

HPC and the Web Birds-of-a-Feather Discussion

November 19, 2013, Denver, Colorado; Moderators: Annette Greiner, NERSC; Rion Dooley, TACC; and Shreyas Cholia, NERSC

The BOF consisted of guided discussion about the role of the web in high-performance computing. The three moderators introduced the session and then each shepherded one general topic area. The scope was intentionally broad, including science gateways but also APIs, UIs, and any idea that contributes toward bridging the worlds of high-performance computing and web technologies.

What is the state of the art?

The commercial, consumer software world has embraced the web, but HPC has lagged in leveraging of web technologies. There is general interest among attendees in making better use of the web to make science and analytics easier. Many are currently using data generated at the command line and creating simple web portals for users to query the data. Doing more than that is a struggle, even for those who are nominally infrastructure providers. The necessary funding and resources to make a successful HPC portal are often underestimated and unavailable. One way to lower the cost per portal is to use APIs, but adoption rates of APIs have been slower than expected. This may be due to poor documentation or failure to target the right users. Targeting HPC web APIs means aiming for a very specialized group of developers: those with expertise in both HPC and web technologies.

There was a call for dialog with developers and creation of new standards. One participant felt that standards made work more difficult, while others pointed out that things like the DOM and libraries that allow its manipulation make their work much easier. Predicting what will become the standard can be a challenge, though.

For identity management, companies like Google, Facebook, Twitter, and Amazon may become the de facto standard. Some attendees favor REST-based protocols like OAuth and OpenID Connect. Another option is XML authentication and encryption, emphasizing data-centric security rather than network security. The overhead of encrypting and decrypting is an issue for HPC, however, though users are increasingly avoiding shipping data.

Several tools for onboarding many users or enabling authentication across many portals were mentioned: Eduperson profile (<http://macedir.org/specs/eduperson/>), SimpleSAMLphp (<http://simplesamlphp.org>), and Shibboleth (<http://shibboleth.net>). As for defining a global profile, participants felt that this was not the right model, since it assumes a global mechanism to manage identities. Centers need to decide what identity providers they trust. TAGPMA (<http://www.tagpma.org>) does something like this, accrediting authentication providers.

Possible Future Capabilities and Low-Hanging Fruit

Job submission is something all centers deal with. Many feel that in exposing such services, the more you can hide the HPC/cluster semantics the more adoption you'll get. But there is some disagreement about what the user base really needs. We enable domain sciences, so users have little reason to learn the HPC domain, yet users do ask for training on how to use HPC systems, and sometimes their lack of knowledge leads to problems, like backed up queues. Allowing users to concentrate on their science may be a good value proposition: recapture the talent.

So, one question is how we can best help each other to implement new capabilities. Do we work together to build shareable versions of tools that will make development easier, or do we put our effort into trying to get agreement on how things should be done by everyone? What many of us would like to see is a JQuery for science, enabling scientists to more easily build their own apps without building the back end themselves. HTML5 widgets, R Studio, and IPython notebook can be useful building blocks. Another question is how we can build new web tools at reasonable cost. The way forward seems to be service based. Some APIs already exist (e.g., Agave, <http://agaveapi.co>; NEWT, <http://newt.nersc.gov>), but people need help building their own.

Some participants just need help determining how to get started on projects. People want details about how others have implemented successful tools. It would be nice to have practitioners on call who can help others get started. Workshops would also be helpful. (XSEDE did one, which resulted in the Science Gateway Cookbook: <https://www.xsede.org/web/gateways/gateways-cookbook>.) We could use a lot more of these things, more engagement, and more outreach.

Outreach

Particularly useful items for getting started would be how-tos and hello worlds. TodoMVC shows how to make the same app in many javascript frameworks; it would be neat to have the equivalent for science gateways.

Outreach to users is in part dependent on answering the question "why do HPC on the web?" Existing users are often content to work at the command line, but there are clear problems with keeping only that model. There is a huge community of potential users who are left out. There are few resources to help them, and their time would be better spent doing science. We need to figure out how to communicate the value proposition to funding agencies. We also need to make interfaces that are commonly needed, like visualization or data upload (though Globus is a great option), or find a "killer app." Getting domain experts working with UI and back-end developers/designers could be helpful in getting traction. Different user groups may have different needs. One trend we're seeing is movement toward data driven science. This may point us toward a different value proposition and different user expectations.