

April 2011

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

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

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.

This report is closed.

```
-----
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:

Angelo gave a talk about massive device deployment and Karl and Marcel did a tutorial about ACE in the Cloud at
EclipseCon / OSGi DevCon.
We're gathering EclipseCon / OSGi DevCon feedback to populate the roadmap and raise the corresponding Jira
issues.

Software:

A lot of work has been done on the build to fully support Maven 3.
We started to work on the documentation also.
Use of Apache Karaf as default runtime is in progress.
Cloud extensions based on jclouds have been added, including a new launcher.

Licensing and other issues:

None at the moment.

Things to resolve prior to graduation:

We created a SVN maven repo waiting for some Felix dependencies release.
Make a release.
Grow the community some more.

Signed off by mentor: pauls, bdelacretaz, cziegeler, hadrian

-----
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
* test the stability of system on a certain scale.
* add the function of recording video in the teacher part.

The following items are planned for the next reporting period:
* go on testing the stability of system
(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.

During the EclipseCon a talk was given over Celix, also spoke with several OSGi people who are interested in the progress, especially related to Universal OSGi (which details specification and requirements for OSGi in other languages).

During the last few weeks work has been done to get rid of memory leaks and null pointers.

Currently we are looking at the requirements/details for bundle deployment and updating, possibly in combination with Apache Ace.

Most important issues are:

- Define a status overview of what is and isn't working.
- Move all "sub" project to the new structure.
- Use APR for all file handling, threading etc.
- Generate awareness and grow a community

Signed off by mentor: marrs

Chukwa

Chukwa is a distributed log collection and processing system built on top of Hadoop and HBase. Chukwa is waiting for a official working release of Hadoop 0.2x and HBase 0.9x. There are certain API incompatibility in metrics collection system to prevent them working together. This is a blocker for Chukwa 0.5 release.

Chukwa entered incubation on August 5th, 2010.

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

issues/agenda:

- 1 ICLA missing. The person does not answer since Feb 4th 2011, when he indicated he was working on getting the ICLA. The svn logs have been gathered and shared between the PPMC members in preparation for a possible cleanup of the svn tree of that person's contribution.

activity since last report:

- the web site <http://incubator.apache.org/easyant/> is published via svnpubsub
- while reviewing the licences of the software included in the svn tree of Easyant, some dependency on external LGPL library was found. There are somehow optional to EasyAnt core, but not for the "plugin" dedicated to the integration of that third library. A consensus was found by the developer community to not release such software under the ASF.

Signed off by mentor: antoine, 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:

- We have successfully completed and published our 2.1.0 release with few but major improvements.
- We have increased your community by taking on a new long-time contributor as a new committer.

Important issues to address for graduation:

We still need more contributors and more attention in the Java community. Recently we experience increased traffic on the user lists and there are plenty of plans on the development side. The challenge will be to keep users interested and let them participate in the development process.

Signed off by mentor: dashorst

HCatalog

HCatalog is a table and storage management service for data created using Apache Hadoop.

HCatalog has been incubating since March 2011.

This is the first HCatalog Incubator status report. The project is still setting up infrastructure, etc. Mailing lists and JIRA have been set up. The initial code has been received from Yahoo. Details of the status are available at <http://incubator.apache.org/projects/hcatalog.html>

Currently there are 14 subscribers on the user list and 15 on the dev list.

Signed off by mentor: gates

Imperius

The Imperius project is currently holding discussions about suspending the project. The project did create an SPL implementation and created a release. However, it has not been able to build a community around the project. The consensus of the project appears to be to suspend work on the project. The community should be holding a vote on this soon.

Signed off by mentor: kevan

Jena

Jena was accepted into the Apache Incubator November 2010 (http://mail-archives.apache.org/mod_mbox/incubator-general/201011.mbox/%3C4CEC31E4.9080401@apache.org%3E). It is an existing project migrating to ASF. It has a large codebase that provides a semantic web framework in Java that implements the key W3C recommendations for the core semantic web technologies of RDF and SPARQL.

Three most important issues to address for graduation:

- * Create a code base in Apache with legal clearance.
- * Establish the Apache-based community for developers, contributors and users
- * Create and communicate open project processes aligned to Apache

Issues for the Incubator PMC or ASF Board:

None.

Community development:

We have started to advertise the existence of Apache-hosted mailing lists. Traffic on these lists continues to grow.

A new work area on enhancing the native storage system with transaction support is being conducted in public on Apache cwiki and mailing lists.

Project development:

We are still awaiting the software grant from HP. As this is the majority of the codebase, we are waiting for this before migrating any code into Apache SVN. We understand that HP has progressed this and has decided on the details for signing.

Signed off by mentor: rgardler

JSPWiki

(project add text here)

Signed off by mentor:

Kitty

Community building efforts have been initiated within the Groovy community to help us get to our first release. We have succeeded in generating interest and Groovy contributors have offered to assist on the project by contributing patches/expertise.

No further development advances have been made.

Signed off by mentor: kevan

Lucene.NET

Lucene.Net was accepted into the Apache Incubator in February 2011. Originally it was a sub project of the Lucene Project.

Lucene.Net is a port of the Lucene search engine library, written in C# and targeted at .NET runtime users.

Recent Activity:

- * New website has been launched
- * All Java Lucene changes for 2.9.4 have been ported and pre-release testing has begun.

Current Activities:

- * Complete testing of 2.9.4 version for release
- * Implement a process for continuous integration, unit testing, and code coverage
- * Develop a process to automatically (as much as possible) convert the Java Lucene code to C# (to maintain our line by line port)

Goals for graduation:

- * Have a nearly fully automated process to convert Java Lucene to C#.
- * Release Lucene.Net 3.0.3 (port of Java Lucene 3.0.3)
- * Have a new .NET version of Lucene utilizing .NET constructs and idioms

Signed off by mentor: bodewig

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 committer David Wheeler.
- * All but two dependencies with problematic licenses eliminated.
- * RAT report integrated.
- * Release guide expanded.
- * Continuous minor developments (e.g. ported more Perl code to C, binding compactness improved, Windows and C++ compatibility restored, miscellaneous bugfixes and improvements).

Top priorities prior to graduation:

- * Disperse knowledge of codebase among developers.
- * Make a release (0.1-incubating, mattmann volunteered to RM)

Issues for Incubator PMC or ASF Board:

- * None at this time.

Signed off by mentor: mattmann

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 actively being used at Twitter, Conviva, UCSF and UC Berkeley.
- * We are working on importing the source code to Apache, likely after we finish merging in a few features that have been developed in the past few months.

Project development

- * Since the last report, we have worked on migrating communication to use protocol buffers (to allow rolling upgrades), improving fault recovery, and improving event logging and monitoring

Most important issues to address

- * Import source code to Apache SVN
- * Move discussion of new features and patches to JIRA
- * Put together a first release

Signed off by mentor:

MRUnit

MRUnit - a library to support unit testing of Hadoop MapReduce jobs.

MRUnit entered incubation on March 8th, 2011.

Community

- * Mailing lists have been created
- * New committer accounts have been setup
- * SVN and git mirror have been created, karma granted for initial committers
- * Initial codebase established - clone of current Hadoop mrunit contrib
- * JIRA has been initialized for MRUNIT
- * Wiki space created on cwiki
- * A discussion on when we should cut a first release is ongoing.

Issues before graduation

- * Create the MRUnit web site
- * Make an incubating release
- * Grow the community size and diversity

Licensing and other issues

- * none - MRUnit was originally a subproject of Hadoop

Signed off by mentor: phunt, mattmann, nigel

Olio

2011-April Olio Incubator status report

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.

Things have been fairly quiet on both the user and developer list this quarter.

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.

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.

We are still working on our first release. All the planned testing in our test plan is done now, we will prepare Release Candidate 6 in the next days and start the vote to release it.

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

- * Resolve potential IP issues around releasing training models
- * Do a release
- * Establish open regression tests

Signed off by mentor: Grant Ingersoll

Rave

Apache Rave is A new WEb And SOcial Mashup Engine. It will provide an out-of-the-box as well as an extendible lightweight Java platform to host, serve and aggregate (Open)Social Gadgets and services through a highly customizable and Web 2.0 friendly front-end.

Rave just entered incubation on 2011-03-01, this is our first board report.

Mailing lists, SVN and user accounts have been created, Issue tracking (JIRA) and continuous integration (Jenkins) have been setup.

All the three initial code donations have been provided and imported in SVN.

Website creation and management is underway, using cms.apache.org, but currently still in staging only. Getting the website ready for publishing is high priority.

Branding requirements will be taken care of together with the website.
Discussion on creating and selecting a Rave logo is underway on the developer mailing list.

The committers voted on and selected Maven as build engine for Rave.
A basic Maven project is committed. Discussion about the new code base and features to implement has started. The project code base will be created from scratch, not be based on or directly incorporate the initial code donations.
Under discussion on the developer list also is choice of development methodology and style.

Next steps:

- * Get the Rave website ready for publishing, select a Rave logo and fulfill the branding requirements.
- * Reach out to the community and build up awareness of Apache Rave.
- * Proceed with building up the new code base and work towards a first (incubator) release.

Signed off by mentor: ate, rgardler, hadrian

Socialsite

(project add text here)

Signed off by mentor:

Tashi

2011-April 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:-

- * Importing bug fixes from CMU deployment
- * Don't display exceptions if failure is expected
- * Validate fields and clean invalid instance entries
- * Eliminating inconsistent state between node and cluster manager
- * Allow hosts to be soft powered off via ACPI
- * Refactor primitive scheduler to be less convoluted (ongoing)
- * Define a stable branch for tested code
- * Example support for dynamic scratch space disks
- * Enforce designated copy-on-write data filesystem for Qemu

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.
- * Scratch space integration would be desirable, accounting integration is probably a necessity
- * 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 progress due to external commitments of the current developers and minimal community size

Top issues before graduation:

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

Signed off by mentor: