

StreamsProposal

Abstract

Apache Streams will be a lightweight server for [ActivityStreams](#). The role of Apache Streams is to provide a central point of aggregation, filtering and querying for Activities that have been submitted by disparate systems. Apache Streams also intends to include a mechanism for intelligent filtering and recommendation to reduce the noise to end users.

Proposal

Apache Streams will bring together individuals who are or are looking to increase and centralize the production, consumption and federation of [ActivityStreams](#) throughout enterprise organizations and the Internet as a whole. The target features include:

- Publication of Activities from multiple systems via HTTP
- Aggregation and syndication of streams
- Support for security trimming of streams by social graph
- Noise reduction and intelligent filtering
- Federation of streams across disparate systems
- Provide libraries for easy integration in source systems

Background

The [ActivityStreams](#) specification standardizes a generic format for describing event-based data. Many social web companies have adopted the format and it has been included in the [OpenSocial](#) specification as the preferred method for delivery of activity data. During discussions of [ActivityStreams](#) at [Open Social](#) events earlier in the year, it became clear that multiple organizations are facing similar issues as to the publication and filtering of their activities and would benefit from a commonly-developed, open-source [ActivityStreams](#) server.

Rationale

In deployment, activities are generated from multiple sources. This is particularly true within the enterprise where disparate systems create and manage activities in stove pipes. What is needed is a central point for consumption, aggregation and exposure of activities generated within these systems.

Initial Goals

The initial goal of the project is to survey donated code and develop a common, high-level architecture for an initial release. The project will then work toward that release in a new code base, pulling in pieces of donated code as necessary.

Current Status

The MITRE Corporation will donate its prototype code to the project. Aside from this, the project is new and will need bootstrap time to develop an initial architecture and roadmap.

Meritocracy

As a new project with many team members who are seasoned Apache veterans, Apache Streams is prepared to build the project under the Apache Way.

Community

The Apache Streams community combines multiple individuals from different organizations, many of which have collaborated before in an open environment. Apache Streams is committed to expanding this community and adhering to Apache principles of openness and collaboration.

Known Risks

An exercise in self-knowledge. Risks don't mean that a project is unacceptable. If they are recognized and noted then they can be addressed during incubation.

Inexperience with Open Source

Most of the initial committers have worked in open source and many are familiar with the ASF and the Apache Way; but, not all. Additionally, many of the committers have not worked on a software project together and will need time to familiarize themselves with each other.

Speed of Development

This project is dependent upon contributions made on company time. For this approach to succeed, the project must deliver a workable system in a timeframe acceptable to those companies. The initial parties have the intention of releasing a first version within 6 months after starting the Incubator. Failure to do so could prevent the project reaching critical mass, and could prevent the project from being in a position to attract new developers.

Reliance on Salaried Developers

At present, the vast majority of contributors will be doing so as a part of their day jobs. Therefore, as already alluded to, there is a risk that the project won't gain enough traction to be of use to their employers. However, given the centrality of these codebases to the participating companies, it is clearly in their best interests to transition to an openly developed alternative.

Relationships with Other Apache Products

Many of the initial committers will be integrating this software with Apache Rave & Apache Shindig. A possibility exists that code is developed by either or both of the projects to support integration with Apache Streams.

An Excessive Fascination with the Apache Brand

The rationale for the community to propose the project to Apache is that Apache offers a meritocratic governance style that is free from direct influence by any single company. This allows all of the participants in this proposal to collaborate regardless of which organization pays them.

Initial Source and Submission Plan

The MITRE Corporation has developed a prototype using Apache Shindig, Apache Rave and the [PubSubHubbub] Enterprise Engine (PuSH-EE) [1]. The source for this will be donated to the ASF upon successful execution of MITRE's open source technology transfer process.

[1] <http://code.google.com/p/pushee/>

Cryptography

The only cryptography included by the project will be via library inclusion. As it is currently under discussion at legal-discuss as to whether or not this qualifies for BIS registration, we propose that any actions to arise regarding cryptographic libraries be deferred until suitable resolution to the discussion is achieved.

Required Resources

Resources that infrastructure will be asked to supply for this project.

Mailing lists

- Apache Streams-commits
- Apache Streams-private (moderated subscriptions)
- Apache Streams-dev

Subversion Directory

<https://svn.apache.org/incubator/streams>

Issue Tracking

JIRA Apache Streams

Other Resources

No special resources or infrastructure required

Initial Committers

- Matt Franklin (mfranklin at apache dot org) - Apache Member
- Craig [McClanahan](#) (craigmc at apache dot org) - Apache Member
- James Snell (jmsnell at apache dot org?) - Apache Committer
- Ate Douma (ate at apache dot org) - Apache Member
- Ryan Baxter (rbaxter at apache dot org) - Apache Committer
- Jason Letourneau (jletourneau at gmail dot com)
- Beth Lavender
- Mark Weitzel
- Steve Blackmon
- Andrew Hart (ahart at apache dot org) - Apache Member
- Davide Palmisano

Affiliations

- The MITRE Corporation (Matt Franklin, Jason Letourneau, Beth Lavender)
- IBM (James Snell, Ryan Baxter)
- Jive Software (Craig [McClanahan](#), Mark Weitzel)
- Hippo (Ate Douma)
- WCG (Steve Blackmon)
- NASA Jet Propulsion Laboratory (Andrew Hart)

Sponsors

Champion

- Matt Franklin

Nominated Mentors

- Matt Franklin
- Ate Douma
- Craig [McClanahan](#)

Sponsoring Entity

- Apache Incubator