

# May2011

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

These reports were due here by **Wednesday, 11 May 2011** so that the Incubator PMC could 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.

```
-----
Amber

Amber has been incubating since July 2010. Amber is a project to develop a Java library which provides an API
specification for, and an unconditionally compliant implementation of the OAuth v1.0, v1.0a and v2.0
specifications. OAuth is a mechanism that allows users to authenticate and authorise access by another party to
resources they control while avoiding the need to share their username and password credentials.

The most important issues that must be addressed before graduation are:

- Attract users and developers
- Generate a release

The Incubator PMC / ASF Board should be aware that:

- Community activity is relatively low

How has the community developed since the last report

- Some users have started asking for help or getting started guides on the mailing list

How has the project developed since the last report

- Fixed some bugs
- Removed LGPL licensed files from website

Signed off by mentor: Stefan Seelmann

-----
Bluesky

(project add text here)

Signed off by mentor:

-----
Clerezza

(incubating since November 27th, 2009) is an OSGi-based modular application and set of components (bundles) for
building RESTful Semantic Web applications and services.

There are currently no issues requiring board attention.

Recent activity:

- Added WebId test suite
- Added Clerezza-UIMA CAS Consumer
- Released a new website version with much more documentation

Next steps:

- First release
```

- Improve documentation and tutorial to website

Top 2/3 Issues before graduation:

- Prepare some easy-to-run demos to get people interested in Clerezza

Signed off by mentor: rgardler

-----  
Deltacloud

(Incubating since July 2010)

Deltacloud is a cross cloud abstraction API that allows management of resources (compute and storage) across different IaaS cloud computing providers using a single, open API. At present Deltacloud supports management of Amazon EC2, Rackspace, Gogrid, Rimuhosting, Terremark, RHEV-M, IBM SBC, Eucalyptus, Opennebula compute resources, as well as Amazon S3, Microsoft Azure, Eucalyptus and Rackspace Cloudfiles storage resources.

Recently released v 0.3.0 - main changes for this release:

Server:

- \* Dynamic driver switching: select driver on a per-request basis; new toplevel 'drivers' collection describing drivers supported by server
- \* Create images from running instances (EC2, Mock, GoGrid, Rackspace); advertised as action 'create\_image' in instance details when possible
- \* New 'user\_files' feature for create\_instance to advertise RAX-style injection of user data
- \* Return status 204 after successful DELETE operation
- \* Return status 401 when authentication fails because of invalid credentials
- \* Blobs: support user metadata as key/value pairs passed through X-Deltacloud-Blobmeta-KEY: VALUE headers
- \* Support HEAD requests to retrieve the operations and methods supported by a collection
- \* Support for OPTIONS request to retrieve optional and required parameters for operations
- \* Advertise 'create\_instance' action for each image
- \* Drivers
  - + EC2
    - instance\_count feature to allow creating multiple instances at once
    - run commands inside an instance via ssh
    - by default, list images owned by 'amazon', when passing in empty owner\_id, list \_all\_ images (thousands)
  - + Eucalyptus
    - new driver for Eucalyptus (Sang-Min Park)
  - + Gogrid
    - run commands inside an instance via ssh
    - allow creating sandbox instances
  - + Rackspace
    - report root password after instance creation
  - + SBC
    - new driver for IBM SBC cloud (Eric Woods)

Client:

- \* run: new method to run commands via ssh
- \* drivers: list drivers supported by server
- \* properly list blobs in a bucket when showing bucket details
- \* full support for managing blobs and buckets

Signed off by mentor:

-----  
Droids

Droids is an Incubator project arrived from Apache Labs. Droids entered incubation on October, 2008.

It's an intelligent standalone robot framework that allows one to create and extend existing web robots.

The last three months have been very productive for the project since mailing and issue tracking activity have been very good. Our two new committer have driven our first release which is planned for this month. More people appeared on the mailing list and especially one dev is driving constantly enhancement in our issue tracker.

Issues before graduation :

- \* IP clearance

Signed off by mentor: Thorsten Scherler

-----

HCatalog

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

The most important issues in moving the project to graduation are expanding the community of developers and producing a release of the software.

Since the last report we have:

- \* got the initial code drop checked in
- \* got the build working for building, testing, and the documents
- \* branched in preparation for a 0.1 release
- \* started development on new features for 0.2

Currently there are 29 subscribers to the user list and 27 on the dev list. There were 14 and 15 respectively last month.

Signed off by mentor: Alan Gates

-----

HISE

(project add text here)

Signed off by mentor:

-----

Jena

Jena was accepted into the Apache Incubator November 2010. 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 have received the software grant from HP; HP has copyright on the majority of the codebase. We can now migrate the code to Apache infrastructure, and contact previous contributors for the remaining areas (these are small in number and scope).

Signed off by mentor:  
rgardler

## ----- Libcloud

Libcloud is a unified interface into various cloud service providers, written in Python. Libcloud joined the Incubator on 3 November 2009.

During this month we have addressed all of the open graduation issues and started a TLP voting thread on the incubator mailing list.

## Community -----

The Libcloud community continues to grow at a nice steady pace. In the past few months we have received multiple community contributions, some of them are listed bellow:

- compute driver for Gandi.net
- compute driver for Bluebox.net
- compute driver for the OpSource cloud
- compute driver for OpenStack (not merged yet, still working with the original author on polishing it and resolving some open issues)
- libvirt proof-of-concept compute driver (not merged yet, still debating on how we can fit it into our project)

During this month (May) we have also agreed upon a new chair (Tomaz Muraus).

To make our website more easy to update and more relevant to the users we have ported it to the Apache CMS and added two new sections - "Community resources" and "Contributing".

Some of the Libcloud developers have also been present at PyCon in Atlanta in March where we have held a sprint.

As part of our ongoing effort to engage the community and inspire new contributors, one of the committers will give a talk about Libcloud at EuroPython 2011 in Italy where we will also held a sprint.

## Releases -----

The last stable release of Libcloud is 0.4.2 which was released on January 17th. We are currently working towards 0.5.0 which is considered as a big milestone, because it will include many new features, fixes and improvements:

- new Storage API
- new Load-balancer API
- new compute drivers
- fixes in the existing compute drivers
- increased test coverage

This version is planned to be released before the end of May.

Signed off by mentor:

## ----- 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. Lucene.Net has three primary goals:

- \* Maintain the existing line-by-line port from Java to C#, fully automating and commoditizing the process such that the project can easily synchronize with the Java Lucene release schedule.
- \* High-performance C# search engine library.
- \* Maximize usability and power when used within the .NET runtime. To that end, it will present a highly idiomatic, carefully tailored API that takes advantage of many of the special features of the .NET runtime.

#### Recent Activity:

- \* Lucene.Net 2.9.4 is in the trunk and currently being tested
- \* Continuous Integration, Unit Testing, and Code Coverage has been fully vetted

#### Current Activities:

- \* Release Lucene.Net 2.9.4
- \* Refactoring of project files for the main trunk, allowing us to fully implement CI, 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

-----

#### 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.

#### Project development

- \* We are still working on importing the source code to Apache, after having merged in some major features this month (e.g. protocol buffers for forward-compatible messaging, and improved fault recovery features).

#### 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: tomwhite

-----

#### MRUnit

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

MRUnit entered incubation on March 8th, 2011.

#### Community

- \* Eric Sammer has been leading discussions for a 0.5 release and a release roadmap beyond that.

#### 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

-----

#### NPanday

NPanday allows projects using the .NET framework to be built with Apache Maven. NPanday allows .NET projects to be converted into Maven projects thus allowing them to fully utilize the other technologies driven by Maven. NPanday has been incubating since August 2010.

In our previous report (February 2011), we stated that we've staged the NPanday-1.3-incubating release. However, there were a couple of issues that we needed to address so we decided to not push through with it.

After addressing those issues, we've staged the NPanday-1.4.0-incubating release for vote which includes the following issues:

- \* Visual Studio 2010 and .Net 4.0 Framework support
- \* Removal of UAC and PAB directories
- \* Updated groupId of all NPanday components including plugins to "org.apache.npanday"

This release is still waiting for other committers and mentors to test and get the required binding votes.

Updated documentation was successfully deployed in <http://incubator.apache.org/npanday/docs/1.4.0-incubating/index.html>

Code licenses and license headers have all been reviewed and pass RAT.

IP clearance was completed.

We have voted a new committer in the person of John Fallows and we also have a new mentor, Matthias Wessendorf.

The top priorities towards graduation are:

- \* complete 1.4.0-incubating release
- \* guide regular patch contributors towards becoming committers
- \* improve engagement of existing committers

There are no issues for the Incubator PMC or board at this time.

Signed off by mentor: Dennis Lundberg

-----

Nuvm

Apache Nuvm will define an open application programming interface for common cloud application services, allowing applications to be easily ported across the most popular cloud platforms.

Nuvm was accepted for Incubation on June, 2010.

- \* John Pradeep elected as Nuvm Committer
- \* Pasindu Madubashana starting to work on Data Services project as part of GSoC 2011
- \* Community has started to be more active, with about 3 to 4 active committers + the new GSoC student.

Top things to resolve prior to graduation:

- \* Increase the number of active committers.

Signed off by mentor:

-----

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 just joined the Incubator on April 26, 2011.

What is needed to start, is resolving all the Issues enlisted on INFRA-3597, resolution is in progress, actions have been completed are:

- \* PMC ML creation in discussion;
- \* Discussed position of PPMC after graduating into Commons;
- \* Created podling status page;

What is still in progress, instead:

- \* Grant Karma to SVN;
- \* Import codebase;

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 just did our first Apache release on 2nd of May and are now planning which features/work will go into the next release. For instance we started to discuss/plan a big refactoring of our machine learning code, custom feature generation for the name finder, dictionary support for the name finder, etc.

There is still good traffic on both the user and dev mailing list. A couple of users asked questions about how OpenNLP can be trained for new languages.

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

- \* Resolve potential IP issues around releasing training models
- \* Establish open regression tests for the parser and coreference component

Signed off by mentor: Grant Ingersoll

-----

#### PhotArk

Apache PhotArk will be a complete open source photo gallery application including a content repository for the images, a display piece, an access control layer, and upload capabilities.

PhotArk was accepted for Incubation on August 19, 2008.

- \* Subash Chaturanga elected as Photark committer
- \* Subash Chaturanga started working on "Face Recognition" project as part of GSoC 2011
- \* Umashanthi Pavalanathan started working on making Photark more "Social" based on Apache Shindig as part of GSoC 2011

\* Issues before graduation :

\* PhotArk started as a project with no initial code-base, and we have grown the community to the minimal 3 independent committer size required for graduation and have been seeing slow but continuous interest in the project. The community started discussion graduation and based on mentors feedback we are going to recruit couple committers before we start the graduation process.

Signed off by mentor: lresende

-----

#### 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 entered incubation on 2011-03-01.

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

Current Status:

- \* The website has been published and continues to be updated regularly
- \* The community voted on and accepted a logo
- \* Issues have been added to JIRA and a release plan is beginning to form
- \* The community has begun working on implementation features and test cases for the 0.1 incubator release within the project build structure.
- \* Community members have submitted proposals for ApacheCon to discuss the Rave project
- \* Mailing list activity remains high

Next steps:

- \* Continue to build up awareness of Rave and grow the community

- \* Complete a sufficient number of features for a 0.1 incubator release

Issues before graduation:

- \* Complete 1.0 release
- \* Expand the community/user base

Signed off by mentor: ate, rgardler

-----  
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

Since SIS released its first release 0.1-incubating in November 2010, development has slowed down. This is a small core set of committers and community though. We haven't had a lot of user mailing list activity.

We should try and target a 0.2-incubating release with a layer service and some basic integration with Apache Tika (a GDAL parser is being worked on in that community with an eye towards integration into SIS). Also we are looking at Apache OODT as a mechanism for geospatial storage and persistence.

There has been renewed interest from ESRI and we are pursuing this as well.

- \* Project progress since last report

SIS was voted into the Incubator by the IPMC on February 21, 2010.

Development progress since the 0.1-incubating release has been modest, focusing primarily on the community development mentioned above. There has been some internal effort at JPL to integrate SIS into ongoing projects such as the CO2 Portal, the CMDS project and others. We hope to begin the push towards 0.2-incubating as specified above.

Signed off by mentor: mattmann, greddin, kevan

-----  
Stanbol

Stanbol is a modular software stack and reusable set of components for semantic content management. Entered incubation on 2010-11-15.

Community development:

- some new users started playing with Stanbol and asking for guidance on the mailing lists

Project activities:

- lots of package/module refactoring
- added a benchmark tool for Enhancer
- KReS module build and tests fixed
- setup continuous integration with Jenkins
- added JCR module of CMS Adapter

Next steps:

- Make a first release, possibly omitting modules with ASF-incompatible dependencies
- Grow the community

Signed off by mentor: Tommaso Teofili

-----  
Stonehenge

(project add text here)

Signed off by mentor:



-----

## VCL

VCL has been incubating since December 2008. VCL is a cloud computing platform for the management of physical and virtual machines.<https://cwiki.apache.org/VCL/>

### Community Involvement:

- \* As stated in the previous report, the Apache VCL community continues to grow. We're seeing more people start to contribute to our online documentation and file JIRA issues. We're also starting to see more patches attached to JIRA issues.

- \* We continue to see growing interest in the VCL project from a variety of locations.

### Releases:

- \* We cut a new release that was primarily for bugfixes. Apache VCL 2.2.1 was released on April 8th, 2011. In our last report, we stated that we had hoped to get it out by the end of February. Other duties of the developers caused some of the delay. Additional delay was due to a bug found after the release voting took place, requiring a new round of voting.

### Top Issues Before Graduation:

- \* Continue to increase contributors to establish a more diverse development community

Signed off by mentor: kevan

-----

## Whirr

Whirr is a library for running services in the cloud.

Whirr was accepted into the incubator on 11 May 2010. Status information is available at <http://incubator.apache.org/projects/whirr.html>.

Progress since the last report (February 2011): We have made the fourth incubator release (0.4.0-incubating). We continue to get contributions from outside the initial committer base, and development has been steady over the period. Whirr continues to be covered at community events, including an upcoming presentation at OSCON.

### Plans for the next period:

- \* Encourage community growth.
- \* Do a fifth release.
- \* Strengthen the core API to make adding new services easier.

### Top three items to resolve before graduation:

- \* Increase community involvement in the project
- \* Make several incubating releases (Done)
- \* Support at least three services on Whirr (Done, there are now six)

Signed off by mentor: tomwhite

-----

## Wink

Apache Wink is a project that enables development and consumption of REST style web services. The core server runtime is based on the JAX-RS (JSR 311) standard. The project also introduces a client runtime which can leverage certain components of the server-side runtime. Apache Wink will deliver component technology that can be easily integrated into a variety of environments.

Apache Wink has been incubating since 2009-05-27.

### Notable Activity:

- \* Luciano Resende and Raymond Feng elected as Wink committers
- \* Luciano Resende elected as PMC member
- \* Released Apache Wink 1.1.3.
- \* The wink community has added a couple new binary Providers and bug fixes.

### Planned Activity:

- \* JAX-RS 2.0 is starting up in the JCP. May attempt to prototype/develop against the next specification to provide an eventual JAX-RS 2.0 compliant runtime.
- \* Work towards graduation.

Top issues before graduation:

\*

Signed off by mentor: kevan

-----

Wookie

Wookie is a Java server application that allows you to upload and deploy widgets for your applications. Wookie is based on the W3C Widgets specification, but widgets can also be included that use extended APIs such as Google Wave Gadgets and OpenSocial.

Wookie has been incubating since 17th July 2009

Since February 2011:

- \* Release Candidate 0.9 put out for testing and feedback from wookie-dev members
- \* Feedback and issues for Release Candidate 0.9 posted to JIRA
  - \*\* Almost all issues closed, new RC build imminent
- \* Planned event for March took place
  - \*\* Demonstrated how to make widgets for wookie, use of the APIs
  - \*\* Promoted how to get more involved with Wookie, rather than simply using it
  - \*\* New participants at event posted issues/patches to JIRA for first time
- \* Various bug fixes and updates to main code branch
- \* Community members working with Apache Rave (incubating) to explore collaboration options
- \* Community members exploring Apache Amber (incubating) for OAuth support within wookie

Next steps:

- \* Complete 0.9 build
- \* Continue to work and liaise with Apache Rave project
- \* Continue exploration of using Apache Amber for OAuth support within wookie

Issues before graduation:

- \* complete first release
- \* build community/expand user base

Signed off by mentor: rgardler, ate

-----

Zeta Components

Apache Zeta Components are a high quality, general purpose library of loosely coupled components for development of applications based on PHP.

What has been done:

- All INFRA tasks have been implemented
- Integrated two new committers
- Fixed bugs

What is still in progress

- Implementation of a build system in order to replace the current one
- Discussion about a roadmap for a release
- Discussing about how to do "marketing" and make the project more visible

What needs to be done:

- Fix more bugs
- Get a release

Signed off by mentor: grobmeier

-----