Digital Object Registry

GENI, Spiral 2

Milestone D: Collaborate on Security Design for Spiral 2.

1. Introduction

In our Clearinghouse Security Requirements report (CNRI, 2009), submitted to the GPO during Spiral 1 on July 01, 2009, we discussed some of the security requirements for the GENI framework and highlighted potential solutions based on our implementation of the GENI Federated Clearinghouse (GFC). During the current spiral, we analyzed the ways in which the Handle System (Sun, Lannom, & Boesch, 2003) could be adapted to manage GENI-wide user identity and related trust components. We addressed this from the point of view of GENI adopting Shibboleth for user identity management, specifically its InCommon Federation profile. This assumption was validated by our GENI system engineer, Vicraj Thomas. We summarize of our analysis in this report and suggest future actions.

As discussed in our Spiral 1 security requirements report, security in any network-based system is multi-faceted, with challenges ranging from ensuring message integrity and confidentiality, to dealing with DoS attacks, and to keeping accurate traffic logs for auditing purposes. To be successful as an experimental and prototyping infrastructure, the GENI environment must be trusted and reliable. Experiments must be repeatable, so the state conditions must be well known and must exclude any unwanted interference, both malevolent and unintentional. Among the many potential topics, we focus here on the trust model, including its flexibility in meeting GENI requirements, distributed authentication, and privilege revocations.

2. Security Systems Overview

The Shibboleth System (Shibboleth, 2010) is standards-based, open source software for managing single sign-on across organizations over the web. It is middleware usually deployed between an identity provider and one or more service providers. Identity providers authenticate users requesting one or more protected resources or services from a service provider. Shibboleth is an established web-based system that passes necessary information from identity providers to service providers in order for those service providers to make informed decisions on whether or not to authorize user access to protected resources.

The functionality provided by Shibboleth overlaps or incorporates similar systems with similar capabilities and functionalities, e.g., OpenID, OASIS SAML, Microsoft's InfoCard, etc. It is, for example, SAML compliant, but SAML can be, and is, used outside of Shibboleth. Other similar systems differ primarily in the way in which messages are encoded and transmitted. Shibboleth's primary strength, in our view, is not in its technology per se but in the communities of practice that have embraced the technology and formed federations, where each of the federations is based on real-world trust. Shared community goals can be encoded in policy, and effected through an interoperable service provider pool. The primary example of this is the InCommon Federation, which serves the U.S. Higher Education community and its partners. Federal agencies are also working with

the InCommon Federation to provide the policy and technical base, which would allow federation members to use their campus-issued credentials for accessing government web-based services.

The Handle System is a standards-based, open source system primarily used for managing the current state data of digital objects, where the concept and implementation of digital object is sufficiently broad and abstract that a digital object may represent any of a variety of artifacts ranging from any unit of information stored in digital form on a network system, e.g., files, videos, etc., to computer systems including mobile devices, as well as users and processes that access the digital information or computer systems. The Handle System is a scalable, distributed system that allows assigning persistent, unique identifiers, aka handles, to digital objects, which are globally resolvable. Usually, handles resolve to the current state data of the digital objects they identify. In the case of principals, the current state data may include the public key, role, and related credential information.

The granularity of administration, i.e., the modification of the current state data of an object as reflected in its handle record, is at the individual handle level. That is, if required, each handle may have its own administrator. This facilitates its use in identity management. A typical use of handles for identity management and access control in many systems developed by CNRI, including the GENI Federated Clearinghouse, can be summarized as follows:

1. Each user of the system is assigned a handle, which resolves to the public key of the user identified by that handle,
2. Users trying to access a service, where authentication is managed through the Handle System, present their handle to the service,
3. The given service, using the handle client library, resolves the user handle, retrieves the public key, perhaps one of many depending on the roles and services involved,, and issues a PKI challenge-response to the user to verify that the user has the corresponding private key.

Additionally, in the GENI Federated Clearinghouse implementation, the clearinghouse verifies that the user handles presented indeed belong to an identifier range that has previously been assigned for use in GENI. The general approach is that authentication policies may be attached to identifiers, which at run-time, may be evaluated and the results of the evaluation used to enforce the policies. Any changes to either the policies or the credentials will be immediately effective, since the whole authentication process is based on the concept of resolving to current state data as required. Additional details of the authentication approach herein described are discussed in our Spiral 1 security requirements report.

3. Security Analysis

Over the past few months, we have investigated the potential synergies among Shibboleth, GENI, and the Handle System, including meeting on this topic with Steve Schwab, from the GENI security-working group, at CNRI on May 20th. Our conclusion, detailed below, is that while Shibboleth has much to recommend it is not an ideal fit with GENI. It is focused on single sign-on and shared access to web resources.

Shibboleth can be described as a system that provides systems and services with information required for trusting users. However, since its inception, the technology has evolved around two assumptions: (1) the set of systems and services is the HTTP-centric web, and (2) the user communities have pre-defined and discrete roles and access control

structures. The InCommon Federation, as a prime example, has developed around the typical higher education roles, e.g., faculty, staff, students, etc., where such roles have discrete, and in some cases hierarchical, privileges. Further, its use is primarily for allowing access for services such as library systems and institutional repositories. While the GENI community, in its current state, is heavily weighted to universities and higher education, we believe its objectives are not and should not be confined to those communities. It is not out of the question that GENI will be used and further developed, for example, by industry and the military. Even now, the roles of GENI users cannot be mapped easily to faculty, staff, and student categories. Different sub-systems within GENI, e.g., measurement data, archive system, and federated clearinghouses, may have different data security requirements as mandated by both the users and the policy makers, in which a pre-defined role template may not be readily applicable.

Beyond the potentially limiting web-centric nature of Shibboleth and InCommon is the issue of Shibboleth being based on sharing user attributes with service providers where the attributes that are shared are pre-established for each of the service providers, thereby making ad hoc sharing of user attributes impossible. Systems that rely on the sharing of arbitrary user attributes for authenticating users will be a bad fit with Shibboleth. For various GENI sub-systems, e.g., control framework aggregate managers, archive systems, etc., which may be implemented and hosted by multiple organizations, abiding by this policy requirement may prove to be costly. The legal framework of InCommon, which has been now been accepted by a large number of organizations, would likely increase the difficulty of a significant policy change.

We believe we could ameliorate some of the shortcomings of Shibboleth and InCommon, in their applicability to the GENI framework, through integration with the Handle System.

Both limitations described above, namely the web-centric nature of Shibboleth and the fixed approach to user attribute sharing could potentially be addressed by integrating the Handle System with Shibboleth. Since Shibboleth is a middleware system that allows using different authentication and attribute storage infrastructure, the Handle System could be plugged-in to manage both those aspects. Each GENI user could be assigned a handle with each of those handles resolvable to corresponding attributes, including public keys and other InCommon required attributes. The Shibboleth framework for InCommon could be implemented using the Handle System. Ad hoc user attribute sharing is made effective without requiring substantial changes to the existing Shibboleth/InCommon framework. Further, if needed, the Handle System, which contains its own PKI, may be used as a complete authentication system for some services. Effectively, given the Handle System integration, trust interactions could happen at three levels:

1. Access to service providers enabled through InCommon functions with no change, except for those that require a specific authentication method.
2. GENI sub-systems may employ user attributes externalized through the Handle System for authentication, authorization, or other reasons.
3. GENI sub-systems may use the authentication mechanism implemented in the Handle System.

Note that the interactions defined in Points 2 and 3 above are based on a non-web network environment.

A possible option is for CNRI to join the InCommon Federation and integrate the Handle System as discussed. This would allow GENI members the benefits of both Shibboleth and the Handle System via the user attributes that it externalizes. Another possibility is to

aggregate user identity and user attributes outside of the context of Shibboleth and InCommon. This would enable non-web centric access to user attributes. Note that the GENI Federated Clearinghouse already has aggregated ProtoGENI user identities and attributes.

We would be happy to discuss any of these options with the GPO. Note, however, that while we believe our technical analysis is accurate, we haven't discussed or analyzed any legal implications. The notion of joining InCommon and extending its functions through integration with the Handle System would have to be completely reviewed from a legal perspective.

4. References

Sun, S., Lannom, L., & Boesch, B. (2003). *RFC 3650.*

CNRI. (2009). *Digital Object Registry*. Digital Object Registry: Clearinghouse Security Requirements. http://groups.geni.net/geni/attachment/wiki/DigitalObjectRegistry/ClearinghouseSecurityReqmnts.pdf

Shibboleth. (2010). Shibboleth, A Project of the Internet2 Middleware Initiative. Shibboleth. http://shibboleth.internet2.edu/