

UsergridProposal

Usergrid Proposal

Abstract

Usergrid is a multi-tenant Backend-as-a-Service stack for web & mobile applications, based on RESTful APIs.

Proposal

Usergrid is an open-source Backend-as-a-Service ("BaaS" or "mBaaS") composed of an integrated distributed NoSQL database, application layer and client tier with SDKs for developers looking to rapidly build web and/or mobile applications. It provides elementary services (user registration & management, data storage, file storage, queues) and retrieval features (full text search, geolocation search, joins) to power common app features.

It is a multi-tenant system designed for deployment to public cloud environments (such as Amazon Web Services, Rackspace, etc.) or to run on traditional server infrastructures so that anyone can run their own private BaaS deployment.

For architects and back-end teams, it aims to provide a distributed, easily extendable, operationally predictable and highly scalable solution. For front-end developers, it aims to simplify the development process by enabling them to rapidly build and operate mobile and web applications without requiring backend expertise.

Background

Developing web or mobile applications obviously necessitates writing and maintaining more than just front-end code. Even simple applications can implicitly rely on server code being run to store users, perform database queries, serve images and video files, etc. Developing and maintaining such backend services requires skills not always available or expected of app development teams. Beyond that, the proliferation of apps inside of companies leads to the creation of many different, ad-hoc, unequally maintained backend solutions created by employees and contractors alike and hosted on a wide variety of environments. This is causing poor resource usage, operational issues, as well as security, privacy & compliance concerns.

In response to this problem, companies have long tried to standardize their server-side stack or unify them behind an ESB or API strategy. Backends-as-a-Service follow a similar approach but their unique characteristic is strongly tying 1) a persistence tier (typically a database), 2) a server-side application tier delivering a set of common services and 3) a set of client-side application interface mechanisms. For example, a BaaS could package 1) MongoDB with 2) a node.js application that offers access through 3) [WebSockets](#). In the case of Usergrid, the trifecta is 1) Cassandra, 2) Java + Jersey and 3) a RESTful API.

The Backend-as-a-Service approach has steadily gained popularity in the last few years with cloud providers such Parse.com, Stackmob.com and Kinvey.com, each operating tens of thousands of apps for tens of thousands of developers. The trend has already reached large organizations as well, with global companies such as Korea Telecom internally building a privately-run BaaS platform. But so far, there have been limited options for developers that want a non-proprietary, open option for hosting and providing these services themselves, or for enterprise and government users who want to provide these capabilities from their own data centers, especially on a very large scale.

Rationale

The issue this proposal deals with is implicit in the name. Backend-as-a-Service platforms are usually offered solely as proprietary cloud services. They are typically closed sourced, hosted on public clouds, and require subscription payment. Usergrid opens the playing field, by making a fully-featured BaaS platform freely available to all. This includes developers that previously could not afford them, such as mobile enthusiasts, small boutiques, and cost-sensitive startups. This also includes large companies that benefit from a reference implementation they can deploy in trust, or extend to their needs without losing time writing less-vetted, less-performant boilerplate functionality.

Usergrid has been open source since 2011 and has grown as an independent project, garnering 11 primary committers, 35 total contributors, 260+ participants on its mailing list, with 3,700+ commits, 200+ external contributions, 350+ stars and 100+ forks on Github, not to mention several large scale production deployments at major global companies in the media, retail, telecommunication and government spaces.

The Apache Software Foundation's "Way", by putting community before the code, will help Usergrid establish a vibrant, more diverse community to provide these features freely to downstream users. The incubation process will help build this community and clear out the cobwebs, while vetting the IP to provide a pristine ASLv2 licensed product to users. Under such conditions, our hope is that Usergrid will have a brighter future, and provide better assurances as an independent organic open source project, instead of continuing forward as a project sponsored by a handful of companies.

Initial Goals

We have no serious codebase concerns at the present moment. Besides vetting the IP by making sure the dependencies are Apache License 2.0 compatible, our main initial concern is to grow community and keep adding features to make Usergrid as robust as possible. However some logistics based goals include:

- Move the existing code base to Apache
- Integrate with the Apache development process
- Ensure all dependencies are compliant with Apache License 2.0
- Set up open-source docs and website
- Incremental development and releases per Apache Guidelines

Current Status

Meritocracy

The project team's goals have always been to grow the community by encouraging contributors to participate. The project has grown steadily and smoothly from the efforts of the original creator and project founder (Ed Anuff) to a small circle of committers (at Apigee), to a distributed, multi-vendor community (Apigee and Korea Telecom) that also includes outside committers (Globo and others), as well as non-salaried committers. Together we discuss the project's goals and roadmap openly, making drastic yet positive changes to the project's direction based on everyone's input. Our goal is to drive further community diversification in a way that only a foundation-sponsored project can achieve, rather than what a vendor-led project can accomplish.

Community

We have a public Google Group for support here: <https://groups.google.com/forum/?hl=en#!forum/usergrid> with over 250 active participants, 367 threads and new messages every day. Usergrid also has an active community on Github issues (with over 200 discussions): <https://github.com/apigee/usergrid-stack/issues?state=all>. Some of these users have contributed their open applications back to the community or have built their own new SDKs for Usergrid found here on Github: <https://github.com/search?q=usergrid&source=cc>.

Expanding and nurturing the Usergrid community is our primary concern and one of the main reasons for the decision to apply for incubation. Usergrid has been developed openly on Github for many years and has enjoyed active developer participation by a committers from all over the globe. However, due to the disparate nature and wide variety of the Github repos that comprise the Usergrid project (the Usergrid-stack, 8 separate SDKs, the Admin Portal, and various samples), it has been challenging to grow a unified community. Bringing the entire project under the umbrella of Apache will promote a unification of the Usergrid community and enable all developers to collaborate on the project.

Core Developers

The core developers include Apache Committers, PMC Members, and Members of the ASF. The developers, some of whom have have been involved with the Apache Incubator and Apache Lucene as PMC members, are active mentors and have participated in and contributed to several projects: i.e. Apache, Lucene, Cassandra, Hibernate, Directory, Wicket, Commons, Roller, MINA, Karaf, Felix, Cloud Stack, HCatalog, and Commons projects. Many of them are also active in Open-source beyond code, and have positions on the committees or organization such as OSCON.

Alignment

The initial code base leverages several Apache Software Foundation products. Usergrid leverages Apache Cassandra for its scalable data store, and uses Maven for its build system. Almost half of Usergrid's dependencies are Apache dependencies:

- Apache Cassandra
- Apache Tomcat
- Apache Commons
- Apache [ZooKeeper](#)
- Apache Shiro
- Apache Amber
- Apache Thrift
- Apache ActiveMQ
- Apache [HttpClient](#)
- Apache Lucene
- Apache JClouds

Besides these direct alignments, Usergrid also complements Apache Cordova and could provide several advantages to the mobile application developer community they serve.

Known Risks

Orphaned products

There are now at least two vendors running Usergrid in product. Apigee is an established startup with a large, diversified customer roster and Korea Telecom is a major, national telecommunications company. The continuity of Usergrid, as an open-source, vendor-independent product are in the interest of all parties. Beyond the vendors, Globo.com and many others large companies have been relying on Usergrid for critical applications and as such they are committed to contributing to the effort.

Inexperience with Open Source

The Usergrid project has been open source and under the ALv2 for 2 years on Github and many of its contributors came with previous open-source experience, (as referenced above), including active members of these communities:

- Apache
- Cassandra (& Hector)
- Lucene
- Hibernate
- CouchDB
- [PhoneGap](#)
- jQuery

Development in this open forum has resulted in a growing community of contributors, and the Usergrid project is now ready and eager to embrace and learn from Apache's wealth of experience. Usergrid would like to embrace an even greater culture of open participation as witnessed on so many Apache projects.

Homogenous Developers

The core development team for Usergrid is a geographically and technologically diverse group. Apigee's team is itself distributed, with contributors based in each timezone in the continental US. Additional regular contributors have joined us from India, Asia, Oceania, South America, the Middle East and Europe. While roughly half of our core developers come from a Java background, the other half is comprised of iOS, Ruby, and [JavaScript](#) developers.

Reliance on Salaried Developers

Most of the principal developers are paid by their employers to contribute, but not all. Throughout the life of the project, we've seen passionate, personal commitment from all parties, as evidenced by our commit distribution on weekends (<https://github.com/apigee/usergrid-stack/graphs/punch-card>). We also believe, given the growing interest in mobile API services and the range of individuals and corporations that are eager to participate, that non-salaried contributions will grow. We know the "The Apache Way" will help us further accelerate this process.

Relationships with Other Apache Products

There's much potential for collaboration with Apache Cordova and, of course, the Cassandra community because of the underlying foundations of Usergrid's scalability. In the future there may be more interactions with any of the communities that Usergrid has direct dependencies to.

A Excessive Fascination with the Apache Brand

Although we are aware of the strength of the Apache brand, we are primarily interested in the transforming power of the Apache Way to help guide Usergrid towards a more diversified and meritocratic community. To that end, the brand's primary benefit for us is to help to attract more participants and diversify the community. Having several committers, PMC participants, and members of Apache as developers on Usergrid, there's little infatuation with the brand, and the Usergrid community is actively conscious of this not being a driver for joining the Apache community.

Documentation

Information on Usergrid can be found at: <https://developers.apigee.com/app-services>.

Initial Source

All initial sources can be found here: <https://github.com/usergrid>

Source and Intellectual Property Submission Plan

The IP transfer for Usergrid is trivial due to it's single source and existing ASLv2 licensing.

External Dependencies

Most dependencies are Apache compatible licenses (Category A). A small set of Category B licenses, like the CDDL exists. For more details please see [Dependency Licenses](#).

Cryptography

Not relevant to Usergrid since all code dealing with cryptography already comes from the JDK or from dependencies on Apache Software.

Required Resources

Mailing lists

- private@usergrid.incubator.apache.org (moderated)
- dev@usergrid.incubator.apache.org
- commits@usergrid.incubator.apache.org

Subversion Directory

We prefer to use Git as our source control system: [git://git.apache.org/usergrid/](http://git.apache.org/usergrid/). If possible, we would like to keep leveraging the extremely useful github facilities for workflow using a process much like that employed by the Apache Cordova project (documented here <http://wiki.apache.org/cordova/ContributorWorkflow>).

Issue Tracking

JIRA Usergrid (USERGRID)

Other Resources

None.

Initial Committers

Below are the initial committers for the project. This list will be used by the mentors to setup initial accounts. The podling will learn how to recruit additional committers during the incubation period and all interested contributors are encouraged to show up and start contributing once the podling has infrastructure in place.

- Alberto Leal <albertonb@gmail.com> (Globo.com)
- Dave Johnson <snoopdave@apache.org> (Apigee)
- Ed Anuff <ed@anuff.com> (Apigee)
- Nate [McCall](#) <zznate.m@gmail.com> (The Last Pickle)
- Rod Simpson <rod@rodsimpson.com> (Apigee)
- Scott Ganyo <scottganyo@apache.org> (Apigee)
- Shaozhuang Liu <stliu@hibernate.org>
- Sungju Jin <sungju@softwaregeeks.org> (Korea Telecom)
- Tim Anglade <timanglade@gmail.com> (Apigee)
- Todd Nine <todd.nine@gmail.com> (Apigee)

Affiliations

- Apigee
- Korea Telecom
- Globo.com
- The Last Pickle

Sponsors

Champion

- Dave Johnson <snoopdave@apache.org>

Nominated Mentors

- Dave Johnson <snoopdave@apache.org>
- Jake Farrell <jfarrell@apache.org>
- Jim Jagielski <jim@apache.org>
- Lewis John McGibbney <lewis.mcgebby@gmail.com>
- Luciano Resende <lresende@apache.org>

Sponsoring Entity

Incubator PMC