OSG XSEDE Identity Management Collaboration Meeting Report

Meet dates: June 7-8, 2012

Von Welch (vwelch@indiana.edu)

July 18th, 2012

1 Introduction

Over the course of two days, representatives from the OSG and XSEDE projects met at NCSA to discuss collaboration in identity management, specifically opportunities for collaboration with ongoing efforts by both projects to establish certificate authorities (CAs) and opportunities for longer-term collaboration in identity management (IdM).

In-person participants were: Mine Altunay, James Basney, Lothar Bauerdick, Randal Butler, Brooklin Gore, Miron Livny, and Von Welch. John Hoover and James Marsteller participated briefly via phone.

2 Summary of Project IdM Approaches

Differences between the approaches to identity management between the two projects were a frequent topic of exploration. In summary:

XSEDE has a process built on top of its allocations process that it has inherited from TeraGrid and the Supercomputing centers going back decades. A PI submits a proposal for computational resources and, assuming acceptance, is entered into the XSEDE central database (XCDB). A password is created for them during this process, which is also stored in the XCDB. Their allocation is referred to as a *project* and they have the authority to add other users to their project, who are added to the XCDB and similarly given passwords. XSEDE organizations (NCSA, TACC, etc.) operate a number of CAs and they all leverage the XCDB (or a local database derived from it) to authenticate users. As a result, these CAs all follow either the MICS or SLCS profiles of the IGTF. In general, users in XSEDE are not required to use certificates, with passwords and SSH keys being other options. Certificates are integrated into the user portal via a MyProxy CA.

Science Gateways are a relatively recent introduction to XSEDE, coming into being during the last phase of TeraGrid, where user access is facilitated through their use of gateways, which authorize users and provide identity information through non-traditional means to resource providers.

Identity Management in the OSG is driven by the OSG Community, which is comprised of virtual organizations (VOs). VOs set IdM requirements and drive the selection of technologies. The OSG Project provides a set of supporting services to

support the community, including registration authorities and the process of leveraging the DOE Grids PKI. User enrollment in OSG is initiated by VOs, they act as registration authorities to obtain certificates for users who do not already have one from the DOE Grids PKI, and authorize users to access the OSG by enrolling them in the VO's virtual organization membership service (VOMS). Access to OSG resources has historically solely been via certificates, though recent pilot job systems, similar at least in concept to Science Gateways, are changing this.

How advancements are accomplished:

There were further relevant differences between the projects, including in how future needs and solutions are assessed. John Towns described XSEDE's planning process based on gathering use cases representing desired user actions, and then distilling those through a process of selecting a set to be supported, and then into an architecture.

Miron Livny described an approach based on understanding principals about the relevant topic, engaging with specific communities to understand their needs, and then making incremental advancements through experimental deployments ("scenarios" in XSEDE terminology).

3 XSEDE CA Status

As described previously, members of the XSEDE project operate a number of CAs, both traditional and issuing short-lived certificates. All of these CAs utilize XSEDE's allocation process and the XSEDE Central Database as their registration authority process. The challenge faced by XSEDE is the difficulty maintaining the existing CAs and needing to support new paradigms such as campus bridging. XSEDE is currently gathering use cases and requirements for its traditional use cases and the new paradigms before deciding on a course of action.

4 OSG CA Status

OSG is in process of transitioning from the DOE Grids PKI operated by ESnet to an internal service provided by the OSG project by February of 2013. A planning phase has been completed, DigiCert LTD. has been contracted to provide the CA service and development on an OSG-operated front end, policies and procedures has commenced. An initial version of the front-end connected to a test CA at DigiCert is available for early friendly testers.

5 Next Steps

During the meeting, no opportunities for collaboration on current CA activities were identified. This was primarily due to two factors:

1. Differences in timeline: OSG is 6+ months into a development path with a firm deadline before the DOE Grids PKI ceases providing service while XSEDE is just beginning their planning process.
2. Differences in IdM Architecture: As described previously, OSG and XSEDE have difference approaches to IdM that drive differences in CA implementation and policy that are not immediately obvious how to reconcile.

A number of topics for potential future work and/or collaboration came up during the meeting.

5.1 A Central User Database for OSG?

There was some discussion if OSG would benefit from a central database similar to the XCDB.

- What is the effort level? Hard for XSEDE to say since the XCDB handles many aspects other than IdM (e.g., accounting).
- OIM will be approximating a central user database after the PKI transition, with an increasing fraction of OSG users registered in it. However, it will be lacking OSG-specific user authenticators; all authentication will be X.509 and no mechanism for authenticating users outside of that exists in OSG today.
- What would be the advantages to OSG for doing this? It would allow OSG to implement CA relying on authentication databases, for example, a MyProxy-based SLCS services.

5.2 Comparison of OSG VO and XSEDE Science Gateway

Conceptually OSG VOs and XSEDE Science Gateways seem very similar. A question that came up was "Could one swap the two?" This would seem to be the way to sort out the differences. However it's not clear this is more than an intellectual exercise at this time.

5.3 A National Host Certificate Service

The two projects, at least pending any surprises in XSEDE's ongoing analysis, seem very much aligned in their need for host certificates. There may be potential for collaboration in either establishing such as service or working with an existing service (DigiCert, InCommon Certificate Service, etc.) to allow it to fill those needs.

5.4 Leverage Campus IdM Interactions

Both projects are increasingly interacting with campuses through federated identity and other interoperability projects (BOSCO, Campus Bridging). However a constant challenge is getting research needs heard over competing requirements. Does it make sense to collaborate on campus interactions to have a single voice with more weight?

5.5 “Everything is a VO”

A resource provider granting access is to varying degree delegating of authority to further grant access. This is formalized in many cases, for example, science gateways select and authenticate their users and give them access to limited applications on the resource provider. XSEDE PIs have the ability to add users to their projects, though the resource provider maintains the responsibility for authentication.

In general though, resource providers are trusting users to grant access according to their rules. Only policy keeps a users from launching glide-ins, SSH deamons, etc. that give direct access to other users of their choosing.

5.6 The Intersection of IdM Domains

When one domain relies on the IdM of another domain, a trust relationship is established which has many aspects. The most obvious is the providing of identity from the identity provider to the service provider and the understanding of the acceptable use of that information (InCommon is focused on this aspect). This is a space that is not completely explored and other aspects include incident response, access to historical information about users, etc.

6 Closing Thoughts

The meeting topics flowed rather organically and it took a great deal of effort by the participants to keep things out of technical weeds and on topics relevant to both projects. Maintaining discussions at a level that was relevant to both projects did however have the benefit of discussions concepts and goals rather than implementation details.