

The Road to Brazil

aka Design and Architecture of the Brazil Web Application Framework

Stephen A. Uhler

Introduction by Stephen A. Uhler

The Brazil project defines a new concept in web service architecture. It views the web, not just as a sea of browsers talking to large content sites, but instead as a network of services and proxies where content passes through a sequence of processing steps throughout the network, is filtered, combined with other content and modified along the way. This architecture encourages small content sources to become "web enabled", such as devices with attached sensors and actuators. The content is not intended to be consumed directly by user agents, but combined with content from other sources by powerful "meta servers".

The Brazil project has only been around for a short time; however many of the concepts embodied by Brazil have been cultivated from seeds, sown as part of other research projects by the members of the Brazil project team. Somewhere along the way, while trying to convince the IT organization of an important customer that, yes, it is possible to build enterprise applications using a web browser as the only user interface, we were reminded of a scene from a movie involving the Ministry of Information, and the name, Brazil, like the movie, stuck.

The Brazil project is a sample implementation of this architecture that provides the infrastructure to create three classes of systems. The first class of systems, "micro-servers", enable web-based access to small content sources by defining a UPI, or URL Programming Interface that maps device capabilities into HTTP URLs. "Meta-servers", the next class of Brazil systems, combine and filter content from other content sources. This allows bits of data and control instructions from devices, via "micro-servers", to be integrated seamlessly into the content of traditional web sites. Finally, The Brazil project may be used to create traditional web sites, but with the flexibility to add dynamic filtering, content aggregation and personal preferences.

The Brazil project has advanced the state of the art in web services in three key areas:

- the separation of presentation from content
- the architectural reformulation from a monolithic client-server style system to a dynamic data-flow component architecture
- the retrieval and semantic extraction of content from existing and legacy systems, and its subsequent integration to create new presentation

As the variety and sophistication of web users and services increases, so does the coupling between the data presentation and the programming logic required to extract and generate the raw data. The Brazil project implementation defines a novel way to separate content from presentation, using an XML-based scripting language, that allows the same content to be reformatted to suit a variety of user needs, without requiring any modification of the content generation code. As an example, the same content can be dynamically reformatted for use by a desktop browser, a web enabled cellphone, or another Web service that requires XML.

Motivated by the wide range of computing systems that participate on the web, from \$50 web-enabled sensors to \$10M+ servers, the Brazil project defines a scalable component model of web services. These services are created by combining simple components in consistent ways, thus eliminating the complex API's and Interfaces that are typical of today's monolithic systems. Only

the required components need to be included in any particular web service, allowing deployment to small environments, while retaining the richness of capabilities for large ones.

While the Brazil project architecture looks toward the future of web services, it provides facilities to web-enable existing and legacy applications now, thus extending their useful lifetimes. This is accomplished by adding Web service front ends to the existing systems, and using the dynamic reformatting capability mentioned above to convert the operations of existing systems into web services.

The Brazil Project has already made significant contributions to Sun's products. The prototype version of sun.net, especially the reverse proxy – a key component of the sun.net system – was built using the Brazil project component architecture. The initial sun.net applications, such as the calendar, demonstrated the feasibility of creating as web services, applications previously considered only viable as desktop applications.

As the web services infrastructure matures, we expect the Brazil project to continue to contribute to Sun's success, by providing a model by which the existing commercial systems can evolve, both to provide better separation of content from presentation, more data integration and aggregation from different sources, and methodologies that will permit these systems to scale down as well as up.

A sample website built using Brazil technology, demonstrating some of its features, including complete documentation, and containing both runnable code and Java™ language source code is available at: <http://www.experimentalstuff.com>.