Facilitating Scientific Collaborations by Delegating Identity Management

Reducing Barriers & Roadmap for Incremental Implementation

Robert Cowles
BrightLight Information Security
1660 Gordon St., Suite 39
Redwood City, CA 94061
+1 650-544-9590
bob.cowles@gmail.com

Craig Jackson
Indiana University/CACR
2719 E. 10th Street, Suite 201
Bloomington, IN 47408
+1 812-856-5621
scjackso@indiana.edu

Von Welch
Indiana University/CACR
2719 E. 10th Street, Suite 201
Bloomington, IN 47408
+1 812-856-0363.
vwelch@iu.edu

ABSTRACT

DOE Labs are often presented with conflicting requirements for providing services to scientific collaboratories. An identity management model involving transitive trust is increasingly common. We show how existing policies allow for increased delegation of identity management within an acceptable risk management framework. Specific topics addressed include deemed exports, DOE orders, Inertia and Risk, Traceability, and Technology Limitations. Real life examples of an incremental approach to implementing transitive trust are presented.

Categories and Subject Descriptors

K.6.5 [Security and Protection]: Authentication.

General Terms

Management, Security, Verification

Keywords

Identity Management, Transitive Trust, Access Control, Identity, Cyber Security, Delegation, Risk Management

1. THE PROBLEM OF SUPPORTING SCIENTIFIC COLLABORATORIES

Computing centers, including those at laboratories operated on behalf of the US Department of Energy Office of Science (DOE Labs), often are faced with seemingly contradictory demands when trying to provide services to scientific collaboratories. Traditional processes and policies push DOE Labs to strongly identify, authenticate, and

authorize their user communities, but contemporary large distributed scientific collaboratories call for more distributed trust models. Collaboratories exhibit scale and dynamic changes in their user communities that place stress on the identity management (IdM) systems of the computing centers. Additionally, collaboratories may span multiple resource providers (RPs) and often advocate ease of use and management of their own user communities, including single sign-on across the RPs.

Transitive trust approaches are increasingly common and are often a desirable approach for IdM in a scientific collaboratory. In this approach, a virtual organization (VO) established by the scientific collaboratory manages its community, and RPs trust the VO to do so with little-to-no cognizance of the individual users. As such transitive trust approaches represent a relatively advanced version of delegated IdM.

The eXtreme Scale Identity Management (XSIM) project¹ at Indiana University, funded by the DOE's Next Generation Networks for Science program, has been studying how collaboratories and RPs work together to implement IdM. We have seen, for example with the work done at NERSC^[3] (discussed in more detail subsequently), that that it is possible to make incremental steps toward transitive trust and evaluate the risks along the way. IdM delegation can establish a simpler workflow with greatly reduced administrative overhead (for both RPs and users) and provide greater assurance by placing the responsibility with the entity best suited to address it.

In this paper, we discuss the work of the XSIM project in identifying and mitigating the barriers that traditional processes and policies may present to supporting scientific collaboratories in the context of IdM, some mitigations to those barriers, and how traditional computing centers as RPs can manage the delegation of IdM to a collaboratory that is acting as a VO.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CLHS'15, June 15-16 2015, Portland, OR, USA

Copyright is held by the owner/author(s). Publication rights licensed to ACM. ACM978-1-4503-3572-0/15/06...\$15.00

DOI: <http://dx.doi.org/10.1145/2752499.2752501>

¹ <http://cacr.iu.edu/collab-idm>

2. POTENTIAL BARRIERS TO IDM DELEGATION AND MITIGATIONS

IT architectures at DOE Labs developed at a time when the activities of scientific collaborations (e.g., instrument design and construction, code development, collaboration, data collection and analysis) could be handled at a single lab. In that environment, a researcher had a computing account with access to a multitude of lab-provided services. Hence, authentication, traceability, auditing, and accounting were the responsibility of that RP, and IdM requirements were specified in that context. Demand for computing resources then outgrew the capacity of the local collaboratory or even site-level RPs, so resources had to be obtained from multiple providers. Collaboratories distributed across multiple RPs bring the challenge of how to handle IdM across those RPs. The single-RP designs for IdM do not scale to multiple sites and requiring researchers to obtain credentials at multiple sites is not only an inconvenience but opens up a number of security problems. While a partial solution was implemented in the previous decade for grid computing, it is widely considered difficult to use since researchers often need to have authentication credentials on a variety of mobile devices. As mentioned previously, having the RPs delegate IdM to the collaboratory is an option, but the existing policies and procedures typically are seen as a barrier to that delegation. In this section we examine these barriers and discuss how, under close examination, mitigations exist for these barriers to VO-IdM delegation.

2.1. Regulation and Compliance

“Deemed Export” and “Unclassified Foreign Visits and Assignments” are often cited as the reasons for a DOE Lab to know the citizenship and birthplace of remote users of computer systems; therefore, they have a requirement to manage identities for their information systems.

2.1.1. Deemed Export

In general, export regulations are complex (for instance OFAC rules apply to economic and trade sanctions targeted against foreign countries and individuals) and sites will have an Export Control Officer who is cognizant in the applicability of the myriad rules. For sites running supercomputer-class systems, the “deemed export” regulations are often seen as a reason for collecting identity information.

Within the US Department of Commerce, the Bureau of Industry and Security has documented the policies surrounding “deemed export.” The definition of a deemed export is “the release of controlled technology to a foreign person.”^[4] Generally, an export license must be obtained prior to the release of controlled technology to a foreign person.

There are many exclusions to the requirement to obtain an export license. For instance, research involving publicly available information is excluded. Also, “fundamental

research”, defined as “basic and applied research in science and engineering where the resulting information is ordinarily published and shared broadly within the scientific community” is also excluded from the requirement for licensing.^[4]

In January 2009, the Office of National Security and Technology Transfer Controls issued an Advisory Opinion concerning the potential requirement for suppliers of grid or cloud computing to obtain export licenses: “Since the service of providing computational capacity through grid or cloud computing is not subject to the EAR, the service provider is not required to inquire about the nationality of the customer,”^[6] because the relevant EAR section regulates “activities in or by” specific countries and individual nationality is not directly relevant.

While there are certain exceptions to the exclusions (e.g., for users involved in missile activities), it appears worthwhile for an organization to have its legal department review the relevant regulations and referenced documents with an eye toward eliminating some of the administrative procedures around identity proofing often found in traditional processes.

2.1.2. Unclassified Foreign Visits and Assignments

The DOE Labs view the contract order DOE O142.3A^[7] titled “Unclassified Foreign Visits and Assignments” as an additional requirement for the local collection of identity information.

DOE O 142.3A spells out the procedures to be followed when foreign nationals come on-site at DOE Labs. Attachment 1 of the order is the Contractor Requirements Document (CRD) and contains the procedures to be followed by the DOE Labs that, while owned by DOE, are operated by a contractor. Attachment 2 of the order contains charts outlining a “graded approach” to the level of approvals required in various circumstances. Sites that perform some classified work are required to document visits or assignments of individuals from sensitive countries in the DOE’s Foreign Access Central Tracking System (FACTS). Visits or assignments of individuals from countries sponsoring terrorism require high-level advance approval.

Access to computing resources at DOE Labs is only directly addressed in the main body of the order. In the Duties and Responsibilities section of the order, it states:

Chief Information Officer (CIO). Drafts policy for issuance by the Secretary of Energy or Deputy Secretary regarding protective measures required for foreign national cyber security access approval, whether onsite or by remote access.^[7, pg. 13]

In the years since the order was issued in 2010, no DOE CIO has drafted any such policy. The CRD in Attachment 1 is silent on the issue of foreign national access to computing resources. Some DOE Labs have stated in their cyber security plans that there is no increased risk resulting

from foreign national access. [personal knowledge] Attachment 2 of the order is also silent on the issue of access to computing resources. The closest approach is the statement in a footnote:

These requirements apply only to those visits and assignments conducted in support of DOE missions and goals, or that otherwise involve access to DOE information or technologies.^[7]

Attachment 2, pp2,3]

Note that “DOE information” is a specific term relating to government owned data requiring Systems of Records notice to be filed in the Federal Register pursuant to the Privacy Act^[5, Attachment 1]; the thrust of this requirement is the identification of systems containing personally identifiable information (PII) and to ensure it is adequately protected. It only rarely is interpreted by the DOE Labs to include data from scientific research expected to be made public or data from work performed in cooperation with an industry partner that is proprietary (“Work For Others”). Similarly, the computing technologies made available to researchers at the DOE Labs are typically commercially available and are not technologies specific to DOE.

2.2. Historical Inertia and Introduction of Risk.

RP with a history of doing their own IdM may have policies, cultures, and reputations around information security, as well as a risk profile developed over many years of experience. Our recent interviews with computing centers supporting collaboratories have reinforced the importance of these factors as potential barriers to delegating IdM. However, we observe at least some of these RPs taking conservative steps to delegate IdM to VOs.

Introducing the possibility of a significant cultural change for DOE Labs, the requirements for a cyber security program were placed on a risk-based footing with the approval of order DOE 205.1B in May, 2011. The key paragraph in that order was the following:

The contractor must ensure all information systems operate within the processes defined and approved by the Federal Authorized Official, and that all systems maintain an acceptable level of risk pursuant to (1) the agreed upon risk profile defined by Site and Federal management, and (2) approved oversight and assurance systems.^[8]

Therefore, the cyber security program (including IdM) can be quite flexible so long as risks are documented, along with any mitigations, and the residual risks accepted both by the DOE Lab management and the DOE official monitoring the contract. Note that acceptance of the risks may require significant work to provide a comparison of the real risk and cost of a local IdM solution with the risk and cost associated with a transitive trust model.

2.3. Traceability

In a transitive trust environment the RP has no ability to contact individual VO members. However, it is natural that during the course of normal operation, unexpected or adverse events will happen at the RP servicing requests from the VO. RPs have long felt a requirement to ensure an individual causing an incident could either be contacted or that the RP has the ability to disable access. The RP and VO should have an agreement in place to handle these events. The OSG Traceability Project documented specific requirements around the issue of traceability and investigated the traceability of grid jobs submitted through GlideinWMS in HTCondor when there is no authentication of the end-user by the RP.^[1] The authors found that in all but a few improbable cases, sufficient information was available to trace a grid job back to the submitting user.

2.4. Technology Limitations

Many tools used by collaboratories assume a traditional model of user authentication and pose challenges when operated by an RP in a transitive trust scenario where IdM functions are delegated to the VO. Wikis, source code library systems, and line-mode interactive tools, to name a few, require identity information to function, but only allow for: authenticated individual access; access by an undifferentiated group to an individual user account; or anonymous access. Some technologies have been extended to allow access by a group to an individual user account while carrying information about the individual user to the RP, but they are relatively new and require some expertise and effort to deploy.

However, technology advances are quickly reducing barriers to IdM delegation. The computing grid services were implemented to serve, among other things, the LHC experiments. The development of improved virtualization technologies has resulted in various cloud computing models; these cloud computing models enable increased sharing within a group across physical and network boundaries and simultaneously increased isolation and *security* from other groups sharing the same physical resources.

3. XSIM ROADMAP FOR IDM DELEGATION

3.1. IdM Functions for Delegation

While transitive trust is a frequent model for implementation, IdM delegation from RP to VO is not an all-or-nothing affair. The XSIM project has identified a number of delegable functions that are either a direct part of or are enabled by IdM: Identification; Authentication; Authorization; Accounting; Auditing; User Support; and Incident Response. Delegating these functions can establish a simpler workflow with greatly reduced administrative overhead (for both RPs and users) and provide greater assurance by placing the responsibility with the entity best suited to address it. These benefits derived from delegation

require trust in the VO and the acceptance of risk by the RP. RPs can incrementally delegate these functions; examples of implementation variations exist as described in the following section.

3.2. Examples of IdM Delegation

3.2.1. XSEDE Science Gateways

In July 2011, the eXtreme Science and Engineering Discovery Environment (XSEDE) was announced as the follow-on to TeraGrid.² XSEDE continued the concept of a Science Gateway developed in TeraGrid.

A Science Gateway is a community-developed set of tools, applications, and data that are integrated via a portal or a suite of applications, usually in a graphical user interface, that is further customized to meet the needs of a specific community. Gateways enable entire communities of users associated with a common discipline to use national resources through a common interface that is configured for optimal use.^[10]

Nearly a decade ago as part of TeraGrid, the concept of “community accounts” in science gateways was proposed along as an example of a transitive trust model addressing the “four A’s” now understood to be a subset of the functions XSIM has identified: Authentication, Authorization, Auditing, and Accounting.^[12] There were a number of potential advantages outlined, but also some issues to overcome. A few years later, Basney *et. al.*, described the experience of implementing the community account model for TeraGrid gateways^[2].

The need to record and report resource utilization by different members of the collaboration was a key issue with the transitive trust model. At the time, the most accurate method of preserving utilization data was for the gateway to provide user identifying information to be recorded in the RP’s accounting records. More recently, virtualization and cloud computing models have forced VOs to take responsibility for resource tracking since the RPs have no knowledge of the individual users.

Initially, there was a strong requirement for RPs to be able to blacklist individual users to reduce the possibility of having to block an entire VO. In practice, the collaborations were willing to have the VO blocked in the rare circumstances it was necessary. Dropping this requirement enabled significant simplification of the code to implement the model.

The paper concluded that web portals worked best in situations where the users are only allowed to perform a limited set of actions and are not provided a general execution environment. Again, this was before significant use of cloud-based or virtual environments.

3.2.2. NERSC Scientific Gateways

The National Energy Research Scientific Computing Center (NERSC) is implementing a special type of login account called a “collaboration account” to facilitate collaboration and sharing of data. It is similar to the XSEDE science gateway described in the previous section in that it supports modern scientific computing that is increasingly performed by teams, meaning researchers need to be able to share access to compute and data resources in a secure, scalable manner. However, it represents a more limited delegation in that NERSC maintains control of authentication, auditing and accounting functionality, delegating only authorization to the VO.

Through a web portal operated by NERSC, termed a *scientific gateway*, a VO’s principal investigator adds users authorized to access the VO’s collaboration account.^[3] The portal allows NERSC to delegate the authorization of users for access to the collaboration account to the VO. The VO determines user privileges to access and manipulate files and jobs run by other members of their VO while NERSC still controls the authentication of the user performing the actions and the subsequent auditing and accounting.

3.2.3. ATLAS LHC Experiment Use of PanDA Without Identity Change

The Production and Distributed Analysis (PanDA) system is a set of services around distributed job submission and execution in a computing grid. It was originally developed to support the job execution requirements of US ATLAS and was then adopted by the entire ATLAS experiment in 2008.^[11] PanDA employs a pilot job submission framework and once an actual job is selected for execution, it can optionally change the identity under which the job is run.^[9] Although the capability exists, US ATLAS does not exercise the identity changing option; all ATLAS jobs run at sites in the US (including DOE Labs) execute under the identity of the pilot job submitted from the ATLAS VO.

This design represents a transitive trust relationship with complete delegation of IdM (Identification; Authentication; Authorization; Accounting; Auditing; User Support; and Incident Response) from the RP to the VO with policies to ensure the ability to trace from jobs back to users. Sites that wish to contact a user directly concerning a job must first contact the VO to obtain the user’s identity and contact information.

4. CONCLUSION

Computing centers, including those operated by DOE Labs, often are faced with seemingly contradictory demands when trying to provide services to large, virtual scientific collaborations. Traditional practices and policies may appear to be in conflict with the collaboration-based IdM; however, the delegation of IdM responsibilities to VOs is increasingly being leveraged to aid usability and streamline support of dynamic user communities.

² <https://en.wikipedia.org/wiki/TeraGrid>

In this paper we've presented how existing policies allow the delegation of IdM functions to laboratories within the context of acceptable risk management. Based on the work of the DOE-funded XSIM project, we have suggested strategies that allow for the incremental increase of trust and delegation of IdM functionality. Examples from XSEDE, NERSC, and ATLAS were provided.

5. ACKNOWLEDGEMENTS

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

6. REFERENCES

- [1] Altunay, M. (2013) *Traceability Requirements for end user jobs without certificates*. Retrieved from OSG Document Database February 27, 2015. DOI=<http://osg-docdb.opensciencegrid.org/cgi-bin/ShowDocument?docid=1149>.
- [2] Basney, J., Welch, V., and Wilkins-Diehr, N. (2010) TeraGrid Science Gateway AAAA Model: Implementation and Lessons Learned. TeraGrid Conference. DOI=<http://www.vonwelch.com/pubs/tg2010>.
- [3] Cowles, R., Jackson, C., and Welch, V. (2014) A Model for Identity Management in Future Scientific Collaboratories. *International Symposium on Grids and Clouds*. DOI=http://pos.sissa.it/archive/conferences/210/026/ISGC2014_026.pdf.
- [4] Deemed Exports. (n. d.) *U. S. Department of Commerce, Bureau of Industry and Security*. DOI=<http://www.bis.doc.gov/index.php/policy-guidance/deemed-exports>.
- [5] Department of Energy Privacy Program. (2009) *U. S. Department of Energy*. DOI=https://emcbc.doe.gov/ContentOffice/doe_o_206.1.pdf.
- [6] Doc: Application of EAR to Grid and Cloud Computing Services. (2009) *Fismapedia*. DOI=http://fismapedia.org/index.php?title=Doc:Application_of_EAR_to_Grid_and_Cloud_Computing_Services.
- [7] DOE O 142.3A. Unclassified Visits and Assignments Program. (2010) *U. S. Department of Energy, Office of Information Resources*. DOI=<https://www.directives.doe.gov/directives-documents/100-series/0142.3-BOrder-a>.
- [8] DOE O 205.1B Chg 2, Department of Energy Cyber Security Program (2011) *U. S. Department of Energy, Office of Information Resources*. DOI=<https://www.directives.doe.gov/directives-documents/200-series/0205.1-BOrder-b-chg2>.
- [9] PanDA Security (2009) *Twiki PanDAWeb Atlas Distributed Computing PanDA*. DOI=<https://twiki.cern.ch/twiki/bin/view/PanDA/PandaSecurity>.
- [10] Science Gateways. (n.d.) XSEDE, *Extreme Science and Engineering Discovery Environment*. DOI=<https://www.xsede.org/gateways-overview>.
- [11] The PanDA Production and Distributed Analysis System. (2014) *Twiki PanDAWeb Atlas Distributed Computing PanDA*. DOI=<https://twiki.cern.ch/twiki/bin/view/PanDA/PanDA>.
- [12] Welch, V., Barlow, J., Marcusiu, D., and Wilkins-Diehr, N. (2006) A AAAA Model to Support Science Gateways with Community Accounts. *Concurrency and Computation: Practice and Experience*. DOI=<http://www.vonwelch.com/pubs/SGW>.