and Logan Bell.
- * Made substantial improvements in committer diversity.
- * Released version 0.2.2
- * Technical developments:
 - * Added several new classes relating to Unicode text processing.
 - * Improved clustering support.
 - * Finished eliminating Perl-licensed dependencies, resolving LEGAL-86.
 - * Set up per-commit build-and-test buildbot.
 - * Finished porting Clownfish compiler to C.

Top priorities prior to graduation:

- * None. We will likely be discussing graduation as soon as version 0.3.0 is released.

Issues for Incubator PMC or ASF Board:

- * None at this time.

Signed off by mentor: mattmann

Mesos

Mesos is a scalable cluster manager that can dynamically share resources between multiple computing frameworks, including Hadoop, Spark and MPI.

Mesos entered incubation on December 23, 2010.

Progress since the last report:

- * Shooting for an end of January release. Hope to call a release vote by early next week.
- * Internal Twitter repo almost eliminated (all code being developed through Apache).
- * Documentation added (in anticipation of end of month release).
- * Support for Mac OS X Lion almost complete (expected to be included in end of month release).
- * New autotools support almost complete (expected to be included in of end of month release).
- * Mesos provided replicated log deployed in production (and optimized).
- * New web UI, including task history viewing, committed.
- * Features and bug fixes contributed by 3+ new contributors.

Top priorities prior to graduation:

- * Making our first release (we want to do that this month).
- * Adding more committers to the project from the various developers that have contributed to Mesos.

Issues for Incubator PMC or ASF Board:

- * None at this time.

Signed off by mentor: tomwhite

ODFToolkit

The ODF Toolkit is a set of Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents. Unlike other approaches which rely on runtime manipulation of heavy-weight editors via an automation interface, the ODF Toolkit is lightweight and ideal for server use.

* ODF Toolkit entered incubation on Aug 1st, 2011.

* Most important issues to address.

- 1) Successful podling release.
- 2) Growing the community, increasing diversity of committers

* Any issues that the Incubator PMC or ASF Board might wish/need to be aware of
None at this time.

* How has the community developed since the last report
We have a new committer: Oliver Rau.

* How has the project developed since the last report.
1) With the help of mentor, Jenkins support for the build is ready.
2) We are voting for the candidate (RC7) of the first release.

Signed off by mentor: nick

Oozie

Oozie is a workflow management and scheduler primarily for Hadoop based jobs.

Oozie entered the incubation on July 11, 2011.

* A list of the three most important issues to address in the move towards graduation:

- * Make the first Oozie release from Apache incubation.
- * Further improve the documentations: user, development for quicker adoption
- * Establish the formal release process.

* Any issues that the Incubator PMC or ASF Board might wish/need to be aware of:
* No issues.

* How has the community developed since the last report:
* Oozie dev/users are active in the email lists (1152 Email in oozie-dev and in oozie-user 212).

- * Around 7 new developers are contributing to the process and coding.
- * 84 JIRAs were created since last report.

* How has the project developed since the last report.

* A lot of new features are added. Fixed a lot of bugs too. Nearly 85 JIRAs were closed/resolved in last three months.

* Oozie product web page is improved and more usable. Further improvement is ongoing on.

* The process has started for cutting the first release from Apache.

Signed off by mentor: Alan Gates, cdouglas

Openmeetings

* Openmeetings is a free browser-based software that allows you to set up instantly a conference on the web.

* A list of the three most important issues to address in the move towards graduation are:

- * Migrating infrastructure
- * Code Audit (Apache RAT is now integrated which should help us)
- * Packaging first incubator release

* A question for PMC: Should we start a demo server on Apache, so people can benefit from communication?

* Community status: 1 committer (German Grekhov) joined.

* How has the project developed since the last report.

- * Jira, SVN and committers migrated
- * Web site created <http://incubator.apache.org/openmeetings/>
- * Half-duplex functionality was being reviewed

Signed off by mentor: jim

Tashi

2012-January Tashi Incubator Status Report

Tashi has been incubating since September 2008.

The Tashi project aims to build a software infrastructure for cloud computing on massive internet-scale datasets (what we call Big Data). The idea is to build a cluster management system that enables the Big Data that are stored in a cluster/data center to be accessed, shared, manipulated, and computed on by remote users in a convenient, efficient, and safe manner.

Tashi originally encompassed just the tools to manage virtual machines using Xen and QEMU, but has been merged with Zoni, which manages the physical aspects of a cluster like power control, network settings and handing out physical machines.

Development activities have included:-

- * Accounting server has been added to the codebase
- * Primitive scheduler changes
 - * bug fixes
 - * Add support for user choice of dense packing or not
 - * Guard against starting more than one VM with persistent disk
- * Client changes
 - * Check syntax of user commands
 - * Add support for querying available images
 - * Add support for querying image size
 - * Add support for copying of images
- * QEMU VMM changes
 - * bug fixes
 - * Reserve some memory for the host itself
 - * Make scratch location configurable
 - * Live migrations take a long time, eliminate some timeout values
- * Cluster manager changes
 - * bug fixes

- * Reduce network traffic
- * Move accounting functions to new accounting server
- * Branched off new stable version and release candidate
- * Audit compliance with Incubator policies

The project is still working toward building a larger user and development community. User groups have been identified in Ireland, Slovenia and Korea, Malaysia, as well as the United States.

Items to be resolved before graduation:

- * A stable branch exists which could be a release candidate, but the codebase is large and test hardware is currently in short supply. We are confident that the code in the stablefix branch will work if running QEMU emulation, Pickle or sqlite data storage, primitive scheduler. Xen, other data stores and schedulers have not been tested recently.
- * Develop community diversity (Committers currently at Telefonica, Google and CMU)

Signed off by mentor:

VXQuery

The VXQuery Project implements a standard compliant XML Query processor. It has been in incubation since 2009-07-06.

Recent activities:

- * minimal activity
- * considering future of the project

Top issues before graduation:

- * Create a release
- * Build community (developer and users)

Signed off by mentor:
