

July 2011

July 2011 Board reports (see [ReportingSchedule](#)).

These reports were due here by **Wednesday, 13 July 2011** so that the Incubator PMC can relay them to the board. **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.
- A 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.

Ace

Apache ACE is a software distribution framework that allows you to centrally manage and distribute software components, configuration data and other artifacts to target systems. ACE started incubation on April 24th 2009.

There are currently no issues requiring board or Incubator PMC attention.

Community:

- * Jean-Baptiste, Bram and Marcel got accepted at ApacheCon NA 2011.
- * Started discussions about a possible donation of ACE IDE to the project (see <https://www.assembla.com/spaces/eclipse-ace/wiki>)

Software:

- * Wrapped up the Maven integration, readied the codebase for doing releases.
- * We did our first release: 0.8.0-incubator!
- * Bugfixes after initial feedback on the release.

Licensing and other issues:

- * None at the moment.

Things to resolve prior to graduation:

- * Not sure, we grew the community, did a release, so this might actually be a good time to start discussing graduation.

Signed off by mentor: cziegeler, bdelacretaz

Airavata

Airavata is a software toolkit which provides features to compose, manage, execute, and monitor large scale applications and workflows on computational resources ranging from local clusters to national grids and computing clouds.

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

1. Remove incompatible license dependencies and implement over ALV2 compatible libraries.
2. Simplify the build process and provide simple to use cases.
3. At least one Apache Incubator release.
4. Develop a strong community with organizational diversity and well aligned with existing ASF projects as outlined here: <http://wiki.apache.org/incubator/AiravataProposal#Alignment>

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

No, not at this time.

How has the community developed since the last report?

A web designer Barbara Hallock has contributed, Airavata Logo, web site banner and CSS. Patanachai Tangchaisin enquired in the mailing list about contributing to the project.

How has the project developed since the last report?

- * Followed all the instructions on the Crypto registration, sent BIS notice and updated the Exports page with cryptographic dependencies.
- * Airavata website is designed and developed in the staging area is now made live.
- * Brainstorming on alternatives for incompatible license dependencies
- * Started development of the various components, importing, cleaning up the donate code by adding ASV2 license headers, changed package names and applying code style templates and formatting.

Signed off by mentor: mattmann, rgardler

Bigtop

Bigtop is a project for the development of packaging and tests of the Hadoop ecosystem.

Bigtop entered incubation on June 20 2011

Community

- * Mailing lists have been created.
- * New committer accounts are being created, some pending ICLA.
- * SVN directory structure has been created, karma granted for initial committers.
- * JIRA has been initialized for BIGTOP.
- * Confluence Wiki space has been created.

Issues before graduation

- * Migrate code to Apache SVN from Cloudera.
- * Create Bigtop web site.
- * Make an incubating release.
- * Grow the community size and diversity.

Licensing and other issues

- * Work to obtain CCLA from Cloudera regarding license grant for existing Bigtop code from Cloudera GitHub repository is underway.

Signed off by mentor: tomwhite, phunt, omalley, stevel

Bluesky

BlueSky has been incubating since 01-12-2008. It is an e-learning solution designed to help solve the disparity in availability of qualified education between well-developed cities and poorer regions of China.

The following items have been performed since the last reporting period

- 1.solve the problem of the bad effect of recording vedio.
- 2.integrate with IPTV function.

The following items are planned for the next reporting period:

- 1.test the system in Linux environment.
- 2.add the access function of mobile terminal.

Signed off by mentor:Bonnie Liu

Celix

Celix is an implementation of the OSGi Specification in C.

Celix entered incubation on November 2, 2010.

Over the past few months work has been done on integration and usage of the Apache Portable Runtime (APR) in Celix.

The work has mostly focusses on stability and robustness of the codebase, but some new features were added:

Log Service and Writer: for logging and reading the log
Service Factories: needed for the log service

Also, in March, a talk was given to the OSGi User Group France. There is a lot interest in the project, but people seem to be waiting for a more stable version.

For community building it might be interesting to work towards a release later this year to make it easier to

download and use Celix.

A talk on Celix has also been accepted for the ApacheCon in November, which is a good opportunity to show Celix to the world.

Also parts of Celix have been ported to Windows. Mostly the utils used by the framework, and not yet the framework itself. This is a lengthy job because Celix is using C99, and Windows (Visual Studio) only support C89 /C90.

Most important issues are:

- Improve robustness (APR, error handling etc)
- Generate awareness and grow a community!

Signed off by mentor: marrs

Chukwa

Chukwa is Hadoop monitoring system

Chukwa is an open source data collection system for monitoring large distributed systems.

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

1. Simplify the build process and provide simple to use cases.
2. At least one Apache Incubator release.
3. Develop a strong community with organizational diversity and well aligned with existing ASF projects as outlined here: <http://wiki.apache.org/incubator/ChukwaProposal#Alignment>

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

No, not at this time.

How has the community developed since the last report?

There are some activity on the user mailing list, and some new users participated in adopting Chukwa for processing and visualizing big data analytics.

How has the project developed since the last report?

* There are several bug fixes, and migration to HBase 0.90.3 integration for data storage.

Signed off by mentor: berndf, cdouglas

EasyAnt

EasyAnt is a tool built on top of Ant and Ivy providing a standard approach to building java projects without locking the users in.

Incubating since January 31st 2011

There is still one ICLA of one of the original contributors missing and the team has started to look into replacing or removing the code contributed, we expect to have a plan and report on it in August.

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.

After having published our 2.1.0 release in April, things went a bit slower during the last quarter. Recently our committers have fixed smaller issues and we have made the Empire-db code base available on git.

We regularly receive user requests on our user list asking for help and or improvements. Some users submitted small patches through Jira.

For the next term we have planned to start a major API improvement which might help making Empire-db attractive for even more users.

In order to graduate we still need to increase awareness of the project through various channels in order to grow our community. This means participating in blogs and write our own articles. We will set up a discussion on how exactly our communication can be improved.

Signed off by mentor:

Flume

Apache Flume is a distributed, reliable, and available system for efficiently collecting, aggregating, and moving large amounts of log data to scalable data storage systems such as Apache Hadoop's HDFS.

Flume entered incubation on June 12th, 2011.

Issues before graduation

- * Migrate code to Apache SVN from Cloudera.
- * Work to obtain CCLA from Cloudera regarding license grant for existing Flume code from Cloudera GitHub repository is underway.
- * Migrate mailing lists from Cloudera infrastructure to Apache mailing lists.
- * Create Flume web site.
- * Make an incubating release.
- * Grow the community size and diversity.
- * Licensing and other issues

Community

- * Mailing lists have been created.
- * Initiated migration of mailing lists from Cloudera infrastructure to Apache mailing lists.
- * Voted to add Nick Verbeck as a committer and PPMC member.
- * New committer accounts have been setup
- * SVN directory structure has been created, karma granted for initial committers.
- * JIRA has been initialized for FLUME
- * Confluence Wiki space has been created.
- * Work on migrating JIRA information from Cloudera hosted JIRA instance to Apache JIRA is underway.

Project developments

- * A Flume 0.9.4 release was made with the intention that it is the last pre-Apache release of the Flume project.

Signed off by mentor: phunt, rgoers

Imperius

The Imperius project has been retired. Some additional administrative shutdown is needed (is on my TODO list)

Signed off by mentor: kevan

JSPWiki

JSPWiki has been incubating since September 2007.

JSPWiki is a JSP-based wiki program.

There were 6 bugs fixed.

And also some progress in the incubation process, the last open issue (JSPWIKI-636 Clean unit test builds) was also fixed.

But still no activity to make an official Apache release.

Very few developers seem to have time to dedicate for JSPWiki.

There appears to be a lack of interest in the new releases, though there is still activity on the users list, which shows that people are using JSPWiki.

The developer list currently has 93 subscribers; and the

user list has 192 subscribers.

Signed off by mentor:

Kitty

Apache Kitty is a command-line, user friendly JMX utility, incubating since November 2010.

* Top graduation blockers: 1) Still resolving major bugs, 2) Have not yet added enough user-friendly / convenience functionality

* Description: Kitty is a lightweight, universal, and user-friendly / fast JMX utility

* The latest developments is the Kitty community has grown after reaching out to the Groovy development community, we have seen much more community involvement than the past. Voted in Alessandro Novarini as a Kitty Committer.

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:

- * Inaugural release 0.1.0 successfully launched.
- * Added new mentor Joe Schaefer.
- * Achieved threshold of knowledge dispersal for one layer of the Lucy codebase (Charmonizer).
- * Technical developments:
 - * Migrated Charmonizer build system to Make.
 - * Numerous portability improvements addressing CPAN testers results for 0.1.0.
 - * Continued Clownfish C porting.
 - * Started on Ruby bindings.

Top priorities prior to graduation:

- * Disperse knowledge of codebase among developers.

Issues for Incubator PMC or ASF Board:

- * None at this time.

Signed off by mentor: mattmann, hossman

Mesos

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

Mesos entered incubation on December 23rd, 2010.

Community

- * Mesos is being used at Twitter, Conviva, UCSF and UC Berkeley.
- * We moved the source code into Apache SVN and started tracking development on JIRA.
- * We accepted patches from several contributors through JIRA and GitHub.
- * We've also made various improvements to the website and documentation to make the project easier to set up and use.

Project development

- * Since the last report, we completed a migration to protocol buffer based communication (to enable rolling upgrades), several fault tolerance improvements, several fixes to Linux Container based isolation, a new client API based on protocol buffers to allow applications to also work on future versions, and porting existing applications (e.g. Hadoop) to use the new API.

Most important issues to address

- * Put together a first incubating release of Mesos.
- * Grow the community and add more committers.

Signed off by mentor: tomwhite

OGNL

Apache OGNL is a Java development framework for Object-Graph Navigation Language, plus other extras such as list projection and selection and lambda expressions.

The project joined the Incubator on April 26, 2011.

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

- make at least a big release that shows to Commons PMC the compliancy level of OGNL to Commons policies;
- attract more users, at the moment only Apache Struts people are interested because they heavily use OGNL;
- attract new committers.

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

- no news about Marc Andrew Davidson's even if was requested twice;
- not every PPMC member joined the private ML yet;
- due to OpenSymphony dismissing, we need importing old issues to ASF jira. We finally contacted someone from OpenSymphony and obtained the huge Issue dump/XML backup, there is someone that is working on it, but due to security policy (the dump contains also users and password hashes), before importing it, an open-heart surgery is needed.

* How has the community developed since the last report

Community activity decreased since codebase has been imported & polished; we finally reached Luc Blanchard, the other OGNL original author, that signed & submitted the OGNL SoftwareGrant & ICLA - we also voted him as a new committer since he should have been included in the initial committers list

* How has the project developed since the last report.

A reported above, codebase has been imported & polished, we are in the middle of a big refactoring stripping off generated code by JavaCC that was versioned in the SCM and modified.

Signed off by mentor: grobmeier

Olio

Olio is a web 2.0 toolkit to help developers evaluate the suitability, functionality and performance of various web technologies by implementing a reasonably complex application in several different technologies.

There is not much activity on this project. No major user requests have been received.

Graduation From Incubation:

The user base is not showing growth. Since this is a test application/workload, it is being used as such for short-term projects. As such, we need to decide how to handle this project. It is still useful as is.

Signed off by mentor:

OpenNLP

OpenNLP is a machine learning based toolkit for the processing of natural language text. It supports the most common NLP tasks, such as tokenization, sentence segmentation, part-of-speech tagging, named entity extraction, chunking, parsing, and coreference resolution. These tasks are usually required to build more advanced text processing services.

OpenNLP entered incubation on 11/23 2010.

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

- * Establish open regression tests for the parser and coreference component
- * Identify and encourage new contributors on the path to committership

Signed off by mentor:

OpenOffice.org

* OpenOffice.org entered incubation 2011-06-13.

OpenOffice.org is an open-source, office-document productivity suite providing six productivity applications based around the OpenDocument Format (ODF). OpenOffice.org is released on multiple platforms. Its localizations support 110 languages worldwide.

* Most important to address

1) Migration of the legacy OpenOffice.org website's content and services to Apache infrastructure, including defect tracking and wiki instances which have undergone significant customizations.

2) Completion of the IP-review portions of the incubation checklist, which will require getting an amended SGA from Oracle to cover additional source files.

3) A succesful Podling Release

* Issues

None at this time.

* Community development progress

Initial work toward establishing committers and a PPMC has, as of 2011-07-08, 56 committers including 2 mentors and 38 PPMC members.

There are tentative steps toward engaging the existing OpenOffice.org community and users. There is limited discussion with related Apache Projects such as POI and the possible incubation of the ODF Toolkit project.

* Project development progress

Addressing issues on migration of code repositories and other pieces of project infrastructure. Discussing migration of existing OpenOffice.org as-is versus "house cleaning" for a streamlined initial website at Apache.

Created a podling project web site: <http://incubator.apache.org/openofficeorg/>

Started more detailed planning discussion on public wiki: <https://cwiki.apache.org/confluence/display/OOOUSERS/Project+Planning>

Signed off by mentor: jim

Rave

(project add text here)

Signed off by mentor:

Socialsite

(project add text here)

Signed off by mentor:

Sqoop

A tool for efficiently transferring bulk data between Apache Hadoop and structured datastores such as relational databases.

Sqoop entered incubation on June 11, 2011.

Community

- * Mailing lists have been created.
- * New committer accounts have been setup except for Guy le Mar whose account request has been submitted.
- * SVN directory structure has been created, karma granted for initial committers.
- * JIRA has been initialized for SQOOP.
- * Confluence Wiki space has been created.
- * Work on migrating JIRA information from Cloudera hosted JIRA instance to Apache JIRA is underway.

Issues before graduation

- * Migrate code to Apache SVN from Cloudera.
- * Migrate mailing lists from Cloudera infrastructure to Apache mailing lists.
- * Create Sqoop web site.
- * Make an incubating release.
- * Grow the community size and diversity.

Licensing and other issues

- * Work to obtain CCLA from Cloudera regarding license grant for existing Sqoop code from Cloudera GitHub repository is underway.

Signed off by mentor: Olivier Lamy (olamy), phunt

Tashi

2011-July 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 has previously encompassed just the tools to manage virtual machines using Xen and KVM, but is gaining the facility to hand out physical machines as well.

Development activities have included:-

- * Refactor primitive scheduler to be less convoluted
- * Ensure that an old CM handle expires to not talk to a dead CM
- * Use virtio networking by default for performance
- * Enable config option for Miha Stopar's auto host registration
- * Clean unused and untested modules from stable branch
- * Reduce VM startup time when using scratch (old sparse files)
- * Conversion of sparse file scratch space to Linux LVM2
- * Work on migrating VMs between hosts
- * Resource usage messages sent to clustermanager for accounting

The project is still working toward building a larger user and development community. User groups have been identified in Ireland, Slovenia and Korea, as well as at Georgia Tech. CMU usage is growing as other groups hear about the availability of the resource. Intel has restructured its research division and folded some operations into adjoining academic sites.

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.

- * Should have example accounting code
- * 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:

- * no progress due to external commitments of the current developers and minimal community size
- * considering future of the project

Top issues before graduation:

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

Signed off by mentor:
