

# April2012

April 2012 Board reports (see [ReportingSchedule](#)).

This report is **CLOSED**.

The biggest Incubator news since our last report is the CloudStack project seeking to enter the ASF. This and other topics as summarized below have kept us busy lately. At times the related discussions have gotten a bit heated, but for the most part the Incubator is functioning acceptably.

The overall trend currently is for more podlings to graduate than are coming in as newly proposed projects. Given the backlog of stuck projects and the recent focus on pushing them forward this trend can be expected to continue at least for the next few months.

## o Community

Chris Hostetter, Leo Simons, Michael McCandless, Ted Leung and William Rowe resigned from the Incubator PMC since our last report.

The following podlings are requesting graduation to Apache TLPs:

- Apache Creadur
- Apache Jena

The Incubator PMC recommends the board to accept the respective resolutions.

We received a proposal to accept the CloudStack project as a new podling. The proposal has received quite a bit of excitement and many volunteer mentors. On the other hand there were some concerns about the publicity around the proposal (current Incubator policy prohibits publicity-seeking in the proposal phase). Another issue of note is the handling of patents that was raised as a potential issue for the proposed podling. The vote to accept CloudStack for incubation is currently in progress.

## o Legal / Trademarks

The Flex podling continues working on a custom trademark licensing deal with Adobe. See the Flex report for more details.

A point was raised about including the project branding checklist in podling graduation criteria. This seems like a useful change, though the details are still to be worked out.

## o Releases

The following incubating releases were made since our last report:

- March 12th, 2012: Apache Tashi 201203-incubating
- March 20th, 2012: Apache Rave 0.9-incubating
- March 23rd, 2012: Apache Jena Fuseki 0.2.1-incubating
- March 27th, 2012: Apache Flume 1.1.0-incubating
- April 2nd, 2012: Apache Bigtop 0.3.0-incubating

Apache OpenOffice used Apache mirrors to distribute binary security patches for fixing the CVE-2012-0037 issue in OpenOffice.org 3.3 and 3.4 beta. Even though the /dist/incubator area was used, this wasn't considered an official Apache release and the distribution was not formally blessed by the Incubator PMC. See the OpenOffice report about the extraordinary circumstances surrounding the patch.

A notable discussion came up in relation to a release candidate by the ManifoldCF podling. Based on more or less standard practice in many Ant-based Java projects in and outside the ASF and related mentor guidance the podling was including libraries from upstream projects as binary dependencies inside the source release package. It was pointed out that an Apache release must contain nothing that isn't included in source form, and rough consensus from the

resulting discussion agreed with this view. The ManifoldCF release candidate was withdrawn and is being revised for this, but the broader issue of many Apache projects shipping binary dependencies inside their source releases still remains. Board guidance on how projects should deal with this regarding both current and past releases would be appreciated.

Followup discussion and related feedback from podlings suggests that our existing release documentation is not up to the task of providing clear and consistent guidance to support our projects. More work in this area is clearly needed.

#### o Infrastructure

The Flex project continues to have trouble with their migration to Apache infrastructure as explained in more detail in their report. A suggestion was made that, if possible, they could/should use external infrastructure to avoid blocking development until the migration issues have been resolved. This is in line with other podlings that have migrated their infrastructure to Apache gradually in smaller steps.

More generally with complex podlings like OpenOffice, Flex and the proposed CloudStack project stressing infra resources and the steady flow of various Incubator-related infrastructure issues (especially around podlings entering and existing the Incubator) we feel that the Incubator is putting quite a bit of load on the infrastructure team. It would be great if the foundation could help through increased funding or other help to infrastructure.

One proposal that was discussed for helping the life of both podlings and infrastructure was about allowing new podlings to set up their infrastructure directly at the expected final TLP locations. The details of that proposal are still open.

Due to spamming the Incubator wiki was configured to require explicit authorization of new contributors.

#### ----- Summary of podling reports -----

We reviewed all podlings reporting in this quarter and categorized them according to their progress through the Incubator and the most pressing issues that are currently blocking progress.

##### Still getting started at the Incubator (2 podlings)

DeviceMap, Syncope

These projects are still getting started, so no immediate progress towards graduation is yet expected.

##### Not yet ready to graduate (17 podlings)

IP clearance: -

No release: Any23, Cordova, DirectMemory, Flex, JSPWiki, Mesos, OpenOffice

Low activity: AWF, Celix, EasyAnt, Kitty, VXQuery

Low diversity: Chukwa, Kafka, ODF Toolkit, Oozie, Tashi

We expect the next quarterly report of projects in this category to include a summary of their actions and progress in solving these issues.

##### Ready to graduate (1 podlings)

Giraph

We expect these projects to graduate within the next quarter.

#### ----- Detailed Reports -----

##### Any23

Anything To Triples (any23) is a library, a web service and a command line

tool that extracts structured data in RDF format from a variety of Web documents. Currently it supports the following input formats:

- RDF/XML, Turtle, Notation 3
- RDFa with RDFa1.1 prefix mechanism
- Microformats: Adr, Geo, hCalendar, hCard, hListing, hResume, hReview, License, XFN and Species
- HTML5 Microdata: (such as Schema.org)
- CSV: Comma Separated Values with separator auto-detection.

Any23 was voted into the Incubator by the IPMC on October 1, 2011.

The community has seen steady levels of traffic with the dev list receiving 281 hits in March; an increase from the previous month.

One very interesting thread related to the webdatacommons.org project, which recently used Apache Any23 to extract embedded structured data from web pages. The project scales to cover several billion web pages, so we are very confident that the Any23 community is growing interest, this news reflects that view.

We were recently contacted by the W3C's Semantic Web Activity Lead (Ivan Herman), and have been invited to provide an EARL report of Any23's test results with regards to the RDFa test suite, so that we could include it into <http://rdfa.info/earl-reports/> We have been given until the end of April to submit.

We are very close to VOTE'ing on the 0.7.0-incubating release (1st during incubation), and Simone Tripodi has stepped up as release manager. We are currently discussing when to push for the RC.

Chris Mattmann is investigating the use of Any23 as one potential component in implementing GeoSPARQL, along with Jena, Tika and Apache SIS.

Signed off by mentor: mattmann

-----  
AWF

AWF is a non-blocking, asynchronous, event driven high performance web framework running on the JVM. The project entered incubation on 2011-07-08.

The project has been renamed from DEFT to AWF, all the mailing-lists, web site, JIRA etc have been migrated.

We have had few activity this quarter:

- Still working on getting our first release out
- Some bugs have been fixed

How has the community developed since the last report?

- On one hand Roger does not have any more time to work on the project
- But on the other hand there is a new potential committer. His recommendation is being under discussion

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

- Not at this time.

How has the project developed since the last report?

- Still working on getting our first release out

NOTE: It was questioned if this project was viable as the activity remains low, but the team thinks that it's a temporary situation due to individuals being swamped into day job and change in family situation, so the decision to pursue the effort has been agreed on.

Signed off by mentors: mnour, elecharny

-----  
Cordova

Apache Cordova is a platform for building native mobile applications using HTML, CSS and JavaScript. The project entered incubation as Apache Callback in October, 2011, before changing its name to Cordova.

- continued code migration to Cordova namespace
- documentation updated for migration inc getting started guides
- cordovajs migration for: ios, android, blackberry complete;
- cordovajs migration remains for: wp7, bada, webos and qt
- Shaz crushed crazy iOS local storage bug; his solution getting a tonne of attention in iOS community
- Tim Kim voted as committer
- Intel contribution of initial Tizen src identified

Graduation concerns:

- official Apache release remains to be verified

Signed off by mentor: jukka

-----  
Celix

Celix is an implementation of the OSGi Specification in C.

Celix entered incubation on November 2, 2010.

The last month we received our first large code donation, the code still has to be added to the project. This code is an implementation of the OSGi Device Access specification for Celix and has been made by Thales Netherlands. Together with this donation a new committer (Pepijn Noltes) is accepted. Pepijn has developed and will maintain the Device Access code.

We have also been working on a graduation plan which is included below.

Most important issues are:

- Improve robustness (APR, error handling etc), resulting in a first release
- Update/Implement remote services for interoperability with Java OSGi (Apache Felix)
- Generate awareness and grow a community!

Graduation Plan

Celix is in incubation since November 2010. During the first one and a half year talks were given at several conferences (EclipseCon, ApacheCon, OSGi User Group meetings, etc). Even though there seems to be an interest in the project, two important questions keep coming up:

- What is the state of the project?
- Why no support for C++?

Trying to answer/solve these two questions might make it able to attract more community members. So this plan will focus mostly on these two items.

State of the project

- Releases

Celix entered incubation in its early stage. There was only a proof of concept, but no complete implementation. This is an important reason for people to hold back and not yet use/improve Celix, on the other hand, being hesitant also keeps Celix from growing towards a more stable/robust solution. To be able to use Celix the implementation has to reach, at least, a more stable state. Over the past year lots of effort has been put into this. Within the next half year a release has to be made of the core component of Celix. Hopefully this will attract more users/testers (and potentially committers). Since a formal release takes quite some effort, it might also make sense to provide snapshots (with documentation)

to be able to reach more people.

#### - Committers

During the last months there has been an interest from Thales Netherlands to use Celix in its middleware. In a research project they are working on an implementation of the Device Access specification. This implementation is donated to Celix, and the main developer has expressed the intention to maintain the code base. Via this path a new committer has been added to Celix [1][2]. But to be able to have a diverse community more committers are needed. Having a release makes it easier for people to use and improve Celix. This is one step towards more committers.

[1]: <http://markmail.org/thread/hfwuspl3swoiseng>  
[2]: <http://markmail.org/message/q4n7562jvngd33s5>

#### - Technical state

One of the important aspects of Celix is interoperability with Java OSGi through remote services. Currently Celix has basic support for Celix to Celix remote services, following the Remote Service Admin specification of OSGi. This implementation has to be improved and extended to comply better to the specification. Also a Java OSGi implementation has to be made which can interact with the Celix implementation. Some existing open source solutions are available, but are either too large for our intended target platforms or rely on too many other libraries (for example XML handling etc). To be able to have an implementation which fits the environment ((de)serialization and protocol) it makes sense to implement a simple solution ourselves. Having functional remote services makes it easier to use Celix in a mixed Java/C environment. This solution can also be positioned as an alternative to JNI with the benefit that the Java and C components are separate processes. If either one crashes the other part is kept running, resulting in a more robust solution.

### C++ Support

#### - Technical Scope

Currently Celix is limited to C only. This was a deliberate choice since Celix tries to target embedded/constrained platforms. But during talks people also seem to be interested in C++ support. Extending the technical scope of the project might attract more users and committers. Over the next half year we will work out a plan how C++ support can be added without impacting the current supported platforms. A start with the discussions has been made on the mailing list, see [2] for more information.

[3]: <http://markmail.org/thread/a3qltqhsocmrnerd>

#### - Cooperate with existing C++ OSGi like implementations

In [3] a list of similar projects is mentioned. Reaching out to these projects and trying to find a common ground on requirements/API etc could benefit Celix (and those projects as well). To see if there is a common ground we need to contact those projects and plan a meeting.

Signed off by mentor:

-----  
Chukwa

Chukwa is an open source data collection system for monitoring large distributed systems. Chukwa is built on top of the Hadoop Distributed File System (HDFS), HBase and Map/Reduce framework and inherits Hadoop's scalability and robustness. Chukwa also includes a flexible and powerful toolkit for displaying, monitoring and analyzing results to make the best use of the collected data.

- Chukwa 0.5.0 has been released
- New committer Ahmed Fathalla has been voted in and granted proper karma
- Mentor William A. Rowe Jr resigned
- Alan D. Cabrera volunteer to become new mentor

- For Chukwa 0.5.0 release, the majority of code has been contributed from single individual.
- There are questions raised by mentors that there should be more diversity of the contribution from the community.
- There are increased activities on user mailing list in the past month, but momentum and traction is still a concern.
- More to the point, we had a discussion whether or not to continue Incubation for Chukwa. The consensus was that it's worthwhile to wait for the project to get traction a little longer.

Signed off by mentor: berndf, cdouglas

-----

#### DeviceMap

Apache DeviceMap is a data repository containing device information, images and other relevant information for all sorts of mobile devices, e.g. smartphones and tablets.

Entered incubation on January 3rd, 2012.

There are no issues that require the Incubator PMC's or the board's attention.

The project is still slowly ramping up. We have started to flesh out our website, created some initial prototypes for data collection of mobile device data and setup a VM at <http://devicemap-vm.apache.org/> to experiment with those prototypes.

Signed off by mentor: bdelacretaz, kevan

-----

#### DirectMemory

(incubating since October 2011)

Apache DirectMemory is a multi layered cache implementation featuring off-heap memory management (a-la BigMemory) to enable efficient handling of a large number of java objects without affecting JVM garbage collection performance.

There is only one important issue to address in the move towards graduation

- understanding process/decision making guidelines (new committer process is undergoing testing, release process still yet to be worked out)

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

- None

How has the community developed since the last report

- Tasks and proposals contributed by non-committers
- The team has up-voted a new committer

How has the project developed since the last report.

- An important contribution (the integration with EHCache, the most used caching system) could spread adoption and community attention
- A talk of DirectMemory should be presented at DevoXX France (thanks to Olivier Lamy)

Signed off by mentor: twilliams

-----

#### EasyAnt

EasyAnt is a build system based on Apache Ant and Apache Ivy. Incubating since 2011-01-31.

Towards graduation, we need to:

- Create a release
- Build a community

Since the last report:

In the last report we described that the activity of the podling was quite low. Some discussion with developers and mentors raised about this concern. Developers acknowledged that the current code refactoring make it's build break too much to keep people interested in contributing. Developers are still motivated even if most lack of time. The conclusion is that the focus should be on doing a release ASAP to keep things going on and bring back stability.

Some discussion started about the last points to tackle before preparing a release.

Signed off by mentor: bodewig

-----  
Flex

Apache Flex is an application framework for easily building Flash-based applications for mobile devices, the browser and desktop.

Summary: A highly energetic community waits for Adobe to complete the source code transition and infrastructure.

Date of entry to the Incubator: December 31, 2011

Top three items to resolve before graduation:

- Resolve trademark donation or licensing
- Complete code and bug database donation
- Make at least one release

Is there anything that the Incubator PMC or ASF Board specifically needs to address?

The ASF legal team may have to agree to a custom Trademark licensing. Adobe is currently evaluating the parameters of a custom Trademark license proposed by an ASF Board member. If approved, a custom agreement will be drafted and signed by both parties. Apache is waiting on Adobe to approve the latest Apache proposal.

Are there any legal, infrastructure, cross-project or personal issues that need to be addressed? (Are there any stumbling blocks that impede the podling?)

Besides the trademark issue mentioned previously, the import of existing JIRA bugs is blocked by an issue with Apache JIRA's import utility. It has failed to import the data file supplied by Adobe. Infra has opened a support ticket with Atlassian, but the problem remains unresolved one month later.

Check that the project's Incubation Status file up to date.

<http://incubator.apache.org/projects/flex>

What has been done (releases, milestones, etc.) since the last report?

- The Flex compiler source code, including history, was donated to Apache. It is pending import into the SVN repository
- Other initial committers have donated code to the whiteboard.
- Martin Heidegger was accepted as committer to the podling.

What are the plans and expectations for the next period?

Adobe expects to complete the transfer the mustella test suite. We hope to make progress on an initial release.

Are there any recommendations for how incubation could run more smoothly for you?

Our mentors are great. We are wondering if Apache has considered giving large projects like Flex their own SVN and JIRA servers. Recently, we pushed

really hard to get a software grant signed (An Adobe VP even came in to sign it while he was taking time off to move to a new house) in order to try to make the weekend deadline for getting code imported into SVN. We are in a hurry because we have seen signs we are losing momentum (mailing list traffic is much lower) and a key Flex conference is coming up on April 15. Being able to announce that new code was in SVN would be a good boost and good timing.

We submitted a software grant and code on Friday afternoon, but were too late to have the secretary record the software grant and that meant that Infra couldn't import the code since they only do that on weekends, and so we are now waiting for the next weekend. Having our own separate SVN repository would make weeknight imports possible and maybe not tax a single database as much. Large projects like Flex have tons of files.

The same holds true for the bug database as well. We submitted the JIRA import file on Feb 1. Infra has been unable to get the import to succeed (it worked on a test instance on my computer) and has had to file a bug with Atlassian (JIRA's maker). We asked about using the SOAP interface to import bugs but Infra is concerned that some problem in that process would screw up JIRA for everyone and/or cause performance issues. Having our own JIRA instance would allow us to try alternate means of getting around the JIRA issue. A SOAP import utility exists and was used to successfully import all 30,000 issues into a test instance but there is risk of having some other issue show up when we add that many issues to an existing database.

Signed off by mentor: bdelacretaz, greddin, wavw

-----  
Giraph

Giraph is a large-scale, fault-tolerant, Bulk Synchronous Parallel (BSP)-based graph processing framework that runs on Hadoop. Giraph entered the incubator in August 2011.

Project developments:

- First incubator release (version 0.1)
- Increased traffic on user list from new users
- Committer/PPMC member Claudio gave talk on Giraph at FOSDEM (<http://bit.ly/ClaudioGiraphTalk>)
- "Processing over a billion edges on Apache Giraph", a talk by Committer/PPMC member Avery accepted for Hadoop Summit in June.
- Community working with students at UC Irvine who are exploring porting various algorithms to Giraph.
- 53 new JIRA issues opened since last report
- 36 JIRA issues closed since last report

Next steps:

- Another incubator release
- Adding new committers

Signed off by mentor: omalley, mattmann

-----  
JSPWiki

JSPWiki has been incubating since September 2007.

JSPWiki is a JSP-based wiki program.

As a result of discussing in the dev-list, the objectives for graduation have changed: the first Apache release will be based on 2.8 (package-renamed), which is battle-tested, being the plan to evolve from there.

4 JIRA issues were fixed since last period, and tests are now part of the main build.

The main issues blocking graduation now are resolution acceptance, and



community and IPMC graduation vote.

Signed off by mentor:

-----

Kafka

(introduced to Apache incubator on Jul 4, 2011)

Kafka provides an extremely high throughput distributed publish/subscribe messaging system. Additionally, it supports relatively long term persistence of messages to support a wide variety of consumers, partitioning of the message stream across servers and consumers, and functionality for loading data into Apache Hadoop for offline, batch processing.

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

- Invite diverse new active committers

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

- None

How has the community developed since the last report?

- Added a new committer. Continued contribution to source code in fixes and features and continued activity on the mailing lists.

How has the project developed since the last report?

- Work was begun on 0.8 release which includes our replication feature and also requires a breaking change for existing consumers/producer to the wiring protocol.

Signed off by mentor: cdouglas, adc

-----

Kitty

The Kitty Project is a lightweight, production focused, Java-based application server performance diagnostic and management utility. Kitty is written in Groovy.

The project entered incubation on 2010-10-03.

In order to get graduated, Kitty needs the following features:

- The ability to save profiles of commonly connected to jmx servers including groups of tomcat servers
- The ability to collect metrics on an aggregate group of JMX servers

There are no issues that require the Incubator PMC's or the board's attention.

The community hasn't grown since the last report.

The project had some contribution as far as code is concerned, but the mailing list had no new threads started nor old threads had new contributions. Changes were related to the management of the console output, and some re-engineering in order to begin working on the missing features.

Signed off by mentor: kevan

-----

Mesos

Mesos is a scalable cluster manager that can dynamically share resources between multiple computing frameworks, including Hadoop, Spark and MPI.

Mesos entered incubation on December 23, 2010.

Progress since the last report:

- Begun the process for our \*first release!\* (vote on mesos-dev list has passed, now moving to vote on general@incubator)!
- Discussion/vote thread started about adding two new committers
- Hadoop patch working and new tutorial written
- API improvements declineOffer
- Updates to cluster Deployment scripts
- Committed major build system update (migrated to Autotools, build sped up by >2x)

Top priorities prior to graduation:

- Finalizing addition of new committers to the project!

Issues for Incubator PMC or ASF Board:

- None at this time.

Signed off by mentor: tomwhite

-----  
ODF Toolkit

The ODF Toolkit is a set of Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents. Unlike other approaches which rely on runtime manipulation of heavy-weight editors via an automation interface, the ODF Toolkit is lightweight and ideal for server use.

ODF Toolkit entered incubation on Aug 1st, 2011.

Most important issues to address.

- Growing the community, increasing diversity of committers

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

- None at this time.

How has the community developed since the last report

- We have two GSoC mentors and two proposals and are discussing these with students.
- Most recent new committer was voted in November 16th, 2011

How has the project developed since the last report.

- We did our first, successful podling release on January 12th
- Researched new crypto reporting requirements in preparation for checking in encryption/digital signature support.
- Working on next release, to feature the new crypto support, for later this Spring.

Signed off by mentor: nick

-----  
Oozie

Oozie is a workflow management and scheduler primarily for Hadoop based jobs.

Oozie entered the incubation on July 11, 2011.

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

- Add new committers.
- Further improve the documentations: user, development for quicker adoption.
- Automate the formal release process.

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

- No issues.

How has the community developed since the last report:

- Oozie dev/users are active in the email lists (around 2000+ Emails in oozie-dev and 300+ in oozie-user).
- In last one month, 67 JIRAs were created and 49 were resolved.
- 2 new developers from 2 new companies are contributing to coding.

How has the project developed since the last report:

- The first Oozie release (3.1.3) from Apache incubator has been successfully released.
- A new release (3.2) process is started. The new branch is created.
- Cookbook like documentation was added to Oozie product web page.

Signed off by mentor: Devaraj Das

-----  
OpenOffice

Out-of-Band Report

On March 22nd we issued a security bulletin and patch for a vulnerability reported in OpenOffice.org 3.3, the last released pre-Apache version of OpenOffice.

Technical details of the issue and patch can be found in this announcement:

[http://mail-archives.apache.org/mod\\_mbox/incubator-ooo-announce/201203.mbox/%3CCAP-ksoj7o5%2B2YH-E4XzR044V0e3YZfZvuef7eJuNGhdy%2Bk9kyA%40mail.gmail.com%3E](http://mail-archives.apache.org/mod_mbox/incubator-ooo-announce/201203.mbox/%3CCAP-ksoj7o5%2B2YH-E4XzR044V0e3YZfZvuef7eJuNGhdy%2Bk9kyA%40mail.gmail.com%3E)

The IPMC and Board should note that this extraordinary patch was made available, as a courtesy to the ecosystem, based on the severity of the reported vulnerability and the ease of exploiting it. The patch was made available under ALv2, and distribution was done via the Apache mirrors, although this did not constitute an official release.

Because of the required secrecy around the preparation of such security patches, a minimum number of Apache members were involved in vetting this release, though we did try to touch all bases by involving mentors, Infra and Legal Affairs.

Signed off by mentor: joes, rgardler

-----  
Syncope

Apache Syncope is an Open Source system for managing digital identities in enterprise environments, implemented in JEE technology. Syncope joined the incubator on February 10th 2012.

Community Development:

Older Syncope users-base is (not really quickly) migrating to the new ASF infrastructure, user@ ML activity is low.

Project Development:

Commits activity is healthy, commits@ ML activity is high.

Committers are working towards the first release.

Web Site/Communication Development:

The website is kept up-to-date on <http://incubator.apache.org/syncope/>  
Status page <http://incubator.apache.org/projects/syncope.html> is updated

It introduces Syncope and provide first content.

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

None so far.

Signed off by mentor: simonetripodi

-----  
Tashi

Tashi has been incubating since September 2008.

The Tashi project aims to build a software infrastructure for cloud computing on massive internet-scale data sets (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 originally encompassed just the tools to manage virtual machines using Xen and QEMU, but has been merged with Zoni, which manages the physical aspects of a cluster like power control, network settings and handing out physical machines.

In the period from January to April, the project had received permission to publish a release. Shortly thereafter, a further release was approved, incorporating several bug fixes detected during deployment of the first release and from Jira reports.

In the process of making our first release, contact to two of our three mentors was re-established. The third mentor has not been heard from (also in other parts of Apache) for quite a while. We obtained sufficient administrative access to Jira to manage our problem reports.

Development efforts this period have mainly been in adding resilience to Tashi components, as well as returning more helpful messages in case of errors. Some parts of the code base that were relevant only to Thrift have been removed.

Upcoming software goals are to investigate what is needed to support IPv6, considering replacement for RPyC and providing the ability to hand out server slices (operating system level virtualization).

The project has a user community, but it is small. Growth mostly has happened by word of mouth. To show potential users at large the utility of this project, the author of this report will apply some of the advice posted to [general@apache.incubator.org](mailto:general@apache.incubator.org), as well as create web pages demonstrating the project's utility, provide sample VM images and disseminate other information relating to the deployments close to him. He will also urge others to make similar information publicly available.

One of our users has requested the creation of a private branch for his team to work on. Perhaps this will result in a valuable feature addition to the project.

Items to be resolved before graduation:

- Generate more publicity for the project.
- Develop members of the user community to submit feature extensions.

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:

- change of the focus of the project (the previous focus was to be flexible wrt the representation of the data model, the new focus is parallel processing of large amounts of XML data)
- renewed development activity

- 2 proposals for GSoC (1 person expressed interest in working on VXQuery for GSoC)

Top issues before graduation:

- Build community

VXQuery's community is still the initial (very small) community  
The goal of the change of focus to attract more interest in the  
project and to increase the size and diversity of the community

- Create a release

Signed off by mentor:

-----