

October2012

Incubator board report, October 2012

The biggest news this month that of the OpenOffice project, along with Cordova and Isis, becoming ready to graduate from the Incubator. As noted below, we feel that these projects are capable of governing themselves as standalone TLPs according to Apache policies and the Apache Way, and recommend the board to accept the respective resolutions.

In other news, Jukka Zitting announced his intention to resign as the IPMC chair due to changes in personal life. Discussion about and nominations for the next IPMC chair are open, and we expect to have a related resolution ready in time for the November board meeting.

o Community

There were no changes to the Incubator PMC since our last report. The IPMC currently has 168 members.

The following podlings are requesting graduation to Apache TLPs:

- Apache Cordova
- Apache Isis
- Apache OpenOffice

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

The following proposals for new incubating projects are being voted on:

- Apache Helix
- Apache Ripple

In addition there seems to be some interest in reviving the earlier idea of bringing the BeanShell project (<http://beanshell.org/>) to the ASF through the Incubator.

Meanwhile the AWF podling was retired due to inactivity, and the Kitty podling is voting to retire itself for the same reason.

JSPWiki is our oldest podling with over five years in the Incubator. Activity in JSPWiki was very low for a few years and they've yet to create their first Apache release. Earlier this year they discussed leaving the Incubator and the ASF since they clearly weren't making much progress. That discussion led to some revival of activity and the decision to continue in the Incubator. Unfortunately the podling no longer has enough active mentors, which has led to some trouble with premature attempts at cutting releases or graduating. Despite these troubles the podling is making progress, and with sufficient help from the IPMC they might well become ready to graduate within a few quarters.

o Releases

The following incubating releases were made since our last report:

- September 20th, 2012: Apache Crunch 0.3.0-incubating
- September 24th, 2012: Apache Ambari 0.9-incubating
- October 1st, 2012: Apache Cordova 2.1.0-incubating
- October 8th, 2012: Apache Syncope 1.0.2-incubating
- October 12th, 2012: Apache Kafka 0.7.2-incubating

The Wookie podling has an open crypto export issue that's blocking their release. Since it looks like the issue can be resolved in time and the podling can demonstrate ability to cut releases, they can proceed towards graduation already before the resolution of that issue.

We discussed key signing and various ways of getting more podling release managers into the Apache web of trust. No concrete actions have yet come out of that discussion, but there's been plenty of good ideas that may end up being implemented after a while.

o Legal / Trademarks

The DeviceMap podling encountered a tricky issue with unclear ownership of a data file they planned to use. The resolution for now is to steer clear of trouble by not including that file. See the DeviceMap report and the mentioned issue for more details.

A somewhat related question came up from CloudStack about how to handle files that we can't redistribute due to upstream licensing or Apache policies. They had a few such files in imported version history, and the consensus was that it's fine as long as such files are removed from the heads of all active branches (so they won't be included in releases) and that there's no need to modify older revisions as long as something like that is not explicitly requested.

o Infrastructure

The infra team started looking at ways to streamline the various tasks related to podling setup and graduation. In practice this would mean making podling infra resemble more that of a standalone TLP (for example with separate podling.incubator.apache.org or even podling.apache.org domains), which seems like a reasonable thing to do as long as the incubation disclaimers remain prominent.

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

Still getting started at the Incubator (2 podlings)

Blur, Drill

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

Not yet ready to graduate (11 podlings)

No release: Celix, Cloudstack, JSPWiki, VXQuery
Low activity: Chukwa, DeviceMap, Kitty, Tashi
Low diversity: EasyAnt, Mesos, ODF Toolkit

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 podling)

Kafka

We expect this project to graduate within the next quarter.

Blur

(incubating since August 2012)

Apache Blur is a search platform capable of searching massive amounts of data in a cloud computing environment.

Issues needing Board/Incubator PMC Attention:

- None.

Key Activities:

- Git repo is just now in place - code cloned from Github and updated to latest revision so the code cleanup (package naming, licensing, etc.) are in the works.

- We decided to add a wiki and thanks to infra@ that's now up.
- Website is created and CMS-ified.
- Public Blur hack sessions/meetups are being held (Mondays) at Near Infinity with summaries posted on dev list.

Community:

- Issues are being created and worked with discussions of new features and modifications on the mail-list.
- Subscriptions: users@ - 19[+2]; dev@ - 21[+3]

Signed-off-by: twilliams

Celix

Celix is an implementation of the OSGi Specification in C.

Celix entered incubation on November 2, 2010.

Over the last months we have focused on creating a first release for Celix. As is mentioned in the graduation plan [1] we hope that a release will create more community. We have prepared the code base and website for release and we are very close to making a release.

The previous report [2] mentioned that we are working on a "Native-OSGi", which could attract more community. This discussion continued in July on the mailing list, but got quiet after that. There will be a talk about Native-OSGi for ApacheCon EU and EclipseCon EU and we expect that this will create attention for Celix.

Most important issues to address before we can graduate:

Make a first release, growing the community and attracting more committers.

Any issues the Incubator PMC or ASF board need to be aware of:

None at this time

How has the community developed since the last report:

There has been a small decrease on activity on the mailing list, hopefully this is a temporary decrease due to the holidays. We did receive patches from a contributor.

How has the project developed since the last report:

We are still working towards a first release. The code base has been prepared for release and the website got some attention based on the shepherd's view from last report.

[1]: <http://incubator.apache.org/celix/community/boardreports/boardreports.html#2012-04>

[2]: <http://incubator.apache.org/celix/community/boardreports/boardreports.html#2012-07>

Signed-off-by: marrs, jukka, wave

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.

Incubating since July, 2010.

- Mailing lists have some activities.
- 2 new committers since July.
- 3 new contributors submitted patches since last report in July.

Most important issues to address:

Growing the community, especially a new release manager

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:

Plan to release version 0.6 to fix some of the major bugs in version 0.5.

How has the project developed since the last report:

More patches are received from new committer and contributors. The new patches are better quality to polish Chukwa code base for 0.6 release. Chukwa community need to train new release manager to release Chukwa 0.6 to ensure the community can continue to flourish. The previous concern of inactive development is less concerning than during January to July time frame.

Signed-off-by: cdouglas, berndf, jukka

CloudStack

CloudStack is an IaaS ("Infrastructure as a Service") cloud orchestration platform. CloudStack has been in incubation since 2012-04-16

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

1. Continuing to build the community and increase diversity
2. Shipping a release
3. Transfer trademark, etc. to Apache

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

No issues at this time

How has the community developed since the last report?

The project has continued to improve its diversity, adding Chip Childers, John Kinsella, and Wido den Hollander to the PPMC and 3 new committers since the month of August. (Mice Xia, Jason Bausewein, and Joe Brockmeier.) We have also taken on a new mentor, Noah Slater.

Contributors and committers have been active in promoting CloudStack at events, such as LinuxCon, Ohio LinuxFest, PuppetConf, a get-together at Schuberg Philis, and a number of other events.

The community is participation at LinuxCon EU and ApacheCon EU and a number of other events through the remainder of 2012.

Citrix is acting as primary sponsor of a CloudStack Collaboration Conference for November 30 through December 2 in Las Vegas. The conference is open to the community, and programming for the event will be chosen by a committee that includes members of the CloudStack community outside Citrix.

How has the project developed since the last report?

The pending 4.0 release has been branched. The project has made a great deal of progress towards the 4.0 release, handling a number of technical issues (such as a move to Maven) and resolving almost all known legal issues that would pose an obstacle to a release. The sole remaining

blocker is under discussion and should be resolved shortly.

The Jira instance for CloudStack was stood up in early September. License checks have been automated to ensure that we remain compliant with Apache guidelines going forward.

Signed-off-by: nslater, mfranklin

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.

No new committers or PPMC members, and very low activity so far, but the OpenDDR code that was mentioned in the incubation proposal has just been donated, which should help motivate people to become more active.

The donation included a data file with had been subject to a DMCA takedown request while hosted at Github. After some discussion we decided to exclude that file from the donation, as the ownership of part of that data is hard to establish. Details at <https://issues.apache.org/jira/browse/DMAP-11> which also links to threads on general@incubator.a.o.

The next steps are preparing that code for a first release, and discussing how to best integrate the data that OpenDDR has been collecting themselves and as such can be contributed without problems.

Signed-off-by: bdelacretaz (mentor), jukka

Drill

Drill's goal is to build an open source clone of Dremel with appropriate extensions to foster greater flexibility.

Drill has been incubating since September of 2012.

Since last month, we have been working on infrastructure. A prototype web-site is ready and several code contributions are nearly ready to commit.

Graduation is still very far away, but the community activity has been high and the mailing list has been active. Numerous public presentations have been made and several Drill Users' Groups have been formed and meetings held.

Most important issues to address before we can graduate:

- Get the basics in place, build up a working code base, make releases (that is, everything)

Any issues the Incubator PMC or ASF board need to be aware of:

- None at this time

How has the community developed since the last report:

- Several active contributors outside the current committer group have emerged. We are working to bind these new contributors into the community and several appear likely to become committers over time.

How has the project developed since the last report:

- The community has begun to gel nicely and significant code contributions have moved forward.

Signed-off-by: Ted Dunning acting for Grant Ingersoll

EasyAnt

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

Since the last report, there were discussion about the graduation process of EasyAnt. There was confusion about what would be its final place - either a TLP or a subproject of the Ant TLP - and the implication about the 'graduation' process.

Things have been cleared with a vote: the EasyAnt project members are willing to join the Ant TLP. So there is no actual need for a full classical graduation process.

Then a vote happend in the Ant project about accepting EasyAnt as a subproject. It failed, there were some concerns about the activity of the project. There were actually no veto, but an unsufficent number of binding +1.

So it has been decided to try again in a few month after showing some nice activity on the EasyAnt project.

A little bit later, some Ant PMC members wrote that they were 'unplugged' when the vote happend and were sorry to have missed it. A good sign for a retry of the vote.

About the activity of the project since the last report: one committer did some great job making the code base ready to be released. One other committer was active on the dev mailing list. We have also seen some non committers asking questions about the use of the project.

The next step for the project:

- doing a release: it will show that the EasyAnt committers are dedicated to maintain it and keep it going
- shortly after that, ask the Ant PMC again to accept EasyAnt as a subproject
- if accepted, EasyAnt could then leave the Incubator.

Signed-off-by: ant, bodewig, jukka

JSPWiki

JSPWiki has been incubating since September 2007.

JSPWiki is a JSP-based wiki program.

Still following the objective of making the first Apache release:

A few JIRA issues were raised in this period due to incorrect handling of external dependencies. They have been fixed, so there aren't any other technical issues blocking a release. A new RC and a release vote has been cast in jspwiki-dev, with 10 +1 (6 from PPMC members). However, we don't have any IPMC vote yet, so the release thread was forwarded to general on Friday 26th.

It's likely that we will have our first Apache release during the first weeks of October.

The main issues blocking graduation keep being the same as in the last report: resolution acceptance, community and IPMC graduation vote, the latter requiring at least one ASF release.

The developer list currently has 91 subscribers; and the user list has 186 subscribers.

Signed-off-by: jukka

IPMC comments: See paragraph on JSPWiki in the report summary.

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

None. We started a graduation discussion on our mailing list. This was slightly delayed due to focus on the development of the intra-cluster replication feature.

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

The mailing list continues been very active with bug reports, patch submissions, feature requests and use case discussions. [(Jul, Aug, Sep): kafka-users (155, 164, 126); kafka-dev (445, 532, 589)].

Received and reviewed several major patches especially in the 0.8 branch.

How has the project developed since the last report.

Intra-cluster replication (KAFKA-50): Development has been very steady on this feature. (Hence the double volume of emails on kafka-dev since the last report.) We set up a development dashboard at <https://cwiki.apache.org/confluence/display/KAFKA/Kafka+replication+development> that gives an overview of the remaining work for that release.

The vote for Kafka 0.7.2 is currently underway. This will be the third Kafka release after its introduction to Apache incubator.

Signed-off-by: cdouglas, bmargulies

Kitty

(No report. Kitty is voting to retire itself due to inactivity.)

By Shepherd: Retirement VOTE thread is +3 to retire with no other votes. It has not been closed, but was started on October 1.

Signed-off-by: Wave (shepherd)

Mesos

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

Mesos entered incubation on December 23, 2010.

Progress since the last report:

- Added mesos-0.9.0-incubating.jar to Maven central.
- Improved build to include warnings and errors.
- Fixes related to ZooKeeper usage to deal with network partitions.
- Committed cgroups isolation module and additional updates and fixes.
- Numerous updates to webui, including preliminary access to files of running frameworks.
- Lots of reviews (<https://reviews.apache.org/dashboard/?view=to-group&group=mesos>).

Top priorities prior to graduation:

- Continue to grow community (lots of questions on the mailing list).
- Add more committers
- Establish a 6-week release cadence.

Issues for Incubator PMC or ASF Board:

Not all committers have access to Jenkins, which makes updating tests a slow bottlenecked process.

Signed-off-by: tomwhite, jukka
IPMC comments: Jenkins access granted.

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.
Our last release was January 14th, 2012.

We voted in our most recent committer on September 27th, 2012.

Most important issues to address:

Growing the community, especially attracting new developers. Generally we'd be pleased to attract 2 or 3 additional active committers, have another release, and then we think we'd be well-positioned for graduation.

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 successfully completed two GSoC projects. However neither led to longer term engagement with the project.

We've seen an increase in incoming patches from users, and have a new active contributors. One was just voted in as a Committer/PPMC member.

How has the project developed since the last report:

Working on next release.

Signed-off-by: bmargulies

Tashi

2012-October Tashi Incubator Status Report

Tashi has been incubating since September 2008.

The Tashi project aims to build a software infrastructure for cloud computing on massive internet-scale datasets (what we call Big Data). The idea is to build a cluster management system that enables the Big Data that are stored in a cluster/data center to be accessed, shared, manipulated, and computed on by remote users in a convenient, efficient, and safe manner.

Tashi 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.

Activities July-October:

In the period from July to October, the project did not ask to make another incubating release, but is ready to start the process for a new release incorporating the development efforts of the last 6 months.

Development efforts this period have included providing a separate administration client, allowing addition of users and networks, and host reservations and availability for scheduling.

The project has received code contributions from one non-committer in this period. Diogo Gomes provided support for deriving the IP addresses of guests automatically, without having to scan the subnet. Thanks Diogo!

Additional stability and user experience improvements were also committed.

Upcoming software goals are to investigate what is needed to support IPv6, replace RPyC, and to provide the ability to hand out server slices (operating system level virtualization). Besides CPU and memory, disk storage should also be a schedulable resource.

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 is creating web pages to demonstrate how to accomplish distributed computing tasks. Base images of (free) OS installs will be provided to allow new users to get started quickly. Hopefully this will increase visibility of 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: mfranklin

VXQuery

The VXQuery Project implements a standard compliant XML Query processor for parallel evaluation.
It has been in incubation since 2009-07-06.

VXQuery was able to maintain the increased level of activity that was reported in the July report and made significant progress towards completing the XQuery functionality.

The remaining top issue is a release. Recently the preparations for the first release have started and they are still ongoing. Release artifacts should be available and voted on in the next few weeks.

Signed-off-by: jochen, jukka
