

# August2014

/!\ \*FINAL\* (!)

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

---

## Incubator PMC report for August 2014

### Timeline

Wed August 06	Podling reports due by end of day
Sun August 10	Shepherd reviews due by end of day
Sun August 10	Summary due by end of day
Tue August 12	Mentor signoff due by end of day
Wed August 13	Report submitted to Board
Wed August 20	Board meeting

### Shepherd Assignments

Alan D. Cabrera	Optiq
Alan D. Cabrera	Twill
Andrei Savu	Flink
Andrei Savu	Hadoop Development Tools
John D. Ament	Sirona
John D. Ament	Slider
Justin Mclean	BatchEE
Konstantin Boudnik	Fleece
Konstantin Boudnik	Parquet
Matthew Franklin	Argus
Matthew Franklin	Droids
Roman Shaposhnik	ODF Toolkit
Raphael Bircher	Kalumet
Raphael Bircher	NPanday
Suresh Marru	Blur
Suresh Marru	Brooklyn

### Report content

Incubator PMC report for August 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 31 podlings under incubation this month. No projects are currently voting on graduation. The incubator added Argus as a new incubating project in July. Blur, DeviceMap, HDT, Metamodel, Slider, and Tez all had releases since the last report.

\* Community

New IPMC members:

Jan Iversen

People who left the IPMC:

\* New Podlings

Argus

\* Graduations

The board has motions for the following:

N/A - included in last month.

\* Releases

The following releases were made since the last Incubator report:

Apache Tez-0.4.1-incubating RC0  
Apache Slider 0.40-incubating RC0  
Apache Blur 0.2.3-incubating RC2  
Apache Metamodel incubating 4.2.0  
Apache DeviceMap BrowserMap incubating 1.4.1  
Apache HDT version 0.0.2-incubating (RC1)

\* IP Clearance

Brooklyn CAMP Server  
Sling Query contribution  
BigCouch updates to Apache CouchDB

\* Legal / Trademarks

It is unclear whether Blur can get the Trademark assign document provided by trademarks@ signed by Near Infinity. The company no longer exists.

\* Infrastructure

\* Miscellaneous

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

\* Still getting started at the Incubator

Argus  
Parquet  
Optiq  
Sirona  
Twill

\* Not yet ready to graduate

No release:

Brooklyn  
Fleece  
Flink  
Parquet

Community growth:

BatchEE

Blur  
Hadoop Development Tools  
NPanday  
ODF Toolkit  
Sirona  
Slider  
Twill

\* Require Mentors/IPMC intervention

Droids  
Kalumet

\* Ready to graduate

The Board has motions for the following:

(none)

\* Did not report, expected next month

(none)

---

#### Table of Contents

Argus  
BatchEE  
Blur  
Brooklyn  
Droids  
Fleece  
Flink  
Hadoop Development Tools  
Kalumet  
NPanday  
ODF Toolkit  
Optiq  
Parquet  
Sirona  
Slider  
Twill

---

#### 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
2. Publish initial code drop along with a code grant
3. Get approval for name of Argus (PODLINGNAMESEARCH-54)
4. Create a project website

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

As part of the discussion about entering incubator, we added a second mentor (Daniel Gruno) from outside of Hortonworks.

How has the community developed since the last report?

- \* podling accepted into incubator
- \* all committers have ICLA filed and accounts created
- \* Jira created
- \* status page created

\* git repository created  
\* mailing lists 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](argus) Alan Gates  
[ ](argus) Daniel Gruno  
[X](argus) Devaraj Das  
[X](argus) Jakob Homan  
[X](argus) Owen O'Malley

-----  
BatchEE

BatchEE projects aims to provide a JBatch implementation (JSR-352) and a set of useful extensions for this specification.

BatchEE has been incubating since 2013-10-03.

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

1. Increase community (surely implies communication)
2. Do more release (in progress)
3. Improve our documentation

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?

We start to get some interesting "real life" user feedbacks (mainly through JIRA).

How has the project developed since the last report?

Mainly bugfixes and some feature enhancement (goal is stability and ease of use).

Date of last release:

2014-04-01  
We are currently preparing for a new release.

When were the last committers or PMC members elected?

Signed-off-by:

[ ](batchee) Jean-Baptiste Onofré  
[ ](batchee) Olivier Lamy  
[X](batchee) Mark Struberg

-----  
Blur

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

Blur has been incubating since 2012-07-24.

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

1. It is unclear if we can get the Trademark assign document provided by trademarks@ signed by Near Infinity. The company no longer exists.
- 2.
- 3.

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

Other than the issue with the Trademark assignment listed above, no.

How has the community developed since the last report?

We continue to be small but active.  
- Subscriptions: user@ - 61[+7]; dev@ - 67[+6]

How has the project developed since the last report?

Apache Blur has released a stable version and has started development of a new feature set that should help grow the adoption of the technology. This new feature set is being driven through community interaction and development.

Date of last release:

2014-07-29

When were the last committers or PMC members elected?

2014-07-28

Signed-off-by:

[ ](blur) Doug Cutting  
[X](blur) Patrick Hunt  
[X](blur) Tim Williams

-----  
Brooklyn

Brooklyn is a framework for modelling, 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 organisation)

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?

No major changes - the community is still small, but is attracting a new names on the mailing list and IRC channel, including some code contributions.

How has the project developed since the last report?

Development is in full progress and we have settled in to our new home at the Incubator, and are preparing to make our first release.

Date of last release:

No releases under Apache yet.

When were the last committers or PMC members elected?

No change since last report, when two new committers/PPMC members joined (2014-07-02)

Signed-off-by:

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

-----  
Droids

Droids aims to be an intelligent standalone robot framework that allows to create and extend existing droids (robots).

Droids has been incubating since 2008-10-09.

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

1. Activity
2. Name Search

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

None. The previous goal of conducting the name search was not met due to time availability. Outlook is much more promising this quarter.

How has the community developed since the last report?

No change. Name search is one big outstanding piece to complete, and it is rather intimidating. It is hoped that one of the more experienced individuals will have time to complete it this quarter. That should help with participation once the final name is known.

How has the project developed since the last report?

No change.

Date of last release:

2012-10-15

When were the last committers or PMC members elected?

2012-05-07

Signed-off-by:

[x](droids) Thorsten Scherler  
[X](droids) Richard Frovarp

-----  
Fleece

Implementation of JSR 353: JavaTM API for JSON Processing.

Fleece 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".  
Proposed name is "Johnzon".
2. Expanding the community, increase dev list activity and adding new committers.
3. Publish first incubator release.

The Fleece project status page at <http://incubator.apache.org/projects/fleece> needs updating but that would be pending on #1 due to many items on the status-page need to be re-visited once #1 is settled.

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

- \* Project have to be renamed because of naming conflicts with "Fleece".
- \* Waiting for <https://issues.apache.org/jira/browse/PODLINGNAMESEARCH-53>

How has the community developed since the last report?

- \* There have been several patches submitted from a current non-committer and we're hoping to be able to mentor that person toward becoming a committer.

How has the project developed since the last report?

- \* Several issues were reported and fixed (mainly core performance and API compliance enhancements).
- \* Lots of new unittests added.

Date of last release:

- \* No releases as of yet.

When were the last committers or PMC members elected?

- \* N/A

Signed-off-by:

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

Shepherd/Mentor notes:

Mailing list is active, all seems to be normal and ok.

-----  
Flink

Flink is an open source system for expressive, declarative, fast, and efficient data analysis. Flink combines the scalability and programming flexibility of distributed MapReduce-like platforms with the efficiency, out-of-core execution, and query optimization capabilities found in parallel databases. Flink was originally known as Stratosphere when it entered the Incubator.

Flink has been incubating since 2014-04-14.

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

1. First incubator release
2. Continue with community growth
3. More comfortable with ASF policies and processes

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?

One new committer has been accepted into the project.  
We received a large external code contribution and pending reviews.  
The project started a bi-weekly Hangout.

How has the project developed since the last report?

The project has moved all infrastructure to Apache,  
the source code has been renamed to the "org.apache.flink" namespace.  
We are finalizing the preparations for the first Incubator release

Date of last release:

2014-05-31 (no incubator release yet)

When were the last committers or PMC members elected?

2014-07-10

Signed-off-by:

[ ](flink) Sean Owen  
[ ](flink) Ted Dunning  
[ ](flink) Owen O'Malley  
[X](flink) Henry Saputra  
[ ](flink) Ashutosh Chauhan

-----  
Hadoop Development Tools

Eclipse based tools for developing applications on the Hadoop platform

Hadoop Development Tools has been incubating since 2012-11-09.

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

1. Build Community
2. Get more 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?

- 133 mails exchanged on mailing list

How has the project developed since the last report?

- Migrated all the existing functionality from master which enabled us for next release.
- Released 0.0.2 incubating

Date of last release:

- July'2014

When were the last committers or PMC members elected?

- November 2013

Signed-off-by:



[ ](hadoopdevelopmenttools) Suresh Maru  
[ ](hadoopdevelopmenttools) Chris Mattmann  
[X](hadoopdevelopmenttools) Roman Shaposhnik

-----  
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 July, 14) 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:

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

-----  
NPanday

NPanday allows projects using the .NET framework to be built with Apache Maven.

NPanday has been incubating since 2010-08-13.

Since the last report in July:

- the build issues on the Windows Jenkins slaves were resolved
- further JIRA issues have been cleaned up to prepare for the release
- Konstantin Boudnik was formally added as a mentor for the project

While some progress was made towards getting the release out, there is still more work to do as time allows.

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

1. Ship the changes on trunk as a release and make it easier for new users to get up to speed when interested in the project
2. Get a critical mass of committers/PPMC members that can respond to contributors, apply patches, nominate committers, and vote on releases
3. Provide guidance on how to get involved in contributing

The incubator is already aware of the low level of activity, that many of the PPMC are now disengaged. There continues to be interest in a new release from users, and new interest from some in making contributions that continue to be followed up.

Date of last release:

2011-05-16

The last committer was added on 2011-04-19. There have not been any PPMC additions since inception.

Signed-off-by:

[ ](npanday) Raphael Bircher  
[X](npanday) Konstantin Boudnik

-----  
ODF Toolkit

Java modules that allow programmatic creation, scanning and manipulation of OpenDocument Format (ISO/IEC 26300 == ODF) documents

ODF Toolkit has been incubating since 2011-08-01.

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

1. Grow the community in terms of committers
2. Have regular releases
3. Decide if we would like to be a top level project or join an existing project

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

How has the community developed since the last report?

We issued a "call for volunteers" to our users list and got a good response: <http://s.apache.org/4Vx>

Of course, it will be easier to grow the community if we can reliably conduct a release vote.

How has the project developed since the last report?

2014-06-02 - ODF Toolkit 0.6.1-incubating available

Date of last release:

2014-06-02

When were the last committers or PMC members elected?

2012-10-29

Signed-off-by:

[ ](odftoolkit) Sam Ruby  
[ ](odftoolkit) Nick Burch  
[ ](odftoolkit) Yegor Kozlov  
[x](odftoolkit) Rob Weir

-----  
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: web site and nightly builds).
2. Re-organize code into org.apache.optiq namespace.
3. Regular releases.
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 have added our first post-incubation committer, Vladimir Sitnikov.

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:

- \* Gained access to Nexus and started publishing (on-demand, not nightly) snapshot builds.
- \* Web site in progress.

Working on first release under the Apache incubator, apache-optiq-0.9.0-incubating.

Date of last release:

No releases yet under the Incubator.

Last unofficial release was optiq-0.8, on 2014-06-27.  
First official release is imminent.

When were the last committers or PMC members elected?

Vladimir Sitnikov, elected committer on 2014-07-15.

Signed-off-by:

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

-----  
Parquet

Parquet is a columnar storage format for Hadoop.

Parquet has been incubating since 2014-05-20.

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

1. Expanding the community and adding new committers
2. 1st release
3. Identifying how to ensure timely code reviews by committers

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

None at this time

Latest Additions:

- \* PMC addition: N/A
- \* Contributor addition: N/A

Issue backlog status since last report:

- \* Created: 60
- \* Resolved: 17

Mailing list activity since last report:

- \* dev 212 messages

How has the project developed since the last report?

- \* New commit workflow has been documented and commits have been increasing using the commit script.
- \* Project website is posted: [parquet.incubator.apache.org](http://parquet.incubator.apache.org), working on moving more content from github hosting
- \* Moved to [issues.apache.org](http://issues.apache.org) for all new issues
- \* Planning first release of parquet-format and parquet-mr. Using parquet-format release to identify steps needed to release the larger projects (e.g., parquet-mr)
- \* Adding documentation on reviews and contacts for specific modules

Signed-off-by:

[X](parquet) Jake Farrell  
[ ](parquet) Chris Mattmann  
[ ](parquet) Roman Shaposhnik  
[X](parquet) Tom White  
[X](parquet) Todd Lipcon

Shepherd/Mentor notes:

Mailing list has a healthy traffic, mostly bug reports. Mentors are active and participating in the community.

-----

Sirona

Monitoring Solution.

Sirona has been incubating since 2013-10-15.

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

1. need to increase contributor number
- 2.
- 3.

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?  
Not a lot since last report as the project was a bit quiet.

Mainly some real usage feedbacks

How has the project developed since the last report?  
A release has been cutted and some discussion on reworking the ui.

We felt holidays but some new features are expected soon (path tracking for instance)

Date of last release:

2014-06-26

When were the last committers or PMC members elected?

Signed-off-by:

[x](sirona) Olivier Lamy  
[ ](sirona) Henri Gomez  
[X](sirona) Jean-Baptiste Onofre  
[ ](sirona) Tammo van Lessen  
[ ](sirona) Mark Struberg

Shepherd/Mentor notes:

Traffic on the list is light. There was a recent release, using new rules it looks like. I'm concerned over volume and availability of core devs.

-----

Slider

Slider is a collection of tools and technologies to package, deploy, and manage long running applications on Apache Hadoop YARN clusters.

Slider has been incubating since 2014-04-29.

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

1. Building a user community
2. Getting those users to contribute their work back
3. Improving the application to make it easier to use

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?

We're getting some independent users/developers on our list, including some initial minor contributions and some bug reports. As an example, we have some contributed logos ... which is now going to trigger a competition/vote for a slider logo.

The main external developer adoption is not in actual contributions for the core slider engine, but from a few people writing slider application-packages to deploy their own applications. This is a good first step - we have to work with them to get them through this process, note and address issues they have -as well as encouraging them to contribute a bit themselves. There's a tradeoff there: we don't want to dismiss all issues with "file a patch with the JIRA" -maybe there's an intermediate step of "let's work on this, can you take this branch and verify it works for you?"

How has the project developed since the last report?

1. We're still fast-iterating through releases.
2. We're testing slider against the forthcoming Hadoop 2.5 release, to help validate that release.
3. We're debating when - in August - to make the leap to only building against Hadoop 2.6

The switch to 2.6 will mean that this code would be developer only, as it would only work against non-production Hadoop clusters. Yet we need this for the forthcoming features: the Credential provider, service registry, Yarn node labeling and other features we need to integrate with.

That's a fairly exclusionary move for potential users: the current plan is to make a release in sync with Hadoop 2.5, which will be the last Hadoop 2.4/2.5 compatible version. Yet it not only gives us those new features, it helps tune those features and find bugs in their implementation. As an example, one fix in Hadoop 2.4 is related to an Application Master failure and restart problem which we'd found, a problem which, if we'd switched to Hadoop 2.4-SNAPSHOT earlier, would have been found and fixed in the 2.4 release.

Date of last release:

2014-07-21

When were the last committers or PMC members elected?

2014-05

Signed-off-by:

[ ](slider) Arun C Murthy  
[ ](slider) Devaraj Das  
[X](slider) Jean-Baptiste Onofré  
[ ](slider) Mahadev Konar

-----  
Twill

Twill is an abstraction over Apache Hadoop® YARN to reduce the complexity of developing distributed applications.

Twill has been incubating since 2013-11-14.

Top three items to resolve before graduation:

- More committers from different organizations.
- Regular Releases.
- Increase adoption.

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

- None at this time.

How has the community developed since the last report?

- 22 new JIRA issues filed since last report (2014-05-01)
- 13 JIRA issues resolved since last report (2014-05-01)
- Subscribers to the dev list up from 44 in May to 52 in August.
- 5 new contributors since last report.

How has the project developed since the last report?

- Version 0.3.0-incubating RC1 being voted upon

Date of last release:

- 2014-04-22: 0.2.0-incubating

What are the plans for the next period?

- Increase release frequency
- Engage active contributors and identify potential committers
- Improve documentation and website

When were the last committers or PMC members elected?

- No new committers since incubation.

Signed-off-by:

[ ](twill) Arun C Murthy  
[X](twill) Tom White  
[X](twill) Patrick Hunt  
[ ](twill) Andrei Savu