Identity Management Guidance to OSG Virtual Organizations and Resource Providers

January 2015 (v2)[1]

Von Welch, Robert Cowles, Craig Jackson
Comments to vwelch@iu.edu
Extreme Scale Identity Management Project (XSIM)

Introduction

Identity management (IdM) is the ongoing process of managing members of a virtual organization (VO), how they are authenticated, and what privileges they have in the context of the VO. In the early days of scientific computing, identity management was handled entirely by the organization providing computational resources. As scientific collaborations increased in both number of people and magnitude of computing requirements, they needed to obtain resources from multiple resource-providing organizations. The VO emerged to manage this set of relationships. Identity management was then distributed among the VO and its resource providers.

This distributed approach to IdM is particularly relevant to the Open Science Grid (OSG), which is based firmly on the VO concept. There are many different ways that IdM can be distributed between the VO and the resource providers in the context of OSG, and this document provides guidance to those VOs and resource providers as to how to manage this distribution.

This work is based on the VO IdM model as created by the DOE-funded XSIM project [1,2], and is part of the XSIM project's applied research. Your use of this document and, especially, your feedback to the authors is greatly appreciated.

The Transitive Trust Approach

Figure 1 depicts an increasingly common and often desirable approach for VO identity management (IdM): *transitive trust* [3]. In this approach, the VO manages its community and RPs trust the VO to do so with little-to-no cognizance of the individuals, seeing them only as a members of the VO community.

This approach is often desirable because it produces a clear separation of responsibilities between the VO and resource provider, establishes a simpler workflow, and reduces administrative overhead in relatively low risk environments. The VO does not need to communicate information about all of its members to the resource providers and can utilize any mechanism for managing their identities it desires (i.e., a mechanism that meets any agreed-to assurance level with the resource providers).

[1] Originally published in June 2014. Version 2 updates the Reference #4.

Figure 1: Transitive trust approach with VO managing its members
and resource providers trusting the VO to do so.
No direct trust relationship between members and RP.

However, the transitive trust approach does introduce some issues which must be considered:

- *Lack of persistent personal data storage at resource provider*. Data that is either shared by the VO or temporary to a specific compute job can be stored by an RP readily because the lifetime of the data corresponds to the lifetime of the VO or compute job respectively. However, storing persistent data that is private to individual VO members is a challenge because the RP isn't aware of individual VO members from an IdM perspective. Addressing this challenge typically entails the VO arranging for any member-specific data to be migrated between VO and RP storage before and after a computational job (often referred to *staging* of the data).
- *VO-hosted collaborative services.* Many collaborative services (e.g., source code repositories, discussion forums, data storage) expect user identities to function. Because resource providers are not participating in individual user management, these services need to be hosted by the VO or a party (perhaps a individual resource provider) acting on the VO's behalf. Where the VO provides a portal as its primary user interface, this is a common place to host such services.
- *User support and incident response coordination*. During the course of normal operation, unexpected or adverse events will happen at the resource provider in servicing requests from the VO. Because the resource provider has no ability to contact individual VO members, RP and VO should have an agreement in place to handle these events -- e.g., an expectation that the resource provider can report events to the VO, who will handle them in some reasonable amount of time. OSG Document 1149 explains the security requirements to execute user jobs submitted without an end user certificate. [4]

There are variations of transitive trust implementations. Some examples found in the OSG include the following:

- Grid Laboratory of Wisconsin (GLOW) [5] is a campus distributed computing environment at the University of Wisconsin-Madison. Computing jobs from GLOW can be submitted to the OSG [6]. GLOW in this case is analogous to an OSG VO

and is an example of how the transitive approach can be used to connect different infrastructures.

- OSG and XSEDE offer a service which allows the submission of jobs to the OSG from the XSEDE infrastructure. Similar to the previous example of GLOW, XSEDE is analogous to a VO and the transitive approach connects the two infrastructures.
- OSGConnect [9] is an example of a web portal representing multiple VOs. For identity and group management, it uses CILogon and Globus Nexus, and then uses its trust relationship with OSG to provide computational resources to VO members.
- Site-override of VO Authorization is a modification to a pure transitive trust implementation. It allows an RP to take part in a transitive trust approach, but still veto individual VO members if needed, and does not require the RP to know any of the VO members besides those it decides to veto. The VO, after making its authorization decision, communicates the member identity to the RP, who can decide to veto the VO's authorization. An example of this is the use of the gLExec authorization callout in DIRAC [10].

The Brokered Trust Alternative

Figure 2 shows an alternative to the transitive trust approach for VO identity management (IdM): a *brokered trust* approach. Here, the resource provider is the primary manager of VO member identities, but the VO acts as a mediator, brokering the relationship between the resource provider(s) and members, and establishing who should be a member.

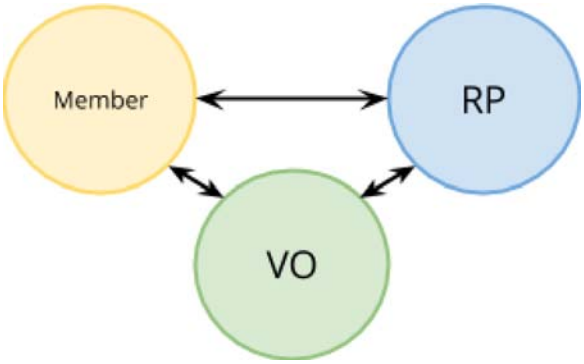

Figure 2: Brokered trust model with resource provider managing members and the VO establishing who is a legitimate members and establishing the relationship.

While this approach offers some advantages under certain circumstances, it presents a significant challenge: The resource provider must have knowledge of each VO member. This, in turn, produces ongoing dependencies between the VO and RPs, including requirements that they use compatible authentication technologies and communicate member identity information. It is only recommended when specifically warranted, including situations where:

- Resource providers have an explicit need to be aware of individual VO member identities.
- There is a need to store persistent data or other state at the resource provider that is differentiated by individual VO members (e.g., the data or systems have access control permissions that vary from member-to-member).

The most common way brokered trust is implemented in the OSG is via the VOMS [11] and VOMRS services [12]. VOs manage their communities in these services and resource providers either regularly pull a list of members from VOMS or members can obtain and provide assertions of their membership in real time.

Feedback and For More Information

For the paper with the full XSIM VO IdM model and the latest information from the XSIM project, please see http://cacr.iu.edu/collab-idm. This document is part of the XSIM project's applied research. Your use of this document and, especially, your feedback to the authors is greatly appreciated. Virtual organizations or resource providers with questions about identity management may contact the XSIM project PI, Von Welch, at vwelch@iu.edu

We thank the US Department of Energy Next-Generation Networks for Science (NGNS) program (Grant No. DE-FG02-12ER26111) for funding this effort.

References

1. Robert Cowles, Craig Jackson and Von Welch. Identity Management for Virtual Organizations: A Survey of Implementations and Model (draft version). 9th IEEE International Conference on eScience, 2013. http://www.vonwelch.com/pubs/VOIdM13
2. Robert Cowles, Craig Jackson, Von Welch and Shreyas Cholia. A Model for Identity Management in Future Scientific Collaboratories. International Symposium on Grids and Clouds (ISGC) 2014, 2014.http://www.vonwelch.com/pubs/XSIMISGC2014
3. NIST Special Publication 800-39, Managing Information Security Risk, March 2011, p.G-1-G-2.
4. OSG Document 1149, Traceability Requirements for end user jobs without certificates. http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=1149
5. GLOW: Grid Laboratory of Wisconsin. http://research.cs.wisc.edu/htcondor/glow/
6. UW-HEP Condor User Info:The Open Science Grid (OSG). http://www.hep.wisc.edu/computing/condor.html#osg
7. XSEDE: The Open Science Grid User Guide. https://www.xsede.org/web/guest/OSG-User-Guide
8. M. Rynge. Open Science Grid: A new XSEDE service Provider. https://www.xsede.org/documents/234989/378230/XSEDE12_HTC_Tutorial_OSG.pdf
9. Welcome to OSG Connect. https://osgconnect.net/
10. Tsaregorodtsev, et al. DIRAC: a community grid solution. 2008 J. Phys.: Conf. Ser. 119 062048 doi:10.1088/1742-6596/119/6/062048
11. OSG: Install VOMS. https://twiki.opensciencegrid.org/bin/view/Documentation/Release3/InstallVoms
12. Fermilab: VOMRS Project. http://www.fnal.gov/docs/products/vomrs/