PROCEEDINGS OF SCIENCE

Building Security and Trust in Inter-Federation

Hannah Short[*†]
CERN
E-mail: hannah.short@cern.ch

Romain Wartel
CERN
E-mail: romain.wartel@cern.ch

The expanding network of Higher Education and Research facilities through inter-federation, whilst generally perceived as extremely valuable for collaboration and online security at large, exposes inviting new possibilities for malicious attacks[1]. A single compromised account may provide an entry point to this global network of resources linking thousands of organisations. How can we, the community, coordinate a response spanning countries and continents? How can trust be built between the organisations, and between the people, in our communities?

REFEDS (the Research and Education FEDerations group)[2], in conjunction with the European Commission funded AARC Project (Authentication and Authorisation for Research and Collaboration)[3], is spearheading the Security Incident Response Trust Framework for Federated Identity (Sirtfi)[4] as a method for mitigating the impact of security incidents to federations. This framework provides a list of statements which an organisation must self-assert to be deemed Sirtfi compliant, spanning best practices in operational security to traceability.

Organic global trust groups already provide a platform for informal alliances within academia, research and industry, however there is a need for heightened transparency, inclusivity and structure to facilitate this process. The lack of centralised governance within this space, in contrast to individual organisations or even national federations, calls for a standard procedure that can be adopted by all participants. What role will individuals play as this network grows in magnitude? This paper, a summary of the presentation given at the International Symposium on Grids and Clouds 2016, explores the practicalities of closing the loop on federated security. A two fold approach is presented, building trust between organisations and between the individuals therein.

International Symposium on Grids and Clouds 2016
13-18 March 2016
Academia Sinica, Taipei, Taiwan

[*]Speaker.

[†]The research leading to these results has received funding from the European Community's Horizon2020 Programme under Grant Agreement No. 653965, AARC (Authentication and Authorisation for Research and Collaboration). This work is based on discussions within multiple communities, chiefly the REFEDS Sirtfi Working Group and AARC NA3.2 Workpackage.

PoS(ISGC 2016)030

1. The Federated Landscape

Identity Federation has become a key enabler for international collaboration within research and education. The ability to use one's home organisational account to access shared services streamlines account maintenance at participating entities and provides a unified experience for end users. Over the past five years, Federated Identity Management (FIM) has established itself as a standard service offered by hundreds of universities and research organisations; over 60 national federations are operational, with more set to follow[2].

A federation can be described succinctly as a group of service providers and identity providers that agree to interoperate under a shared policy set, such that a user from a participating identity provider is able to access any participating service. Federations tend to be geographically bound with policies reflecting the national law and funding structures. To truly enable global research and collaboration FIM has been extended internationally via eduGAIN, the inter-federation service[5], which links these national federations.

Whilst this technology is clearly beneficial for research, it does introduce a new potential strategy for online criminals, commonly termed a vector of attack. Federations are built upon a chain of assurance profiles spanning services, organisations and national federations - the threat of federated incidents has the potential to break this trust model and obstruct the successful adoption of global federated access.

Unlike many single nation federations, the inter-federated world has no centralised help desk[6]. In the event of an incident spanning multiple federation participants, it will be incumbent upon the affected entities themselves to collaborate towards recovery. Currently there is no visibility into each participant's security capability and no guarantee that they will willingly collaborate.

2. Common Vectors of Attack

There has never been a time when cybercrime has offered such high profit at such low risk. Organised crime units propose malware as a service, using robust frameworks that have been developed over many years. At first glance, the data stored by research institutions is typically public and there is relatively little money to be extorted from such organisations. Recent intrusions do, however, demonstrate several areas for concern. Organisations conducting research and innovation may be targeted for their forward looking technologies, tender and purchase strategy, vendor information and technology, employees' personal details, computing accounts or resources, all of which can then be sold on dedicated underground markets. Their financial systems are also commonly targeted to send money to rogue bank accounts by exploiting salary payments or contractor invoices.

The threat of cybercrime to research and education should not be underestimated. Data, no matter the content, has a price on the criminal marketplace.

3. Federated Incident Response

The lifecycle of a security incident can be generally described as follows; preparation, identification, containment, eradication, recovery and lessons learnt. This is a cycle in which lessons learnt from one incident should be incorporated into preparation for future incidents. The traditional lifecycle is further complicated when applied to a federated community. Is an organisation able to know whether adequate preparation has been made at partner organisations? Who can be contacted to help identify the scale of an incident? The community must be able to identify the penetration of a compromised account through the combined logs of participants. Are such logs available and are organisations able to share this user data?

Let us take an example incident in which a service provider identifies suspicious activity from a user's account. The service provider reaches out to the user's home organisation to alert them of the abnormal behaviour. The home organisation then begins forensics and identifies multiple compromised accounts. Through the logs kept by this organisation, they are able to identify all other services accessed by the compromised accounts and alert them. In practice there are many possible points of failure for this simple flow. Firstly, a service provider or identity provider, may be unwilling to share the details of the compromise for fear of damage to their reputation; federation participants may not be bound to confidentiality protocols and disclose sensitive information. It may not be technically possible to identify compromised accounts and their activity. Finally, there may be no way for security contacts to identify each other and initiate collaborative response.

A particular challenge within this federated incident scenario is to identify the party responsible for resolving an incident. Individual players will likely have multiple conflicting tasks at their own organisations and ensuring that an incident is actively managed is non trivial. The issue is scope. EduGAIN is a skeleton service which simply enables the sharing of metadata between federations, hence there is no centrally managed support. Once a security incident has spread outside a national federation there is a strong dependency on voluntary contribution of individuals which extends beyond their contractual remit. Reflecting the structure of the federations themselves, it is essential that distributed responsibility forms the baseline for collaboration.

4. Building Trust Between Organisations

As highlighted in section 1, trust is fundamental for inter-federation to be successful. Without trust, individual federation participants may opt-out of eduGAIN, limit attribute release or choose to block external users. Whilst policies aim to establish formalised trusted relationships, there is no shared inter-federation policy regarding security incident response. An entity may be hesitant to trust in collaborative incident response due to fear for their reputation or uncertainty in the operational security capability and personnel availability of partner entities. There are numerous reasons an organisation may choose to hinder federated access due to security concerns; one solution is to create a trust framework addressing known security requirements.

4.1 The Security Incident Response Trust Framework for Federated Identity

The need for a Security Incident Response Trust Framework for Federated Identity (Sirtfi) was identified in the 2013 paper "A Trust Framework for Security Collaboration among Infrastructures"[7] and was picked up by REFEDS, leading to the creation of a Sirtfi Working Group. Work was included in the Authorisation and Authentication for Research and Collaboration (AARC) Project[3] in 2015, which aims to develop an integrated cross-discipline authentication and authorisation framework. A stated objective is to address requirements for incident response. Consequently, in January 2016, version 1.0 of Sirtfi was published via REFEDS following community consultation[4].

Sirtfi breaks down the identified requirements for trust into four sections; Operational Security, Incident Response, Traceability and Participant Responsibilities. A list of assertions is defined for each of these sections. Operational Security assertions require that a security incident response capability exists with sufficient authority to mitigate, contain the spread of and remediate the effects of an incident. The Incident Response assertions assure the confidentiality of information exchanged, identify trusted contacts and guarantee a response to collaboration requests. Traceability assertions ensure that logs contain relevant information and are kept in accordance with policy. Participant Responsibility assertions confirm that end users are aware of an appropriate Acceptable Use Policy (AUP). An organisation is deemed Sirtfi compliant if they are able to agree to each and every assertion.

Compliance with Sirtfi is expressed in federation metadata and gives a transparent view of those organisations willing and able to engage. The credibility gained by asserting Sirtfi compliance will open doors within eduGAIN as organisations choose to enable authentication based on this enhanced level of trust.

5. Building Trust Between Individuals

Due to the diversity in participating federation entities, it is unrealistic to expect each organisation to employ experts in security response. The ultimate benefit of federated identity management is its ability to facilitate collaboration. This is as true for those running the infrastructure as it is for the end users. Security professionals will be able to leverage Sirtfi to identify external support and be sure of confidential information exchange when requesting help. In addition to the general security contact information, which is required for Sirtfi compliance, identifying trusted individuals at organisations will catalyse successful incident response.

Lack of trust and poor communication between individuals are recognised obstacles to information flow. Individual relationships are able to traverse political, geographical and cultural barriers and establishing these relationships via face to face meetings or effective incident response communication will create the environment required for collaboration. Actively participating in cyber-security threat intelligence trust groups boosts personal credibility and opens doors to increasingly useful circles of trust. It is via these channels that incidents are often identified and continual leverage of such communities will lead to improved incident response for inter-federation. Formalised lever-

age of such networks is expected to feature in future work on federated incident response within REFEDS and the AARC Project.

6. Summary

The increasing scale and connectedness of Research and Education federations clearly exposes an inviting vector of attack. Fundamentally, this is a problem that eclipses any technical solutions and will be solved by establishing a network of trust and shared responsibility between federation participants. The concern for this emerging vector of attack should not be a reason for failure of an international identity federation; the requirements to manage the paradigm shift away from traditional security incident response have been identified and work is ongoing to prepare the community.

It is essential that the community approach federated security proactively and build the trust frameworks required both between organisations and individuals.To formalise the trust between organisations a Security Incident Response Trust Framework for Federated Identity is proposed, see section 4.1 on Sirtfi. Work is ongoing within REFEDS[2] and the AARC Project[3] to promote the adoption of Sirtfi throughout eduGAIN[5].

References

[1] "Federated Identity Management for Research Collaborations", D Broeder et al, April 23 2012, CERN-OPEN-2012-006,
https://cdsweb.cern.ch/record/1442597

[2] REFEDS, the Research and Education Federations Group,
https://www.refeds.org

[3] AARC, Authentication and Authorisation for Research and Collaboration,
https://www.aarc-project.eu

[4] A Security Incident Response Trust Framework for Federated Identity v.1,
https://refeds.org/wp-content/uploads/2016/01/Sirtfi-1.0.pdf

[5] eduGAIN,
https://www.edugain.org

[6] Interfederation Crash Course, SWITCHaai, aai@switch.ch,
https://www.switch.ch/aai/support/presentations/crash-course-2013/InterFed_all_slides.pdf

[7] "A Trust Framework for Security Collaboration among Infrastructures", D. Kelsey et al, in the Proceedings of the International Symposium of Grids and Clouds 2013 (ISGC 2013), Taipei, Taiwan, March 17-22, 2013, PoS (ISGC 2013) 011
https://www.eugridpma.org/sci/