

April2009

April 2009 Board reports (see [ReportingSchedule](#)).

These reports were due here by **Wednesday, 8 April 2009** so that the Incubator PMC could relay them to the board.

THIS PAGE 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
-

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.

Recent Activity: We finished the required work by Bill, our mentor. We'd tried to use Dirac for replacing FFmpeg, while we were waiting for the results of second code view. However, some problems happened and could be resolved. Reading api of FFmpeg and vobistheora is still undergoing. The reason, why it is retarded, is that we had to do several projects simultaneously, besides the classes. The running demo of Bluesky is postponed too, but i can assure that it would be completed next week, after i finish the coding task this week.

Next step:

*Waiting for the second code view results.

*Recording Bluesky running demo.

*Tried to work out a reliable sketch of replacing FFmpeg in two weeks.

Cassandra

Cassandra is a distributed storage system for managing (un)structured data while providing reliability at a massive scale. It has been in incubation since January 2009.

The incubator project had a very slow start. Code started to get moved over from Google Code just about 5 weeks ago. But a very basic web site and the wiki is in place. A new committer was elected. Unfortunately there is some fraction on about contributions, code reviews and branching that needs to be addressed ASAP. This is a community issue and we hope to get this resolved. More on that in the next report.

- Consensus about the development process (code reviews, branching etc)
- JIRA permissions and configuration
- A bit more information on the web site

Empire-db

Empire-db is a relational data persistence component that aims to overcome the difficulties, pitfalls and restrictions inherent in traditional Object Relational Management (ORM) approaches.

Empire-db is on the Apache Incubator since July 2008.

Community:

- Development seems to pick up some momentum. Parts of the code are now unit-tested and two new integration examples have been added (Apache CXF)
- dev@ and user@ traffic is rising, this brought some fixes to PostgreSQL Driver as well as some possible changes for future releases.

Release:

- Preparations for a release are under way, this will be our first maven-based release
- We hope the switch to maven will lower the barrier for adoption

Issues prior to graduation:

- Grow the community,
- Successful maven-based release
- Even more integration example code (Apache Wicket?) + more documentation

Imperius

Imperius has been incubating since November 2007.

Imperius is a rule-based infrastructure management platform.

Web site:

Code: Development continues on the code base.

Community:

- Several bug fixes were provided.
- In the process of adding two new committers (David Wood, Xiping Wang) to the community.
- Discussions underway regarding first release

Issues prior to graduation:

- Grow the community
- Successful completion of the release

JSPWiki

2009 April JSPWiki Incubator status report

JSPWiki has been incubating since September 2007.

JSPWiki is a JSP-based wiki program.

During the past three months, the JSPWiki community has still continued to increase. The developer list now has 82 names, a modest increase from 72 the last time, and the user list is now 167 people strong, ten more than three months ago.

The big push now is for our graduation. The biggest hurdles have been overcome:

The packages are now called "org.apache.wiki" [please see Tomcat bug https://issues.apache.org/bugzilla/show_bug.cgi?id=46462 to see why we couldn't use org.apache.jspwiki]; JSPWiki has switched to the Stripes framework; and we have switched to use the JCR backend (e.g. Apache Jackrabbit). As a result of a vote, the support for webdav has been dropped. The fairly old jsonrpc has been replaced by jabsorb.

However, the trunk is in a fairly immature state. For example, it is not deployable at all, which means that nobody can yet use the trunk for its intended purpose.

Our main concern is to stabilize the code to some state where it would be actually usable before making a release.

The biggest slowdown is simply the fact that JSPWiki is (still) a completely volunteer-developed project, with nobody actually receiving any money for its development. While this is great in the sense that development is relatively easy to join, and communities are not overrun by corporate requirements, it also means that the committer's priorities often conflict with the object of speedy graduation.

We are planning to have an "alpha" grade release as our first official Apache release, after which we should be fairly ready for graduation.

Ki

Ki is a powerful and flexible open-source Java security framework that cleanly handles authentication, authorization, enterprise session management and cryptography.

Ki has been incubating since June 2008.

The project team voted at the end of February to change the project name from JSecurity (its previous name prior to entry into the ASF Incubator) to 'Apache Ki' for a number of reasons, which are documented in full detail here:

<http://markmail.org/thread/zcmi4pjv2bbf4574>

Project infrastructure is being changed to match this name change, but is not yet fully complete. Mailing lists are pending still.

There was a post on the old jsecurity.org website by an individual stating that they would appreciate if we used a different name other than Ki, posted here: <http://www.jsecurity.org/node/1081#comment-289>.

One of our project members responded and have not received any responses further. The project team is now internally debating whether we need to change the project name yet again.

The old jsecurity.org project website has entered an archival state, clearly pointing that all users should be redirected to the new Incubator ki site. A crontab entry has been created to auto-export the cwiki to here: <http://incubator.apache.org/ki>

We're happy to report that development and user list activity has steadily increased since we started incubation last year, with March (last month) being the most active month to date.

The project team is not considering graduation at this point, as the code is not ready for an Apache release. Once IP clearance is complete, we'll attempt our first incubator release.

One mentor, Emmanuel Lecharny, decided to step down from the project. The number of mentors is now down to three, which is the quorum. Finding a few more mentors should be a good idea at this point.

The status is being maintained at <http://svn.apache.org/repos/asf/incubator/security/STATUS>

Lucene.Net

Lucene.Net has been in incubation since April 2006

Lucene.Net is a source code, class-per-class, API-per-API and algorithmic port of the Java Lucene search engine to the C# and .NET platform utilizing Microsoft .NET Framework.

Not much changed from the last report:

- Announced Lucene.Net 2.3.1 as "final" (via SVN).
- Port of 2.3.2 is completed; the code is being reviewed and should be announced soon.
- Work is simultaneously underway on the Lucene.Net 2.4.0 port.
- We are in the process of looking into graduation

Olio

Olio has been incubating since September 2008.

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.

- Olio is continuing to gain users and being used for diverse purposes.
 - One user is using it for testing MySQL performance.
 - Another is using it for testing performance of file systems.
- Microsoft has started a .NET port of Olio.
- VMWare has joined the project and plans to use it for several of their performance testing efforts.
- There are multiple deployments attempting to scale Olio to test performance of system infrastructure, software application stack etc.
- A lot of work has gone to make the product more robust and stable over long benchmark runs - over 50 JIRA issues have been fixed.
- Work is underway to create the first release. The artifacts are tested; the major work is in fixing the source headers, licenses etc.
- Olio has been presented to some Apache and Rails Meetups to get attention.
- We're encouraging users to become active in contributing to the code.

Pivot

Pivot has been in incubation since January 2009.

Pivot is an open-source platform for building rich internet applications in Java.

- Noel Grandin was added as a new committer on 3/29.
- The codebase was forked into the "WTKX" project at Google Code: <http://code.google.com/p/wtkx/>. We didn't get a clear answer from the developer as to why he chose to fork rather than collaborate. From what we can gather, he seems to want to do Pivot "his own way".
- The Pivot PMC is currently voting on a proposed 1.1 release. Work has begun on a 1.1.1 maintenance release and 1.2 feature release.

Top 3 issues prior to graduation:

- Successfully complete the 1.1 release.
- Strengthen the understanding of The Apache Way and how to operate a PMC.
- Even stronger community; currently quite good with 5 active committers from 4 different employers.

RCF

Sanselan

2009-January Sanselan Incubator status report Sanselan has been in incubation since September 2007. Sanselan is a pure-java image library for reading and writing a variety of image formats.

The community hasn't grown much in the past six months. There continues to be only one active committer. Participation level is low. A new version has been released

Barriers to graduation continue to be diversity, size of the community, and overall activity. We will address the possible graduation paths in the next quarter and try to evaluate what the best approach for Sanselan might be.

Stonehenge

Stonehenge has been incubating since December 2008.

Stonehenge a set of example applications for Service Oriented Architecture that spans languages and platforms and demonstrates best practices and interoperability by using currently defined W3C OASIS standard protocols.

In the month of March, the growing Stonehenge community has ramped its activity on several fronts, and has made solid plans for near future::

Admin: The donated code from Microsoft and WS02 was moved from contrib to trunk. In the process various installation and configuration issues were identified, and the community has worked very closely on the DL and offline to advance these issues.

Community: 3 SUN developers have joined and have been working through the various installation and configuration issues, and understanding the code architecture. They plan to contribute a Metro based java implementation in the medium term. Synergy has been developing between the WS02 and Microsoft developers towards identifying and resolving common issues.

Code: The community has had discussions about a first release (M1) in the immediate short term. The release is defined to be executable code from the .NET and PHP implementations of the [StockTrader](#) application, along with complete documentation allowing a new Stonehenge member to demonstrate interoperability by mixing and matching components from the two [StockTrader](#) applications. The two applications are also undergoing necessary code changes to qualify for passing the RAT tool (such as adding the legal headers)

Wider Community: Some members and sponsors of Stonehenge have attended [ApacheCon EU](#) and connected with the wider community about Stonehenge.

Tashi

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.

Development activities have included adding support for using a database to track virtual machine information, modifying code to work with Python 2.4, integration with dynamic DHCP and DNS servers, and adding a "tidy" target to check the source code.

There have been a few questions on tashi-user@i.a.o that centered around getting Tashi up and running. Additionally, submissions on the dev list included patches from a student at CMU, and some scripts from a non-committer at Intel. There is still a lot of work to be done in growing a community.

Items to be resolved before graduation:

- Put more effort into project documentation so that other potential contributors may more easily get involved
- Develop community diversity (currently Intel and CMU committers)
- Prepare and review a release candidate

Thrift

Thrift is a software framework for scalable cross-language services development. It combines a software stack with a code generation engine to build services that work efficiently and seamlessly between a variety of programming languages. Thrift entered the Apache Incubator in May 2008.

Recent Activity:

- Bryan Duxbury and Andrew [McGeachie](#) added as committers
- Automatic "instant releases" to simplify the build process
- First official release candidate scheduled for April 7th, 2009
- Finished migrating all source headers to ASF licensing
- Continued language library improvements
 - Support for Twisted in Python
 - Support for .NET 2.0
 - New compact encoding
- Cross-pollination with other projects continues
 - More [PoweredBy](#) entries
 - Cassandra, now in Incubation, uses Thrift
 - Discussions underway about collaboration with the Avro project
 - Hadoop and HBase work ongoing

UIMA

UIMA is a component framework for the analysis of unstructured content such as text, audio and video. UIMA entered incubation on October 3, 2006.

Some recent activity:

- OASIS (<http://www.oasis-open.org>) members approved the Unstructured Information Management Architecture (UIMA) version 1.0 as an OASIS Standard.

- A new contributor, Rico Landefeld, contributed a component to enable indexing annotations as fields using Lucene.
- Many bug fixes and improvements continue to be focused on UIMA-AS (the asynchronous scaleout add-on to base UIMA) driven by experience with users doing extensive scaleout. The fixes address error handling and performance.
- IBM and MAYO Clinic announced a new consortium around making available open-source UIMA annotators for use by the Medical / Health area; see <http://ohnlp.org>

Items to complete before graduation:

- We still need to attract more new committers with diverse affiliations.