

Web20KitProposal

Web2.0Kit Proposal

New Name Needed

Web2.0Kit sounds too much like [WebKit](#) which this is not related to.

Current names being considered include:

Ketero Olio

Potpourri [MixMatch](#)

Abstract

Web2.0Kit is a web 2.0 toolkit to help developers evaluate the suitability, functionality and performance of various web technologies by implementing a reasonably complex application in several different technologies.

Proposal

Web2.0Kit will develop an example application to understand the benefits, performance, and scalability of popular web technologies. Multiple implementations of the application are planned - each providing the same functionality but staying true to the philosophy of its base language/framework.

Background

Most web 2.0 sites today use open source languages and frameworks such as PHP, Ruby on Rails, and Java EE to develop their applications. Deployments of these applications also use popular open source servers such as Apache httpd, Tomcat, MySQL, Memcache, and Glassfish. Many other servers/technologies such as lighttpd, mogileFS, mongrels, JRuby are also gaining popularity.

With the myriad technologies available, it is not easy to understand how they differ, especially in terms of performance and scalability. With varied levels of documentation available for some open source applications, it is also quite difficult for a web 2.0 startup to understand the correct usage of these technologies so that they don't become a bottleneck as their site grows.

Rationale

Web2.0kit is a toolkit that will attempt to address the above issues.

What it does

Web2.0Kit defines an example web 2.0 application (the initial implementation uses an events site somewhat like yahoo.com/upcoming) and provides three implementations: PHP, Java EE, and Ruby on Rails. The toolkit will also define ways to drive load against the application in order to measure performance.

As developers join the project, they can implement the same application using their favorite web frameworks and compare their implementations to others.

What you can learn from it

a) Understand how to use various web 2.0 technologies such as AJAX, memcached, mogileFS etc. in the creation of your own application. Use the code in the application to understand the subtle complexities involved and how to get around issues with these technologies.

b) Evaluate the differences in the implementations: PHP, Ruby on Rails, Java EE, and other contributed implementations to understand which might best work for your situation.

c) Within each language implementation, evaluate different infrastructure technologies by changing the servers used (e.g: apache vs lighttpd, MySQL vs PostgreSQL, Ruby vs Jruby etc.)

d) Drive load against the application to evaluate the performance and scalability of the chosen platform.

e) Experiment with different algorithms (e.g. memcache locking, a different DB access API) by replacing portions of code in the application.

A robust, community-developed standard implementations of a web 2.0 application using different technologies will enable developers to compare and contrast these technologies in a manner that does not exist today. By providing excellent sample implementations of a concrete application that is available to everyone, we will enable faster and easier application development for users. Although we list three implementations in this proposal, we encourage others to come up with many more using other language stacks and/or frameworks e.g. Spring framework, Python etc.

Current Status

This is a new project with some sample not-ready-for-prime-time code.

Meritocracy

The initial developers are very familiar with meritocratic open source development, both at Apache and elsewhere. Apache was chosen specifically because the initial developers want to encourage this style of development for the project.

Community

Web2.0Kit seeks to create developer and user communities during incubation.

Core Developers

The initial core developers are Sun Microsystems, Inc. employees, and faculty and students at UC Berkeley. We hope to expand this very quickly.

Alignment

The developers of the [Web20Kit](#) want to work with the Apache Software Foundation specifically because Apache has proven to provide a strong foundation and set of practices for community-based development.

Known Risks

Orphaned products

This project has a lot of enthusiasm among the core developers, has ongoing development, and is not orphaned.

Inexperience with Open Source

The initial developers are well-versed in open source methodologies and practices.

Homogenous Developers

The initial group of developers is from two organizations. We would like to expand this and that is a primary reason for bringing this project to Apache.

Reliance on Salaried Developers

Although part of the initial development team are students, the core developers are employed by Sun Microsystems.

Relationships with Other Apache Products

None in particular, except that Apache HTTPD is the most common place to run PHP, and which the initial PHP implementation uses.

A Excessive Fascination with the Apache Brand

We believe in the processes, systems, and framework Apache has put in place. The brand is nice, but is not why we wish to come to Apache.

Documentation

Initial Source

Sun Microsystems Inc. intends to donate code for their PHP implementation of the sample events application as well as code to drive load against the application. UC Berkeley intends to donate code for the Ruby on Rails implementation.

This code is still a work in progress and will be provided primarily as a starting place for a much more robust, community-developed implementation.

External Dependencies

Required Resources

Developer mailing lists web20kit-dev@incubator.apache.org web20kit-commits@incubator.apache.org web20kit-private@incubator.apache.org

A subversion repository

A JIRA issue tracker

Initial Committers

Akara Sucharitakul <akara.sucharitakul@sun.com>
Shanti Subramanyam <shanti.subramanyam@sun.com>
Sheetal Patil <sheetal.patil@sun.com>
Binu John <binu.john@sun.com>
Kim Lichong <kim.lichong@sun.com>
William Sobel <wsobel@eecs.berkeley.edu>
Arthur Klepchukov <avk@berkeley.edu>
Craig Russell <craig.russell@sun.com>

Sponsors

Champion

Craig Russell <craig.russell@sun.com>

Nominated Mentors

Craig Russell <craig.russell@sun.com>
Henning Schmiedehausen <henning@schmiedehausen.org>
Matt Hogstrom <hogstrom@us.ibm.com>

Sponsoring Entity

The Apache Incubator.