CTSC
CENTER FOR TRUSTWORTHY
SCIENTIFIC CYBERINFRASTRUCTURE

InCommon Membership in eduGAIN: the LIGO Perspective

LIGO-G1300767-v1
July 16, 2013
For Public Distribution

Jim Basney and Scott Koranda

About CTSC

The mission of the Center for Trustworthy Scientific Cyberinfrastructure (CTSC, trustedci.org) is to improve the cybersecurity of NSF science and engineering projects, while allowing those projects to focus on their science endeavors. This mission is accomplished through one-on-one engagements with projects to solve their specific problems; broad education, outreach and training to raise the practice-of-security across the community; and looking for opportunities for improvement to bring in research to raise the state-of-practice.

Acknowledgements

The authors would like to thank Stuart Anderson, Warren Anderson, Steven Carmody, Scott Cantor, John Krienke, Mikael Linden, Maria Laura Mantovani, Chris Phillips, Hiroyuki Sato, Tom Scavo, Rhys Smith, Ian Young and all members of the InCommon TAC Subcommittee on Interfederation, and the LIGO Identity and Access Management project team.

This document is a product of the Center for Trustworthy Scientific Cyberinfrastructure (CTSC). CTSC is supported by the National Science Foundation under Grant Number OCI-1234408. For more information about the Center for Trustworthy Scientific Cyberinfrastructure please visit: `http://trustedci.org/`. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

Using & Citing this Work

Cite this work using the following information: J. Basney and S. Koranda, “InCommon Membership in eduGAIN: the LIGO Perspective,” Center for Trustworthy Scientific Cyberinfrastructure, `trustedci.org`, May 2013.

This work is available on the web at the following URL:
[Insert trustedci.org URL when/if applicable.]

1 Introduction

The Laser Interferometer Gravitational-Wave Observatory (LIGO) is a large research project funded by the National Science Foundation. LIGO seeks to make the first direct detection of gravitational waves, use them to explore the fundamental physics of gravity, and develop the emerging field of gravitational wave science as a tool of astronomical discovery. Through a cooperative agreement with NSF, the California Institute of Technology (Caltech) and the Massachusetts Institute of Technology (MIT) jointly operate the LIGO Laboratory and its two observatories, one in Hanford, WA and one in Livingston, LA. The LIGO Scientific Collaboration is the international group of researchers carrying out the science of the LIGO Observatories as well as that of the GEO600 detector in Hannover, Germany. Today LIGO is a worldwide collaboration with more than 1000 members from across five continents.

Over the past few years LIGO invested significantly to develop a SAML-based single sign-on infrastructure. LIGO operates a Shibboleth Identity Provider (IdP) and provisions a LIGO electronic identity (branded as an “albert.einstein@LIGO.ORG” identity) for each collaboration member. The collaboration operates more than 50 Shibboleth service providers (SPs) that host a wide spectrum of services including wikis, document catalogs, event databases, and data investigation tools.

LIGO has planned from the beginning to leverage federated identities to address two primary use cases. First, although LIGO provisions an electronic identity for each collaboration member, many members have a pre-existing federated identity that could in principle be used to access LIGO SPs. By reducing the number and scope of provisioned LIGO identities the collaboration can decrease the burden of having to operate an IdP and the associated help desk services needed to assist users in managing a LIGO electronic identity. Second, the full impact of LIGO science can only be realized with close collaboration between LIGO scientists and astronomers and astrophysicists from other projects. Federated identity helps streamline collaboration between LIGO scientists and other researchers by enabling easier access to resources without the need for provisioning LIGO identities to external collaborators.

To facilitate leveraging federated identity and begin pursuing interoperability LIGO has joined the InCommon identity federation in the United States (US). Through the InCommon identity federation LIGO has enabled federated access to its resources for a large number of researchers in the US. Because LIGO is an international collaboration, however, and the pool of possible external collaborators is global, far more work remains to federate with institutions and projects from around the world.

The number of identity federations of interest to LIGO is large. Today the federations that either intersect directly with LIGO membership or with existing and possible LIGO collaborators includes:

- Australian Access Federation (AAF)
- FederationCAFe (Brazil)
- Canadian Access Federation (CAF)
- CERNET Authentication and Resource Sharing Infrastructure (China)
- DFN-AAI (Germany)
- Fédération Éducation-Recherche (France)
- eduID.hu (Hungary)
- INFLIBNET Access Management Federation (India)

- IDEM (Italy)
- GakuNin (Japan)
- SURFnet (Netherlands)
- Servidor de Identidad de RedIRIS (Spain)
- UK Access Management Federation for Education and Research

Of special interest are collaborators from other interferometric gravitational wave experiments and organizations including the European Gravitational Observatory (EGO), responsible for the computing and networking for the Virgo (French and Italian) interferometer experiment. At this time, to facilitate research, LIGO provisions a LIGO electronic identity for Virgo members who request access to LIGO resources. LIGO would prefer to leverage federated identity instead, since provisioning and managing identities for Virgo members is burdensome.

To realize the promise of federated identity to enable easier collaboration, LIGO is faced with the daunting task of pursuing interoperability with each of the identity federations separately, most likely by having to directly join each federation.

A better path to international federation for LIGO would be to leverage its existing membership in InCommon. Ideally, having already joined InCommon, LIGO would automatically interoperate with the federations listed above, as well as other identity federations throughout the world, through inter-federation agreements, policies, and practices. A vetted research and scholarship organization such as LIGO, after joining InCommon, should find that without further effort its IdP and SPs interoperate with any IdPs and SPs in any of the higher education and research SAML federations worldwide.

No functional inter-federation infrastructure between InCommon and other federations, however, exists today. The eduGAIN service in Europe originally intended to enable federation between the GÉANT (GN3) partners' federations in Europe but has more recently expanded to include federations from around the globe including Canada and Brazil and pending memberships from Japan and Chile. Should the InCommon Federation in the US join eduGAIN, LIGO through its membership in InCommon would gain access to a large number of identity federations, IdPs, and SPs that would greatly enhance collaboration between LIGO scientists and other astronomy and astrophysics researchers throughout the world.

This document discusses InCommon membership in eduGAIN from the LIGO perspective, including the current status of InCommon efforts to join eduGAIN, and highlights issues that LIGO should continue to monitor as that effort evolves. Also included is a discussion of the issues around the release of attributes by eduGAIN IdPs in the European Union (EU) to LIGO SPs.

2 eduGAIN and InCommon: the LIGO perspective

The eduGAIN service was "developed within the GÉANT project. eduGAIN interconnects identity federations around the world, simplifying access to content, services and resources for the global research and education community. eduGAIN enables the trustworthy exchange of information related to identity, authentication and authorisation (AAI) by coordinating elements of the federations' technical infrastructure and providing a policy framework that controls this information exchange."[1]

[1] http://www.geant.net/service/eduGAIN/about_edugain/Pages/AbouteduGAIN.aspx

eduGAIN is not a legal entity that can sign contracts. Rather an identity federation like InCommon declares that it adheres to the eduGAIN policy, signs the eduGAIN Policy Declaration,[2] and then applies for membership. Version 2.0 of the eduGAIN Policy Framework[3] was ratified on June 26, 2013.

Note that membership in eduGAIN does not automatically enable the interoperability of all IdPs and SPs from the member federation. Rather each member federation is free to set its own policy and technical requirements for which IdPs and SPs it exposes to other member federations by inclusion of their SAML metadata in the eduGAIN metadata aggregate feed [4].

From the LIGO perspective the primary questions regarding InCommon's membership in eduGAIN are:

1. Does the metadata for IdPs and SPs with which LIGO would like to interoperate exist or will it exist in the eduGAIN metadata aggregate feed?

2. Is the eduGAIN Policy Framework acceptable to InCommon and can InCommon sign the eduGAIN Policy Declaration?

3. Does the eduGAIN metadata aggregate feed meet the trustworthiness standards for LIGO at this time or in the future?

4. Will InCommon adopt a policy and operational support so that the metadata for LIGO IdPs and SPs can be made available as part of the eduGAIN metadata?

5. Will IdPs from eduGAIN member federations, particularly in the EU, be willing and legally able to release attributes about authenticated users to LIGO SPs?

In anwser of the first question, a cursory examination of the current eduGAIN metadata aggregate shows that a large fraction of the Virgo collaboration would already be represented by IdPs available in the aggregate, including IdPs operated by the Italian Academic and Research Network GARR[5] that represent and include users from the various Istituto Nazionale di Fisica Nucleare (INFN) organizations, as well as researchers at the Research Institute for Particle and Nuclear Physics of the Hungarian Academy of Sciences and the Nikhef National Institute for Nuclear Physics and High Energy Physics in the Netherlands. With the pending inclusion of the Japanese federation GakuNin and the UK Access Federation, there is no question that a large number of the IdPs and SPs of most interest to LIGO either are already or will be available in the eduGAIN metadata aggregate.

3 Status of InCommon joining eduGAIN

As part of the CTSC and LIGO engagement Jim Basney helped launch and then lead an InCommon Technical Advisory Committee (TAC) Interfederation Subcommittee.[6] The mission of the subcommittee was "to promote and pursue interfederation between the InCommon Federation and other SAML federations via a community-based process. The subcommittee makes recommendations to the InCommon Technical Advisory Committee, and members of the subcommittee interact with members and operators

[2] http://www.geant.net/service/eduGAIN/resources/Pages/home.aspx

[3] http://www.geant.net/service/eduGAIN/resources/Pages/home.aspx

[4] http://mds.edugain.org

[5] http://www.garr.it/b/eng

[6] https://spaces.internet2.edu/display/incinterfed/Interfederation+TAC+Subgroup

of other SAML federations to draft agreements and common practices." Scott Koranda joined the subcommittee to represent LIGO's interests. InCommon operations was represented on the subcomittee by John Krienke and Tom Scavo.

The subcommittee completed its work in June 2013 and sent to the TAC a number of recommendations. [7]. For convenience we reproduce here the recommendations related to LIGO's interest in having InCommon participate in eduGAIN:

1. **Establish international interfederation agreements with eduGAIN and UK federation**. Acknowledging that these agreements are not the totality of interfederation but are a concrete step forward.

   (a) **InCommon becoming an eduGAIN member**. Work with InCommon Operations to achieve InCommon membership in eduGAIN. Follow the InCommon governance process to obtain InCommon Steering approval for eduGAIN membership. Sign eduGAIN declaration. Work with Canadian Access Federation on eduGAIN pilot projects.

   (b) **InCommon interfederating with UK federation**. Follow InCommon governance process to sign bilateral agreement with UK federation. UK federation has an agreement ready for InCommon to sign.

2. **Document trust practices and policies for entity registration and publishing**.

   (a) **Metadata exchange**: Perform a due-diligence review of InCommon policies related to metadata exchange with non-InCommon members. Determine policy for which eduGAIN entities would be provided in a metadata aggregate to InCommon members, and which InCommon entities would be provided to eduGAIN (potentially including an opt-in or opt-out process and potentially starting with R&S entities). Communicate with InCommon membership regarding trust issues associated with eduGAIN participation. Determine level of trust required for entities included in InCommon's "import" interfederation metadata aggregate(s). Determine if InCommon should provide "untrusted" interfederation metadata to its members versus only entities that meet baseline trustworthy practice, to help scale the trust. Determine opt-in/opt-out process for InCommon entity inclusion in "export" aggregate(s).

   (b) **Registration practice**: Document InCommon registration practices to a level similar to UK Federation Technical Specifications. This documentation will be useful as input to eduGAIN. REFEDS may develop a template for registration practice statements, and if/when that happens, InCommon should conform to the template. Develop a common InCommon-UK registration practice standard that could be floated for wider adoption. Topics include private key handling, upload of metadata from org to fed operator, key sizes, organizational validation, etc. This can set a criteria for assessing eduGAIN members and other interfederation partners.

3. **Develop and adopt a US-EU Code of Conduct to address privacy and attribute release**. There is a DRAFT of an extension to the CoC that would allow EU-based IDPs to release attributes to SPs that are InCommon members if those SPs were to assert compliance. This draft should be forwarded to the InCommon lawyers for review.

[7] https://spaces.internet2.edu/display/incinterfed/June+2013+Recommendations+to+TAC

4. **Implement improvements and new capabilities for metadata management, publication, aggregation, tagging (i.e., technical work)**. Continue to rely on LIGO as a driver for technical pilot projects, and welcome additional driving use cases.

    (a) **InCommon adding <mdrpi:PublicationInfo> and <mdrpi:RegistrationInfo> elements in metadata**. Addition of <mdrpi:PublicationInfo> to InCommon metadata is now planned. Assuming that goes well, adding <mdrpi:RegistrationInfo> to each entity in InCommon metadata can happen later. This will help with metadata aggregation by clearly identifying the registrationAuthority and publisher for each entity. When an aggregator publishes metadata, the registrationAuthority won't change but the publisher will identify the aggregator.

    (b) **InCommon providing one or more production "import" metadata aggregate(s) for consumption by InCommon members**.

    (c) **InCommon providing one or more production "export" metadata aggregate(s) for consumption by external partners (UK, eduGAIN, etc.)**.

    (d) **InCommon support for additional entity tags**. As REFEDS and other groups develop standard entity tags, indicating (for example) whether an IdP should be included in discovery interfaces or indicating an SP's privacy policy, InCommon should provide the ability for InCommon entities to self-assert these tags. This can also include a tag indicating acceptance of the InCommon membership agreement.

As noted above the subcommittee has formally recommended to TAC that InCommon join eduGAIN. The recommendation came in substantial part because of the changes to the eduGAIN Policy Framework made for version 2.

Open questions do remain that need to be addressed as InCommon considers joining eduGAIN. For example, the new eduGAIN policy does not require entities to consent to metadata exchange and so InCommon is obligated to make a due diligence review of what statements InCommon has made to participants around sharing metadata.

Similarly, discussions by the subcommittee noted that the eduGAIN Policy Framework and the eduGAIN operational details may be necessary but not yet sufficient to accomplish a truly trustworthy exchange of metadata. Particular concerns were raised over registration practices by federations for entities and the documentation of those registration practices. The core question is whether an end entity like a LIGO SP can trust that an IdP entity in the eduGAIN metadata is truly representative of the organization it purports to represent. For example, can a LIGO SP trust that an IdP in the eduGAIN metadata purporting to be operated by Cardiff University is really operated by *the* Cardiff University (this is just an illustrative example and in no way reflects any concern about the trustworthiness of the Cardiff University IdP metadata). As noted above the subcommittee has made particular recommendations about how InCommon can addresses some of these questions both at a technical and operational level.

The subcommittee also learned that it is not unlikely that as InCommon brings up issues and concerns with eduGAIN about particular policy and operational details those issues can be addressed and new versions of the eduGAIN policy framework issued "as long as such a new version did not introduce any mandatory constraints which existing members could not meet (i.e., as long as no member was in danger of being excluded from eduGAIN by a proposed change) then the process to pass such a revised profile towards the end of the year sounds fairly straightforward." [8].

[8] PrivatecommunicationtoIanYoungasamemberofthesubcommitteefromeduGAINrepresentatives

Based on the subcommittee's work and further communications between InCommon and eduGAIN representatives, John Krienke, Chief Operating Officer, sees a path forward for InCommon to join eduGAIN and although he cannot promise a particular time frame it would not be unlikely that InCommon could join eduGAIN in the first half of 2014.

Taken together these developments positively answer questions 2, 3, and 4 above and indicate that LIGO will be able to benefit from InCommon's participation in eduGAIN.

4 Attribute release from EU eduGAIN IdPs

In addition to a trustworthy exchange of metadata, LIGO SPs will benefit most when IdPs in the EU release user attributes containing a minimal level of personally identifiable information (PII), such as name and email address, during a SAML authentication flow. Today the exchange of attributes from an EU IdP to an EU SP in the higher education and research context is governed by the GÉANT Data Protection Code of Conduct (CoC). [9]. The CoC is based on particular EU law and does not and can not by itself apply to the release of attributes to entities in the US.

Instead the community is proposing that the CoC be combined with a "standard contractual clause", based on an amendment to an existing EU law, [10] and use the combined CoC plus the clause as a vehicle that will allow EU IdPs to legally release attributes with minimal PII to SPs in the US including LIGO SPs. A high level overview was given by Mikael Linden at at recent REFEDs meeting. [11] The combined current CoC and the inclusion of a standard contractual clause would constitute a "non-EU CoC" that could be used with VOs like LIGO. A lawyer has recently begun drafting such a non-EU CoC and the related conversations have used LIGO as an example use case. A draft is expected at the end of the summer 2013 [12] Steve Carmody is representing InCommon in the process.

At this time it is envisioned that a single signed non-EU CoC will be sufficient for LIGO to receive attribute assertions from IdPs in the EU, rather than having to sign agreements with each federation. The goal is to "make the federations disappear" so that LIGO need only sign once. In fact a signature might not be necessary and instead the LIGO SPs will be asked to place a specific Entity Category attribute in the SAML metadata to indicate commitment to the non-EU CoC and then refer to the non-EU CoC in the published LIGO privacy policy.

Note that some LIGO SPs already operate within the EU (bugs.ligo.org for example is operated by the LIGO group at the Albert Einstein Institute in Hannover, Germany). It is an open question if those SPs need to be covered by the EU CoC or the non-EU CoC. Most likely since it will be LIGO as an organization that commits to the CoC and not an individual SP operator the SPs would be covered by LIGO committing to a non-EU CoC.

The work outlined above clearly indicates that the answer to question 5 above is that relatively soon IdPs in the EU will be able to legally release some attributes with PII to LIGO SPs in support of the LIGO research mission.

[9] https://refeds.terena.org/images/1/18/GEANT_DP_CoC_ver1.0.pdf

[10] http://eur-lex.europa.eu/LexUriServ/LexUriServ.do?uri=OJ:L:2004:385:0074:0084:EN:PDF

[11] See slide 5 at https://refeds.org/meetings/june13/slides/20130602-ML.pdf

[12] Private communication from Mikael Linden to Scott Koranda.

5 Conclusion and recommendation

The work of the InCommon subcommittee on Interfederation as led by Jim Basney from CTSC clearly helped facilitate and expedite the consideration by InCommon to join eduGAIN. We conclude that it is quite likely that InCommon will be a member of eduGAIN by the end of 2014 and that LIGO through its membership in InCommon will have access to IdP entities in the eduGAIN metadata and will have its own metadata consumed by those same IdPs, leading to much greater internal interfederation and interoperability for LIGO services supporting collaboration between LIGO and other astronomy and astrophysics projects.

We recommend that LIGO continue to provide resources to allow it to be used as a primary use case and driver for further interfederation policy work and as a tester for anticipated operational changes at InCommon in support of the anticipate membership in eduGAIN.