

September2014

/\ *FINAL* (!)

This report has been submitted to the Board. Please do not edit further.

Incubator PMC report for September 2014

Timeline

Wed September 03	Podling reports due by end of day
Sun September 07	Shepherd reviews due by end of day
Sun September 07	Summary due by end of day
Tue September 09	Mentor signoff due by end of day
Wed September 10	Report submitted to Board
Wed September 17	Board meeting

Shepherd Assignments

Alan D. Cabrera	Fleece
Andrei Savu	Drill
Andrei Savu	Streams
Dave Fisher	Optiq
John Ament	Argus
John Ament	log4cxx 2
Justin Mclean	MRQL
Konstantin Boudnik	Sentry
Matthew Franklin	Kalumet
Raphael Bircher	Storm
Roman Shaposhnik	Brooklyn
Ross Gardler	Wave
Suresh Marru	Falcon
Suresh Marru	REEF

Report content

Incubator PMC report for September 2014

The Apache Incubator is the entry path into the ASF for projects and codebases wishing to become part of the Foundation's efforts.

There are 32 podlings under incubation this month. No projects are currently voting on graduation. The incubator added REEF as a new incubating project in August. Aurora, Drill, Sentry, Twill, Flink, and Optiq all had releases since the last report.

* Community

New IPMC members:

No additions

People who left the IPMC:

(none)

* New Podlings

REEF

* Graduations

The board has motions for the following:

Apache Storm

* Releases

The following releases were made since the last Incubator report:

Apache Aurora 0.5.0-incubating
Apache Drill 0.4.0-incubating, 1.0.0-m1-incubating
Apache Sentry 1.4.0-incubating
Apache Twill 0.3.0-incubating
Apache Flink 0.6-incubating
Apache Optiq 0.9.0-incubating

* IP Clearance

Radii8 for Apache Flex
Apache Sqoop Contribution

* Legal / Trademarks

Apache OPTIQ wishes to trademark it's name

* Infrastructure

Nothing to report

* Podlings without mentor sign offs

None

* Miscellaneous

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

* Still getting started at the Incubator

REEF

* Not yet ready to graduate

No release:

Argus
Brooklyn
Fleece
log4cxx2
Wave

Community growth:

Brooklyn
Drill
Falcon
Fleece
MRQL
Optiq

Sentry
Wave

* Ready to graduate

The Board has motions for the following:
Storm

* Did not report, expected next month

Streams

Table of Contents

Argus
Brooklyn
Drill
Falcon
Fleece
Kalumet
log4cxx2
MRQL
Optiq
REEF
Sentry
Storm
Streams
Wave

Argus

The Argus project is a framework to enable, monitor and manage comprehensive data security across the Hadoop platform.

Argus has been incubating since 2014-07-24.

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

1. Get all podling committers and mentors on mailing lists - COMPLETED
2. Publish initial code drop along with a code grant - COMPLETED
3. Get approval for name of Argus (PODLINGNAMESEARCH-54) - IN PROGRESS
4. Create a project website - TO DO

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?

* Developers are showing interest in contribution to Apache Argus

How has the project developed since the last report?

- * Code Grant has been sent to Apache
- * Initial Codebase is uploaded to Apache Argus git repository
- * Review Board for Argus project has been setup for code review process
- * Total of 32 JIRA tickets have been created; out of which 21 of them are RESOLVED.
 - 9 - CLOSED
 - 12 - RESOLVED,
 - 3 - PATCH_AVAILABLE
 - 4 - IN PROGRESS
 - 4 - OPEN

Date of last release:

none yet

When were the last committers or PMC members elected?

none yet

Signed-off-by:

[](argus) Alan Gates
[](argus) Daniel Gruno
[](argus) Devaraj Das
[x](argus) Jakob Homan
[x](argus) Owen O'Malley

Shepherd/Mentor notes:

Argus appears to be brand new. They don't have a working website yet (was unclear if there was a pending ticket for it not), though mailing list is working and let me subscribe. They do have a draft board report floating around their mailing list, which is a great sign for a brand new podling.

Brooklyn

Brooklyn is a framework for modeling, monitoring, and managing applications through autonomic blueprints.

Brooklyn has been incubating since 2014-05-01.

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

1. Performing our first release under Apache
2. Grow the community
3. More diversity of the committers/PPMC (currently biased towards employees of a single organization)

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?

Our community continues to grow slowly but surely, and has received contributions from new community members that are significant and highly interesting to our project.

How has the project developed since the last report?

No major changes from the last report; we are well-settled into our new home and commits are being made at a regular rate. Preparation for our first release continues.

We neglected to mention in last month's report about the completion of a code grant for the "camp-server" component. This is a crucial component of Brooklyn, and deserves its place in Brooklyn, rather than an external dependency where we would have no control over release schedules and bug fixing. So we are grateful to the donor, Cloudsoft, for passing this component to Apache Brooklyn.

Date of last release:

No releases under Apache yet.

When were the last committers or PMC members elected?

No change since 2014-07-02.

Signed-off-by:

[](brooklyn) Matt Hogstrom
[](brooklyn) Alex Karasulu
[X](brooklyn) David Nalley
[](brooklyn) Marcel Offermans
[](brooklyn) Jean-Baptiste Onofré
[](brooklyn) Olivier Lamy
[](brooklyn) Chip Childers
[](brooklyn) Andrei Savu
[X](brooklyn) Joe Brockmeier
[X](brooklyn) Jim Jagielski

Shepherd/Mentor notes:

(rvs) a very strong community, the only thing missing to make it exemplary would be a release ;-)

Drill

Description:

Apache Drill is a distributed system for interactive analysis of large-scale datasets that is based on Google's Dremel. Its goal is to efficiently process nested data, scale to 10,000 servers or more and to be able to process petabytes of data and trillions of records in seconds.

Drill has been incubating since 2012-08-11.

In the previous reports, the following were listed as goals before graduation

1. Complete the feature set
2. Continue to attract new developers/contributors with a variety of skills and viewpoints
3. Continue the outreach activities to build the early user community for the technology

These have been achieved and the podling has made several releases with no more than minor issues that were related to changing requirements for notices in incubator projects. The next release (0.5) is currently being voted on. Subsequent to that, the podling is likely to vote to request the board to graduate Drill to TLP status.

Issues to Call to Attention of PMC or ASF Board:

None

How community has developed since last report:

Community awareness and outreach were strengthened in multiple forums as below

8/7/14	Big Data Analytics Melbourne	MC Srivas
8/13/14	Chicago HUG Chicago	Jim Scott
8/20/14	Pittsburgh HUG Pittsburgh	Andy Pernsteiner
8/21/14	Heartland Big Data Omaha, NE	Neeraja Rentachintala
8/26/14	Data Mining San Francisco, CA	Tomer Shiran

Mailing list discussions:

Activity summary for the user mailing list:

http://mail-archives.apache.org/mod_mbox/incubator-drill-user/

- * June 2014: 79
- * July 2014, 12
- * August 2014, 63

Activity summary for the dev mailing list:

http://mail-archives.apache.org/mod_mbox/incubator-drill-dev/

- * June 2014, 374
- * July 2014, 294
- * August 2014, 247

For details of code commits, see <https://github.com/apache/incubator-drill/graphs/commit-activity> (about 400 commits in the past 3 months)

31 contributors have participated in GitHub code activity; there have been 152 forks.

Community Interactions

Monthly Drill hangout continues, conducted remotely through Google hangouts Tuesday mornings 10am Pacific Time to keep core developers in contact in realtime despite geographical separation.

Community stays in touch through @ApacheDrill Twitter ID, and by postings on various blogs including Apache Drill User <http://drill-user.org/> which has had several updates and through international presentations at conferences.

Articles

Examples of articles or reports on Apache Drill since last report include:

- * Self Service Data Exploration is Here by Neeraja Rentachintala

Social Networking

@ApacheDrill Twitter entity is active and has grown substantially by 20%, to 1057 followers.

How project has developed since last report

Web-site clean slate revamp

Significant progress has been made in performance and stability

New functionality has been added to the product including reading and writing complex types in Parquet, as well as using hadoop 2 API for Parquet

Nearly ~450 bugs filed and ~550 bugs resolved

New docs have been published on Drill wiki (Develop Custom Functions, Querying HBase Tables, Querying Complex Data)

Started monthly releases. 0.4 release at end of July.
Announcement: <http://s.apache.org/t0a>

0.5 release currently up for vote.

Signed-off-by:

[x](drill) Ted Dunning
[x](drill) Grant Ingersoll
[](drill) Isabel Drost-Fromm
[X](drill) Sebastian Schelter

Shepherd/Mentor notes:

----- Falcon

Falcon is a data processing and management solution for Hadoop designed for data motion, coordination of data pipelines, lifecycle management, and data discovery. Falcon enables end consumers to quickly onboard their data and

its associated processing and management tasks on Hadoop clusters.

Falcon has been incubating since 2013-03-27.

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

1. Continue to build community

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

- No

How has the community developed since the last report?

* Three more committers were invited and they have accepted to join the project

* More users & contributors have joined the falcon project and the community continues to grow

How has the project developed since the last report?

* Development activity has been very hectic more than 200 JIRAs have been created and about 120 of them resolved since the last report

* There are more than 100 users subscribed on the dev mailing list

* We have formulated bi weekly sync up to coordinate with developers and contributors across the world

* 0.5-incubating release has been withdrawn due to LICENSE & NOTICE issues and same is intended to be prepared for vote shortly and will be released in Sep 2014.

Date of last release:

2014-02-03 (0.4-incubating)

When were the last committers or PMC members elected?

Aug 28, 2014

Signed-off-by:

[](falcon) Arun Murthy
[X](falcon) Chris Douglas
[](falcon) Owen O'Malley
[](falcon) Devaraj Das
[X](falcon) Alan Gates

Shepherd notes:

Fleece (currently renaming to Johnzon)

Implementation of JSR 353: JavaTM API for JSON Processing.

Fleece/Johnzon has been incubating since 2014-06-09.

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

1. Rename the project due to naming conflicts with "Fleece".
New approved name is "Johnzon". Waiting for INFRA-8269.
2. Expanding the community, increase dev list activity and adding new committers.
3. Publish first incubator release.

The Fleece project status page at <http://s.apache.org/0m> needs updating but that would be pending on #1 due to many items on the statuspage need to be re-visited once #1 is settled. Clutch needs also an update.

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

* N/A

How has the community developed since the last report?

- * Voted in one new committer.

How has the project developed since the last report?

- * Some minor fixes and improvements.

Date of last release:

- * No releases as of yet.

When were the last committers or PMC members elected?

- * New committer Hendrik Saly (salyh) on 2014-08-23

Signed-off-by:

- [X](fleece) Justin Mclean
- [](fleece) Christian Grobmeier
- [](fleece) Daniel Kulp

Shepherd/Mentor notes:

Kalumet

Kalumet a complete environment manager and deployer including J2EE environments (application servers, applications, etc), softwares, and resources.

Kalumet has been incubating since 2011-09-20.

Community Development:

Apache Kalumet 0.6-incubating version has been released. However, due to some misunderstanding, the release vote has not been completed by 3 IPMC. Especially, some legal files issues have been raised. We are preparing a 0.6.1-incubating release to fix the legal files and submit to IPMC vote. We are in the way of promoting the documentation on the website.

Project Development:

We are preparing the 0.6.1-incubating (plan for September, 20) to fix the legal files issues raised on 0.6-incubating release and have a complete IPMC vote. We are preparing the 0.7-incubating release with the development changes.

Local branches have been created containing:

- new model and REST API
- new webconsole (remove of Echo framework)

These local branches will be merged on the 0.7-incubating branch (master).

Before Graduation:

- The documentation has been updated and aligned with the 0.6-incubating release. The documentation will be promoted on the website and "linked" in announcement e-mails as soon as the 0.6.1-incubating release has been voted.
- The 0.7-incubating release will fix the graduation issue, especially in term of license (removal of Echo2 framework for the UI).

Post Graduation:

- Kalumet scripts. It's the extension of the "software" concept globally to all Kalumet resources. It will allow users to create custom deployment/update plan, with a set of fine-grained steps.

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

None so far.

Date of last release:

2013-11-22

Date of next release:

2014-08-24

When were the last committers or PMC members elected?

None

Signed-off-by:

[X](kalumet) Jim Jagielski
[](kalumet) Henri Gomez
[X](kalumet) Jean-Baptiste Onofre
[X](kalumet) Olivier Lamy

log4cxx2

Logging for C++

log4cxx2 has been incubating since 2013-12-09.

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

1. We need to release 0.11.0 as maintenance release.
2. We probably need to decide what's an acceptable platform/compiler list, sort out any remaining hiccups, and then move towards a release.
3. We need to make a broader use of the ASF infrastructure (notably the CI, as well as improved management of the Jira site - notably with triaging), and establish a roadmap for the next releases.

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

N/A

How has the community developed since the last report?

No changes in team.

How has the project developed since the last report?

Mails suggest that there's work going on in the background by some committers especially regarding the build process, but there hasn't been any notable commit in the last weeks. The current trunk now fixes most of the important outstanding bugs and we are actively supporting the project. Some minor changes have been applied for the website.

Date of last release:

2008-04-03 was the official, pre-incubation 0.10.0

Many post-0.10.0 commits exist on trunk which we intend to get out as 0.11.x.

When were the last committers or PMC members elected?

N/A

Signed-off-by:

[X](log4cxx2) Christian Grobmeier
[](log4cxx2) Scott Deboy

Shepherd/Mentor notes:

log4cxx is still getting started. While activity is happening, they need to get their first release out the door to be successful and start to get interest.

MRQL

MRQL is a query processing and optimization system for large-scale, distributed data analysis, built on top of Apache Hadoop, Hama, Spark, and Flink.

MRQL has been incubating since 2013-03-13.

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

1. increase the number of active committers
2. increase adoption, expand user community, and increase user list activity

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?

We have released our second release under Apache incubation. After we added Apache Flink as an evaluation backend for MRQL, there was an interest by the Flink community in our project, which may result in collaboration between the projects and may expand our user community.

How has the project developed since the last report?

We have added support for Apache Flink. Now users can run MRQL queries on a Yarn cluster using 4 different backends (Hadoop map-reduce, Hama, Spark, and Flink), without having to change their queries.

Date of last release:

2014-06-26

When were the last committers or PMC members elected?

2014-04-17

Signed-off-by:

[X](mrql) Alan Cabrera
[](mrql) Anthony Elder
[](mrql) Alex Karasulu
[](mrql) Mohammad Nour El-Din

Shepherd/Mentor notes:

jmclean: Mentors active and project healthy. Project has discussed graduation but feels it needs more active committers.

Optiq

Optiq is a highly customizable engine for parsing and planning queries on data in a wide variety of formats. It allows database-like access, and in particular a SQL interface and advanced query optimization, for data not residing in a traditional database.

Optiq has been incubating since 2014-05-19.

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

1. Migrate fully to Apache infrastructure (next: nightly builds).
2. Re-organize code into org.apache.optiq namespace.
3. Improve web site.
4. Build an ASF community.

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?

We are planning our first hackathon.

There have been community discussions about the goals for a future 1.0 release, in particular whether the project should be reorganized to allow projects to use the optimizer core without using SQL parser and code-generation facilities.

How has the project developed since the last report?

Adopting Apache infrastructure:

- * First release approved by PPMC and Incubator PMC
- * Web site

We will push our and announce the release shortly.

Apache infrastructure is extremely frustrating. The release has taken a lot longer than it should have because it frequently makes this newbie feel like an idiot. Questions are answered promptly but it is difficult to know what questions to ask.

Date of last release:

2014-08-25

When were the last committers or PMC members elected?

2014-07-15

Signed-off-by:

[X](optiq) Ted Dunning
[X](optiq) Alan Gates
[X](optiq) Steven Noels

Shepherd/Mentor notes:

REEF

REEF (Retainable Evaluator Execution Framework) is a scale-out computing fabric that eases the development of Big Data applications on top of resource managers such as Apache YARN and Mesos.

REEF has been incubating since 2014-08-12.

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

1. Set up the basic infrastructure: git repository, mailing lists (INFRA-8190)
2. Publish initial code drop along with a code grant
3. Get approval for name of REEF (PODLINGNAMESEARCH-55)
4. Update the existing project web site

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?

- * All committers have filed ICLAs.
- * JIRA created
- * Status page created

How has the project developed since the last report?

This is the first report.

Date of last release:

None yet.

When were the last committers or PMC members elected?

None yet.

Signed-off-by:

[X](reef) Chris Douglas
[X](reef) Chris Mattmann
[](reef) Ross Gardler
[X](reef) Owen O'Malley

Shepherd/Mentor notes:

Sentry

Sentry is a highly modular system for providing fine grained role based authorization to both data and metadata stored on an Apache Hadoop cluster. Sentry has been incubating since 2013-08-08.

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

1. Continue to grow the Sentry community
2. Initiate graduation discussion with community and address any potential concerns raised
3. Update project status page, website etc.

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?

The community is growing. A number of new members have started actively contributing to the project. New contributors made a significant contribution in patches as well as release activity. Two of contributors are elected to become new committers.

How has the project developed since the last report?

The project did its 3rd release from incubation. Thanks to lot of hard work and initiative from the release manager Tuong Truong, Sravya Tirukkovalur as his mentor and contributions from other community members, version 1.4 was released on 18th August. Several new features and patches are being actively worked on the master branch.

Date of last release:

2014-08-19

When were the last committers or PMC members elected?

Arun Suresh and Tuong Truong have been added as a committer as announced on 9/1/2014. No new PPMC members have been added since the project has entered the incubator.

Signed-off-by:

- (sentry) Arvind Prabhakar
- (sentry) Joe Brockmeier
- (sentry) David Nalley
- (sentry) Olivier Lamy
- (sentry) Patrick Hunt
- (sentry) Thomas White

Shepherd/Mentor notes:

dev@ list is active, although the amount of traffic was quite lower this month, compare to the past periods. Mentors are visible (IPMC report discussion)

Storm

Storm is a distributed, fault-tolerant, and high-performance realtime computation system that provides strong guarantees on the processing of data.

Storm has been incubating since 2013-09-18.

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

N/A - Storm is ready to graduate.

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

Both the Storm PPMC and the IPMC have voted to graduate. The graduation resolution has been sent to the board for approval.

How has the community developed since the last report?

The Storm community continues to grow, and the mailing lists are quite active. We have added 3 additional committers/PPMC members.

How has the project developed since the last report?

We have issued one release and are preparing for the next release.

Date of last release:

2014-06-24

When were the last committers or PMC members elected?

2014-08-26

Signed-off-by:

- (storm) Ted Dunning
- (storm) Arvind Prabhakar
- (storm) Devaraj Das
- (storm) Matt Franklin
- (storm) Benjamin Hindman

Shepherd/Mentor notes:

Wave

A wave is a hosted, live, concurrent data structure for rich communication.

It can be used like email, chat, or a document.

Wave has been incubating since 2010-12-04.

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

1. Make a release.
2. Grow community.
3. Develop new feature/fix bugs.

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

No.

How has the community developed since the last report?

A new committer joined us. There has been increased activity from existing members, and we have started receiving patches from a new developer.

How has the project developed since the last report?

Work has continued on generating our first release. We are now testing out RC6 - after RC5 failed due to a couple of missing licenses, and a bug in the MongoDB implementation. In the mean time, development work on the main branch has added profiling support, JDK 7 support, Solr as a search backend, and various bugfixes.

Date of last release:

No release yet.

When were the last committers or PMC members elected?

July 2014 - Andrew Kaplanov

Signed-off-by:

(wave) Christian Grobmeier
 (wave) Upayavira

Shepherd/Mentor notes: