

April2010

April 2010 Board reports (see [ReportingSchedule](#)).

THIS REPORT IS CLOSED

These reports were due here by **Wednesday, 14 April 2010** so that the Incubator PMC could 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.
-

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

- Karl Pauls and Marcel Offermans did a tutorial on, amongst other things, ACE at the [EclipseCon/OSGi DevCon](#) in March.
- Marcel Offermans presented ACE at the OSGi UK Users Forum end of January.

Software

- Completed the migration of both the target and server nodes from Ant to Maven, we now have a fully functional Maven build.
- Switched to the newly agreed terminology in the user interface (in the code, this still needs to be done).
- Started a discussion about migrating to a common scheduler API, possibly sharing it with Sling and [ServiceMix](#).
- Toni Menzel did some work on setting up a CI build on Hudson, setting it up and successfully deploying to a Nexus instance using his own setup, still a work in progress on getting it up and running on Apache infrastructure.
- Implemented the option to statically or dynamically link bundles to features in the UI.

Licensing and other issues

- None at the moment.

Things to resolve prior to graduation

- Make a release.
- Grow the community some more.

Signed off by mentor: Carsten Ziegeler, bdelacretaz

BeanValidation

Bean Validation will deliver an implementation of the JSR303 Bean Validation 1.0 specification.

There are currently no issues requiring IPMC or Board attention.

Since Bean Validation entered incubation on March 1, 2010, we have accomplished the following:

1. Initial project resources and accounts created
2. Initial code contribution under SGA from Agimatec GmbH imported into svn
3. Source code package names updated from com.agimatec to org.apache.bval
4. Agimatec copyright moved to NOTICE files and removed from source
5. Three committers are already active making code updates
6. Started setup of our Confluence space as our main website
7. Already have one non-committer using the code and submitting patches

Upcoming major goals:

1. Finish setup of website
2. Setup and usage of Nexus
3. First release of artifacts
4. Start TCK testing

Top 2 or 3 things to resolve before graduation:

1. Build community
2. Create at least one release

Signed off by mentor: Kevan

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.

Recently, we've passed [RealClass](#) release vote in dev mailing list. Now, we are waiting for Bill to check the completeness of the release candidate. Also we are considering to add more feature in the next version of [RealClass](#). The following will be included in the new features.

- optimize DTU structure;
- support IPv6 and satellite;

The coding of IPv6 and satellite module has been finished. But we need some time to test the functionality and robustness of the module. After that we would commit that as new version.

Signed off by mentor:

Chemistry

Apache Chemistry is an effort to provide a Java (and possibly others, like [JavaScript](#)) implementation of the upcoming CMIS (Content Management Interoperability Services) specification. Chemistry entered incubation on April 30th, 2009.

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

1. First formal incubator release. Although this is planned, there are several tasks to complete (e.g. documentation) and a learning curve to climb.

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

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

How has the community developed since the last report

1. The OpenCMIS project (originally proposed to the Incubator which also targets a Java implementation of CMIS) has joined Apache Chemistry, with the following new committers - David Ward, Florian Müller, Jens Hübel, Martin Hermes, Paul Goetz and Stephan Klevenz
2. Jeff Potts has contributed his Python CMIS client library and became a committer.
3. Nick Burch has joined as a new mentor.
4. Mailing list traffic has increased 2.5 times since last quarter.

How has the project developed since the last report

1. Development continues at a steady pace, and Chemistry now targets CMIS 1.0 CD07, the version of the specification submitted to OASIS.
2. Command-line Shell has been contributed (from Nuxeo).
3. Hudson builds have been setup and stabilized.
4. An agreement has been met on how to merge the Chemistry and OpenCMIS codebases. The merge will take place in a branch, until stabilized (which should take no longer than 2 weeks).
5. A first formal incubator release of the merged Chemistry/OpenCMIS codebases is planned shortly after the merge is complete.

Signed off by mentor: nick, gianugo

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.

issues to address in the move towards graduation

Empire-db is mature and seems to be used in many productive environments. The release process has been fully implemented according to Apache conventions. User feedback is positive apart from some complaints about the documentation.

Yet our community is still small and we need to do more advertising in order to attract more users. In order to attract more committers we also need to provide a roadmap that shows where Empire-db will be going in the future and how people can participate.

community development since the last report

During the last three months several new users have had questions or suggestions for improvement. The requests have shown that Empire-db is being used in various different environments (OSGi, Spring) and even together with other programming languages than Java (Scala).

One user donated a new example project that demonstrates the use of Empire-db together with Spring that we will publish with our new release.

project development since the last report

We have just finished working on Release 2.0.6 and are currently seeking approval from the Incubator PMC to publish the release. The release contains various new features, improvements and bugfixes. The most important new feature is a reverse engineering component that generates data model mapping files from existing databases either from command line or by a Maven plugin.

Signed off by mentor:
[ThomasFischer](#), Dashorst

Imperius

Imperius has been incubating since November 2007.

Imperius is a rule-based policy evaluation engine based on the CIM-SPL language from Distributed Management Task Force (dtmf.org).

The voting for the first release of Imperius was passed on Jan 18, 2010. The release was initially made available on April 12, 2010 via the Imperius website and mirrors.

Communication continues to be intermittent. However, the community is discussing the next steps for Imperius. The implementation of the CIM-SPL standard appears fairly complete and so, although the community is currently small, we would like to consider promoting Imperius as a formal Apache project.

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

1. Grow community
2. Javadoc needs improvement

How has the community developed since the last report?

Still limited, although we have confirmation that we have 3 independent committers.

How has the project developed since the last report?

Limited work other than release preparation. Some additional bugs have been reported. Currently discussing the need for possible additional features.

Signed off by mentor: Craig L Russell

JSPWiki

JSPWiki has been incubating since September 2007.

JSPWiki is a JSP-based wiki program.

Quite a few bugs were fixed in the JSP tier. Slimbox was upgraded, various media formats are now supported so you can view youtube or facebook videos, other wikipages, or even external web-pages. Also support was added for multi-file uploads. We also added an experimental wysiwyg editor based on mootools.

[ReferenceManager](#) was rewritten to use JCR UUIDs to keep references between pages.

A couple of bugs were fixed and, and the unit test compliance has climbed from 96.8 to 99.3 %.

Of the 28 items (1 was recently added) on the graduation checklist, 19 are complete, so in that area there has been not much progress, with the exception of fixing junit tests.. All open items are documentation and ASF process and infrastructure related. There still is no 3.0.0-incubating-alpha1 release.

It looks like things have slowed down the last few weeks for an unknown reason.

The developer list currently has 88 members, an increase from 82; and the user list has 193 members, an increase from 191.

Signed off by mentor: Craig L Russell

Lucene Connectors Framework

Description

Lucene Connectors Framework is an incremental crawler framework and set of connectors designed to pull documents from various kinds of repositories into search engine indexes or other targets. The current bevy of connectors includes Documentum (EMC), [FileNet](#) (IBM), [LiveLink](#) (OpenText), Patriarch (Memex), Meridio (Autonomy), [SharePoint](#) (Microsoft), RSS feeds, and web content. Lucene Connectors Framework also provides components for individual document security within a target search engine, so that repository security access conventions can be enforced in the search results.

Lucene Connectors Framework has been in incubation since January, 2010.

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

1. End-user documentation needs to be converted into a usable form; this is probably the biggest obstacle to developing a broader community at this point
2. Javadoc and nightly builds need to be set up
3. The first official release needs to be planned and executed

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

1. We'd like to know whether there is any official Apache position on inclusion of NTLM implementations in ASF projects, since we've gotten mixed signals on this from other developers. This represents a crucial piece of functionality needed to support [LiveLink](#), Meridio, [SharePoint](#), RSS, and Web connectors properly.

How has the community developed since the last report?

We've received several queries and comments from non-Apache developers recently. This is a good sign. There is also a Eurocon conference in Prague which will include LCF, where we will have an opportunity to introduce the project to a broader Lucene community.

How has the project developed since the last report?

The LCF site has been fleshed out, and much more extensive developer documentation has been written and linked into the LCF site. Preparations have also been made to pull appropriate build and execution dependencies in using Ivy or Maven. Javadoc is now available to people willing to build the project themselves. The project remains buildable and usable.

Signed off by mentor: Grant Ingersoll

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.

We have so far put out two releases successfully. Most users are now using the 0.2 release. Except for one issue with the Java version, this release seems to be stable. Several new developers are actively working on the Java version to improve and expand on it's functionality and robustness.

Olio seems to be the workload of choice for testing virtual machines. Several researchers as well VMware are using Olio for this purpose.

Graduation From Incubation:

Diversity of committers is the primary issue with the project - although we have users, we haven't been successful in converting them to committers (yet).

We could use the PMC and Board's help in spreading the word about Olio to get better traction.

Signed off by mentor: Craig L Russell

OODT

Description

OODT is a grid middleware framework for science data processing, information integration, and retrieval. OODT is used on a number of successful projects at [NASA's Jet Propulsion Laboratory/ California Institute of Technology](#), and many other research institutions and universities.

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

1. Port OODT code and license headers into ASF license headers
2. OODT contributions from at least 2 other organizations besides JPL
3. At least one OODT incubating release, hopefully in the first few months

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?

Dave Kale, from Children's Hospital Los Angeles (CHLA), is the first external to JPL contributor to the project (besides mentors of course). Dave contributed the patch for [OODT-8](#), fixing a minor bug in referencing a jar dependency for Maven. Cameron Goodale, another JPL'er (but not a committer), input [OODT-14](#) for documentation that he is working on. Much of the other activity continues to be from the mentors and committers.

How has the project developed since the last report?

OODT was [voted into the Incubator by the IPMC](#) on January 22, 2010.

Sean Kelly imported the Python version of the OODT query, profile, product, and webgrid components in [OODT-6](#). Work on the initial import into Apache ([OODT-3](#)) is nearing completion. Sean McCleese and Andrew Hart have been leading the way. Committers have begun logging new issues in Jira and using the Apache SVN for their current development efforts (e.g., see the efforts from Brian Foster in [OODT-10](#), [OODT-11](#), [OODT-12](#) and [OODT-13](#)). Brian Foster also initiated [mailing list discussion](#) regarding [OODT-15](#) which proposes to create one top-level build for the OODT components and a versioning scheme for the software.

Signed off by mentor:

Shiro

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

Shiro has been incubating since June 2008.

The project is just about ready for its first 1.0 release. Cryptography API and implementation adjustments had to be made prior to the 1.0 release, delaying the 'code complete' stage before incubator vote by 2 weeks. That effort is being finished this week.

As soon as this code is complete, and we resolve 4 outstanding Jira bugs, we will go immediately initiate the voting process to clear our first 1.0 release (hopefully next week).

The project team is not considering graduation at this point, but after the first release, the team will decide on a roadmap targeting graduation.

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

Signed off by mentor: Craig L Russell

SIS

Apache SIS is a toolkit that spatial information system builders or users can use to build applications containing location context. This project will look to store reference implementations of spatial algorithms, utilities, services, etc. as well as serve as a sandbox to explore new ideas. Further, the goal is to have Apache SIS grow into a thriving Apache top-level community, where a host of SIS/GIS related software (OGC datastores, REST-ful interfaces, data standards, etc.) can grow from and thrive under the Apache umbrella.

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

Not at this time

Community progress since the last report

Chris Mattmann was contacted by an [ESRI](#) representative inquiring about the direction of the project. This lead will be followed as garnering support from ESRI would be big win for the project and attract a large community.

Project progress since last report

SIS was [voted into the Incubator by the IPMC](#) on February 21, 2010.

Chris Mattmann has finished importing the LocalLucene code into the SIS codebase ([SIS-1](#)) and Patrick O'Leary has compiled around [~50,000 geocoded records](#) for development/testing of SIS spatial functions. Work on creating the SIS incubator website ([SIS-2](#)) has begun led by Sean McCleese as well.

Signed off by mentor: Kevan, greddin

Socialsite

[SocialSite](#) has been incubating since May 2009.

[SocialSite](http://rollerweblogger.org/roller/entry/oracles_social_site_promise) has been stalled for quick a long time, but it appears that there is some new movement at Sun/Oracle. There are some details here: http://rollerweblogger.org/roller/entry/oracles_social_site_promise. If there is no movement by the next report, we should consider mothballing this project.

Three most important issues to address for graduation: 1. Get source code grant from Sun, import to SVN. 1. Migrate codebase: repackage, work out bad deps, etc. 1. Learn Apache way: demonstrate that we have a community

[SocialSite](#) is an open source Social Networking Service based on Apache Shindig (incubating). The software is not simply a "canned" Social Network or Facebook-in-a-box type of web application; it's something different. [SocialSite](#) is designed to add social networking features to existing web applications and web sites. [SocialSite](#) is made up of two parts: a social data server that supports the [OpenSocial](#) APIs and extensions and a set of [OpenSocial](#) gadgets that provide a complete user-interface for social networking.

Signed off by mentor: Dave Johnson

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.

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 fixes to conform to new python programming standards, and a module for Zoni to assign ports on HP blade server switches.

The project is still working toward building a larger user and development community. Michael Ryan, an active committer on the project, has taken a new job and is unable to actively contribute to the project any longer. Richard Gass, who is running a Tashi production environment, has been added as a committer. Richard introduced the Zoni physical hardware management layer to Tashi earlier.

Items to be resolved before graduation:

- Prepare and review a release candidate
- Develop community diversity (currently Intel and CMU committers)

Signed off by mentor: Craig L Russell

Traffic Server

Traffic Server is an HTTP proxy server and cache, similar to Squid and Varnish (but better). Traffic Server has been incubated since July 2009.

Recent activities:

- 2010-03-30 The PPMC has begun the graduation process.
- 2010-03-29 The new home page is launched.
- 2010-03-17 Diane Smith joins the Traffic Server PPMC.
- 2010-03-13 Apache Traffic Server v2.0.0-alpha is released.
- 2010-03-04 The community votes for CTR for trunk, RTC for release branches.
- 2010-03-02 Manjesh Nilange joins the Traffic Server PPMC.
- 2010-02-26 Manjesh Nilange joins the project as a new committer.
- 2010-02-23 2.0.x release branch created, and CI environment setup.
- 2010-02-09 The last RAT issues are resolved, we're clean.
- 2010-02-02 KEYS file added to dist area.
- 2010-02-02 Automatic sync from SVN dist repo to dist servers setup.
- 2010-01-18 George Paul joins the Traffic Server PPMC.

The graduation process is completed, and we've passed the votes in both the PPMC and the IPMC. A resolution proposal has been submitted to the board for the next board meeting.

Signed off by mentor: Jean-Frederic Clere

Thrift

(project add text here)

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:

- more progress towards running the complete XQTS (XQuery Test Suite)
- first signs of community interest

Top issues before graduation:

- Build community
- Create a release

Signed off by mentor: