19-04-2016

Deliverable D15.2 Report on the Achievements of JRA3 Trust and Identity Research Task 2 Trust and Identity Technologies and Recommendations on Future Work

Deliverable D15.2

Contractual Date:	30-04-2016
Actual Date:	19-04-2016
Grant Agreement No.:	691567
Work Package/Activity:	15/JRA3
Task Item:	Task 2
Nature of Deliverable:	R (Report)
Dissemination Level:	PU (Public)
Lead Partner:	SURFnet
Document Code:	D15.2
Authors:	R. Poortinga-van Wijnen (SURFnet)

© GEANT Limited on behalf of the GN4-1 project.

The research leading to these results has received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement No. 691567 (GN4-1).

Abstract

This deliverable reports on the achievements of Joint Research Activity 3 Trust and Identity Research, Task 2 Trust and Identity Technologies during GN4-1 in the areas of Certificate Transparency, identity federation extensions, federation technology test tools (FedLab) and multi-factor authentication. It also makes recommendations on future work.

Table of Contents

Table of Figures

Table of Tables

Executive Summary

This deliverable reports on the achievements of Joint Research Activity 3 Trust and Identity Research, Task 2 Trust and Identity Technologies (JRA3 T2) during GN4-1, and makes recommendations on future work.

The main goal of the Trust and Identity Technologies Task is to increase the quality, interoperability, security and reach of identity federations and of the tools used by those federations. To achieve this, the Task has worked in the areas of Certificate Transparency, protocols used for identity federations, test tools, and security measures in federations, such as multi-factor authentication. To focus the research, the following objectives were set:

- Enhance the trustworthiness of certificates.
- Extend the reach of identity federations.
- Increase the quality and interoperability of identity federations.
- Enhance the security of identity federations.

The work item on the trustworthiness of certificates has resulted in an independent, stable, open source implementation of Certificate Transparency (CT), named Catlfish, the architecture of which has been designed with a multi-organisation Internet-wide deployment environment (such as GÉANT) in mind. Catlfish is currently being used for a log that holds around 6 million certificates. This work item has also created a Gossip protocol, which allows clients of CT logs to detect potential attacks carried out by the log itself. The recommendation for Certificate Transparency is to continue development and work towards pilot and production deployment with GÉANT partners.

The Identity Federation Extensions work item has devised an approach to creating and signing metadata for OpenID Connect (OIDC) identity federations, aiming to fill a gap in current authorisation and authentication standards: OIDC potentially offers a more dynamic solution to managing federations than the de facto standard SAML, but there is currently no support in the base OIDC standard for building federations. The recommendation is to start testing with a limited number of willing federations to gain experience with this approach in a production setting, alongside existing (SAML) federations.

The FedLab team has redesigned the test suite and created a single common framework, called aatest, to be used by all test tools, making it easier to develop, add and maintain test tools for new protocols. The development of a test editor for SAML2 facilitates adding tests created by different people to the total test pool available. The OIDC identity provider test tool has now been in

production for about a year and is the official test tool used for OIDC certification. During that period, 23 installations have been certified for at least one profile, with the total number of tested instances somewhere between 100 and 150. It is recommended that support for FedLab should continue in future, particularly to develop an OIDC service provider-instance testing tool and to enable the user-managed access test tool to support three-entity scenarios, either in the next phase of the GÉANT project or in/by other suitable contexts such as REFEDS.

The multi-factor authentication (MFA) work item has identified various use-cases for MFA and identified two major issues that prevent its widespread adoption in the NREN and R&E community: 1) vendor lock-in by single-vendor solutions, and 2) lack of support for step-up authentication (i.e. single-factor initial authentication, with a second authentication factor required if and when the user requests access to specific sub-services) in standards and implementations. This Task recommends that activities in the field of MFA, and to promote generic solutions and implementations to stimulate adoption by the R&E field and to prevent vendor lock-in, be continued in the next phase of the GÉANT project.

Task members have made significant contributions to standardisation work with regard to Certificate Transparency (IETF), OIDC (the OpenID Foundation) and user-managed access (Kantara working groups).

JRA3 T2 has carried out several high-profile dissemination activities during GN4-1, including at TNC2015 and Internet2 TechExchange, and the delivery of two two-day courses for GÉANT project participants and InCommon members. It is recommended that the level of activity be maintained, to continue raising awareness of the advances and enhancements in trust and identity technologies being made available to the R&E community and to elicit feedback.

1 Introduction

The main goal of GN4-1 Joint Research Activity 3 Trust and Identity Research, Task 2 Trust and Identity Technologies (JRA3 T2) is to increase the quality, interoperability, security and reach of identity federations and of the tools used by those federations. To achieve this, the Task has worked in the areas of Certificate Transparency, protocols used for identity federations, test tools, and security measures in federations, such as multi-factor authentication. To focus the research, the following objectives were set:

- **Enhance the trustworthiness of certificates**, by continuing the work on Certificate Transparency started in GN3plus.
- **Extend the reach of identity federations**, by exploring how new technologies such as OpenID Connect (OIDC) can be introduced into existing identity federations and how federations based on different technologies can be bridged.
- **Increase the quality and interoperability of identity federations**, by providing test and monitoring tools ("FedLab") for various technologies and entities used within a federation.
- **Enhance the security of identity federations**, by investigating vendor-independent two-factor authentication and creating a reference vendor-independent architecture for implementing two-factor authentication.

All (intermediate) results were regularly shared with the Research and Education Federations (REFEDS) community, to obtain feedback.

The following section outlines the achievements of the work items aligned to these objectives in the Trust and Identity Technologies Task.

2 Achievements of the Trust and Identity Technologies Task

This section outlines the achievements of JRA3 T2 in the areas of certificate transparency, identity federation extensions, federation test tools, multi-factor authentication and standardisation work.

2.1 Certificate Transparency

2.1.1 Introduction

Certificate Transparency (CT) is an open framework originally started by Google "to detect SSL certificates that have been mistakenly issued by a certificate authority or maliciously acquired from an otherwise unimpeachable certificate authority" [CT]. In essence CT is an append-only externally verifiable log of certificates, logging the existence of Transport Layer Security (TLS) certificates as they are issued or observed, allowing anyone to inspect the log. The goal is to identify certificates that have been either mistakenly issued or issued by a compromised certificate authority (CA). Note that CT in itself does not prevent the inadvertent or malicious issuance of "wrong" certificates; rather, it allows these occurrences to be more easily detected, thereby solving a blind-spot in the current Public Key Infrastructure (PKI) system, which assumes that CAs are always fully trustworthy and cannot be compromised. Based on historical security compromises, this assumption is both dangerous and wrong. A good introduction to the reasoning behind CT can be found on the Certificate Transparency website [CT-WhatIs].

2.1.2 Certificate Transparency Achievements

2.1.2.1 *Catlfish*

As part of GN4-1, the CT work item team, consisting of people from the Swedish NREN, SUNET, and the Royal Institute of Technology in Stockhom (KTH), participating in the project through NORDUnet, has designed and implemented a distributable implementation of Certificate Transparency, called Catlfish [Catlfish], as specified in [RFC6962].

Catlfish is an implementation of a Certificate Transparency log that is designed to work well in a multi-organisation Internet-wide deployment environment such as GÉANT, as opposed to the first implementation by Google, which is designed to work well in a single (albeit a very distributed) organisation.

Figure 2.1 shows the architecture of the Catlfish CT implementation. The system is divided into front-end, storage, merge and signing nodes. Front-end nodes provide the public HTTPS service for submission and queries. Submitted certificate chains are validated and sent to all storage nodes. When a configured quorum of the storage nodes report that they have stored the data on disk, a signature is requested from a signing node and a Signed Certificate Timestamp (SCT) is returned to the submitter.

The primary merge node periodically collects all new log entries from the storage nodes, orders them, builds a new tree, signs the root hash and distributes the new entries and the signed root hash to the front-end nodes.

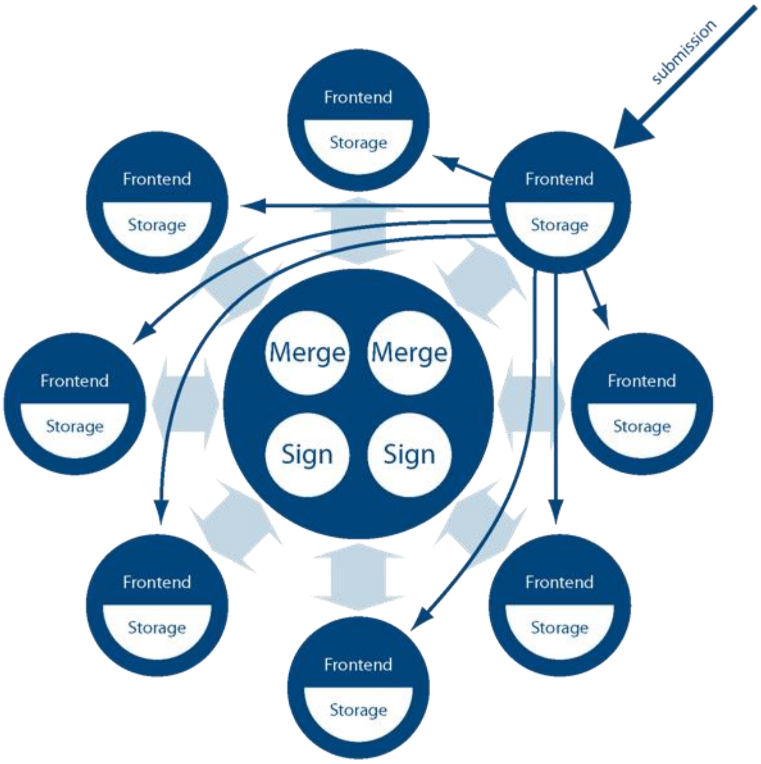

Figure 2.1: Catlfish Certificate Transparency architecture

2.1.2.2 *Gossip*

The purpose of CT is to reduce the need for absolute trust in CAs by making it easier to detect inappropriately issued certificates. However, if this CT solution means that the logs themselves need to be trusted for it to work, then it merely moves the problem around instead of solving it. In effect, the whole CT setup should be such that no component needs to trust any other component (fully).

One potential approach for attacking the CT system is a partitioning attack, where a log provides different views to different clients. Each client would be able to verify the append-only nature of the log but – in the extreme case – each client might see a unique (and seemingly correct) view of the log, leaving targeted groups of clients of that log vulnerable if they inadvertently trust and rely on that information.

To thwart this type of attack, log clients need to talk to each other ("gossip") about their view of a log to verify that the log does not present different views to different clients. One of the difficulties is that while clients need to talk to each other to verify the log's trustworthiness, they have to be careful not to divulge any information that may be privacy sensitive. After all, the act of verifying a specific certificate from the log may be privacy sensitive information in itself.

The CT team has led the effort of standardising Certificate Transparency Gossip [Gossip] as part of the IETF Public Notary Transparency working group [Trans]. The standardisation work has produced a draft that has been accepted by the working group, with a goal of being ready for last call during the second half of 2016.

2.1.3 Conclusions and Recommendations

The CT team has created an independent open source implementation of Certificate Transparency (Catlfish), thereby satisfying the IETF standards process requirement that states that two or more independent interoperable implementations must exist.

The design and implementation of Catlfish have resulted in a stable and useful CT system, currently used for a log that holds around 6 million certificates. The software is also used as a building block for other types of public append-only externally verifiable logs.

The work on Gossip significantly enhances Certificate Transparency goals by providing a means for thwarting potential partitioning attacks on the CT system.

The architecture of Catlfish was designed with a multi-organisation Internet-wide deployment environment such as GÉANT in mind. The CT team recommends that work on Certificate Transparency and Gossip be sustained, by continuing its development and starting a pilot and production deployment in the next phase of the GÉANT project, preferably aligning it with tests of Certificate Transparency browser support being developed by other groups represented in the Public Notary Transparency working group [Trans].

2.2 Identity Federation Extensions

2.2.1 Introduction

In addition to SAML2, which is the current de facto standard within NREN communities, there is a new standard for authorisation and authentication, OpenID Connect (OIDC), which could potentially be an alternative. The current OpenID Connect standard assumes a use case of one service provider talking to one or more identity provider(s); there is no support in the base standard for building identity federations. The goal of the Identity Federation Extensions work item is to fill the gaps in this area, by adding those extra specifications that are needed to construct identity federations based on OIDC.

2.2.2 OpenID Connect Federation Design

Current SAML federations are very static; changes in metadata such as key rollover (a process whereby one key is systematically replaced by another key in SAML metadata) can take as much as a day. Adding new members is time-consuming and labour-intensive, since all present processes demand human involvement. In the typical setup and context of SAML federations, where the relations between service providers and identity providers are relatively static, this makes sense since automating those processes will not provide big gains. However, the envisaged context of OIDC is much more dynamic and therefore OIDC provides the necessary functions for managing federation participants in a correspondingly much more dynamic way. The team has tried to leverage this in its design for OIDC-based federations.

The basis of SAML federations is the signed metadata [Metadata]. Since SAML is an XML-based protocol (using SOAP, an XML RPC technology), the metadata is XML-based as well. A similar approach is needed for OpenID Connect, but since OIDC is JSON-based, the logical approach is to use JSON-based metadata as well. The basic idea is therefore straightforward: use “software statements”. Essentially, software statements are JSON documents containing information about an entity, which can be an OpenID provider (OP, synonymous with identity provider (IdP)) or relying party (RP, synonymous with identity provider (IdP)). This information can be signed by the federation operator before the OP or the RP actually enters the federation or whenever some information in the software statement about those entities has changed.

If at some later time an RP wants to talk to an OP within the federation, the RP does what is called a client registration [OIDC-CReg]; this is completely dynamic and will in the normal case contain information about the client that is not easily verified. In the Identity Federation Extensions team’s design, the most significant client information has been verified and signed by the federation operator, so the OP can have confidence in the correctness of the client information and can act upon it accordingly.

Information about the OP is handled in a symmetric way. The most significant information about the OP is sent to the federation operator for signing. The resulting signed software statement will be used by the OP whenever it distributes information about itself to the client.

As at the end of GN4-1, development of software implementations that support this model [OIDC-Fed] and a SAML2 to OIDC proxy [SATOSA] has progressed to the point where it is ready for testing in a real-life setting. It is expected that federations that expressed willingness to set up an OpenID Connect identity federation in parallel to their existing SAML2 federation will start testing at the beginning of the next phase of the GÉANT project.

2.2.3 Conclusions and Recommendations

The idea of using software statements for OpenID Connect federations has been presented by this work item's team and discussed at a number of workshops and conferences. The development of software implementations using this approach has shown that the basic idea is feasible. Proposed and recommended future work is to start testing OpenID Connect federations in a production setting parallel to existing SAML2 federations.

2.3 FedLab

2.3.1 Introduction

The goal of FedLab [FedLab] is to provide tools that federations and operators of entities within those federations can use to verify that an individual entity (be it an identity provider or a service provider) works properly before it is introduced into a federation. By testing an entity beforehand, any bug or inadvertent misbehaviour is easier to track down than it would be in a live setting. Therefore most (and the most common) bugs and misconfigurations can be detected and solved before an entity enters a federation, in most cases leading to a seamless operational introduction into a federation. Once an entity is in a federation, other tools exist to make sure the setup is still working properly. A notable example is the eduGAIN Connectivity Check Service (ECCS) ECCS. However, while ECCS does check whether an IdP is still present and responsive, it does not check whether it is working properly and conforming to the relevant profile.

2.3.2 Implementation/Configuration Verification Tools

During the project period, an extensive overhaul was carried out by the FedLab team to create a single common framework, called aatest, to be used by all test tools. The common framework provides a foundation, on top of which specific test tools can be developed. This approach simplifies development and maintenance of the various test tools and reduces the time it takes to develop a test tool for a new protocol. Building on this foundation work, it was planned to develop test tools for the protocols of most interest to the GÉANT community:

- SAML2, for both IdP and SP.
- OpenID Connect, for both OP and RP.
- User-managed access (UMA), for resource server, authorisation server and client.

However, during the project, the number of personnel in the team from NORDUnet responsible for a number of work items within this Task decreased significantly as a result of people leaving for other employment. The specialist nature of the work meant that finding replacement resource before the end of the project was unrealistic. Development effort for FedLab was therefore concentrated on the SAML2 and OIDC test tools [SATOSA; OIDC-Fed] (see Section 2.2.2).

2.3.2.1 *SAML2*

For SAML2, working versions of IdP- and SP-instance testing are available. A test editor was developed to make construction of new tests easier. The test editor is intended to encourage and facilitate adding tests created by different people to the total test pool available. Federations can then choose their own set of tests from that pool to create a test suite that best matches their specific setup and configuration.

2.3.2.2 *OpenID Connect*

The OIDC OP test tool has now been running in production for about a year and is the official test tool used for OIDC certification. During that period, 23 installations have been certified for at least one profile, with the total number of tested instances somewhere between 100 and 150 [Certification].

For RPs, the existing test tool – operational since September 2015 – was geared towards testing (OIDC RP) libraries rather than RP instances. Of course, for a production setting, using a tested library is not enough. RP-instance testing is needed for testing proper configuration and (therefore) proper operation of an RP before it enters production, to ensure a smooth operational introduction. Therefore, the development of an RP-instance testing tool started near the end of the project. A fully completed RP-instance testing tool is not expected before the end of GN4-1, but development is expected to continue in the next phase of the project.

2.3.2.3 *User-Managed Access*

The problem with testing user-managed access (UMA) settings is the involvement of three entities in most scenarios: the resource server (RS), the authorisation server (AS) and the client (C). Tests that only involve two parties (RS-AS and C-AS) are catered for by the tool. A common RS API is needed to make testing of scenarios with three entities possible. However, the Kantara UMA working group has not yet reached agreement on an RS API. Therefore, the UMA test tool cannot be developed further to include three-entity scenarios and is currently restricted to testing two-entity scenarios only.

2.3.3 Conclusions and Recommendations

FedLab provides an invaluable test service for the different entities and protocols that exist in the federation space. The complete redesign of the test suite ensures that new protocols and tests can be added as efficiently as possible. It is the Task's recommendation that support and development of FedLab be continued, either in the next phase of the GÉANT project or in other suitable environments such as REFEDS.

2.4 Two-Factor Authentication

2.4.1 Introduction

The NREN community operates several services where key management for secure client access is becoming a critical factor. However, discussions on the continued domination of passwords in this area and/or alternative authentication methods have not yet produced a definitive answer. From a security perspective two- or (in general) multi-factor authentication (MFA) methods – for instance, based on hard- or software tokens – are a good choice, but only if deployed in a homogeneous enterprise-wide environment.

Deployment of those mechanisms in the research and education (R&E) community has been hindered mostly by the vendors in this area being focused on their own single-vendor solutions. Researchers requesting access to a federated service provided by another organisation (i.e. a service provider) and associated use cases are potentially confronted with vendor lock-in, or the need to buy multiple tokens. Interoperability or the option to use MFA in a (inter-)federated way are currently unsupported. Common standards, e.g. PKCS#11 (cryptographic tokens) or ISO/IEC 7816 (smart cards), address those interoperability issues on the service provider side, but each home organisation has to decide on a unified solution (i.e. issuance and rollout of new tokens/smartcards, revocation and re-issuance of stored credentials such as user certificates or even time- or event-counter synchronisation).

Multi-factor authentication is a method of access control in which a user is only granted access to a system or service after successfully presenting several separate pieces (factors) of evidence about her identity to an authentication mechanism. All of these methods can be categorised into sets of factors, based on their specific properties:

- Something the user knows (e.g. password, PIN).
- Something the user owns (e.g. physical token, smartcard, mobile phone).
- Something the user is (e.g. retina, fingerprint, voice).

Authentication is called multi-factored if it includes more than one of these factors (e.g. password combined with a physical token). These factors have to be independent. This means that gaining access to one of the factors used must not trivially grant access to the second one. A counter-example is a software phone accessible using the user’s password or storing a software one-time password (OTP) generator on a device accessible using just the first factor, e.g. user’s password.

Larger social networks such as Google+ and Facebook have begun to introduce additional information into the authentication process that can be classified as context or “something the user does” (e.g. characteristic usage patterns, physical location from where the login has been initiated, or even the usual login time of a user).

The biggest advantage of using a combination of all these factors lies in a higher security level through resilience to targeted impersonation, throttled or un-throttled guessing, and credential theft. Using such a second factor will introduce resistance to phishing attempts, since it will be used regardless of the user's device or location, meaning that the user will either be alerted to the phishing attempt because no additional factor is used, or that successful theft of credentials will not give the criminal access to the real service(s) since the user is still in possession of the additional factor needed for authentication.

The combination of a first and a second factor therefore also mitigates first-factor risks often seen as the result of a successful phishing attack, online-password guessing or even theft of a (single) factor.

2.4.2 Use-Cases

In addition to the organisation-wide scenarios, where MFA helps to better protect critical resources, R&E also has federated or even inter-federated scenarios (across organisation boundaries). For such use-cases it is important to distinguish between:

- **Direct multi-factor authentication**, i.e. a service provider signals MFA as the only acceptable authentication method.
- **Step-up multi-factor authentication**, i.e. a service provider only needs a single-factor authentication to start with, but if the user requests access to specific sub-services then a second authentication based on MFA is required.

The service provider's signalling requires further differentiation, which results in the definition of various use-cases grouped by primary interested party: the end user, the identity provider (IdP) or the service provider (SP). These use-cases are specified in the following sections. T2 has tried to formalise these use cases in order to support future implementation.

2.4.2.1 *End-User Use-Case*

In this use-case, the usage of MFA is based on the user's decision to select MFA at her IdP's login page when accessing a service, for example a service that processes (privacy) sensitive data.

ID	User-1
Short description	User selects to use MFA
Actors	User, IdP
Preconditions	None
Description	1. User requests access to a specific resource provided by a service provider 2. User selects her IdP and is redirected for authentication 3. User's IdP provides different login mechanisms a. Password-based authentication b. Multi-factor-based authentication 4. After successful user authentication the IdP responds to

	the SP, asserting successful MFA-based authentication
Results	User can access requested resource based on MFA
Known issues	An attacker, who is able to acquire the first factor, is able to perform the authentication without the second factor, unless the SP enforces the use of MFA

Table 2.1: User-1 use-case

The IdP has to respond with an appropriate response message. If SAML is used, then this SAML response message contains an *AuthnContext* element referencing the appropriate *AuthnContextClass*.

2.4.2.2 *IdP Use-Case*

In this use-case, the usage of MFA is prescribed by the IdP; it therefore offers only MFA authentication at login.

ID	IdP-1
Short description	IdP requires users to use MFA
Actors	User, IdP
Preconditions	None
Description	1. User requests access to a specific resource provided by a service provider 2. User selects her IdP and is redirected for authentication 3. User’s IdP provides an MFA login mechanism 4. After successful user authentication the IdP responds to the SP, asserting successful MFA-based authentication
Results	User can access requested resource based on MFA
Known issues	None

Table 2.2: IdP-1 use-case

2.4.2.3 *SP Use-Cases*

For SPs, two distinct use-cases exist: one that requires MFA directly at login and one that triggers an additional MFA authentication later, for example because the user wants to access particularly valuable or (privacy sensitive) functions or data.

ID	SP-1
Short description	SP requires users to use MFA at login
Actors	User, IdP, SP
Preconditions	IdP supports MFA. (If IdP does not support MFA then the IdP has to send an appropriate error message.)
Description	1. User requests access to a specific resource provided by a service provider 2. User selects her IdP and is redirected for authentication 3. User's SP requests MFA authentication and the IdP can provide MFA 4. After successful user authentication the IdP responds to the SP, asserting successful MFA-based authentication
Results	User can access requested resource based on MFA
Known issues	None

Table 2.3: SP-1 use-case

The SP sends an authentication request message containing a SAML *AuthnContextClassRef* to one of the MFA classes (only).

ID	SP-2
Short description	SP requires users to use MFA for specific actions
Actors	User, IdP, SP
Preconditions	IdP supports MFA
Description	1. User requests access to a specific resource provided by a service provider 2. User selects her IdP and is redirected for authentication 3. User's IdP provides different login mechanisms a. Password-based authentication b. Multi-factor-based authentication 4. After successful user password-based authentication the IdP responds to the SP, asserting successful authentication 5. The user tries to access a specially protected part of the service 6. The SP wants to step-up the authentication level of the user and requests MFA from the IdP

	7. The user authenticates using MFA at the IdP and the IdP responds to the SP, asserting successful MFA-based authentication
Results	User can access requested resource based on MFA
Known issues	None

Table 2.4: SP-2 use-case

T2 also identified a variant of this use-case in which a service provider prefers, but does not require, MFA. In this case the authentication request message can contain more than one *AuthnContextClassRef* elements referencing one of the preferred MFA classes (listed first) and a base-level authentication indicating, for example, password-based authentication. This variant requires that the SP knows which authentication methods the IdP supports. If the IdP supports neither, then it has to send an error message back enabling the SP to send a second authentication request that does not include a context reference element to allow IdPs to assert a different context.

2.4.3 Protocol-Specific Recommendations

Below are recommendations and configuration examples for specific protocols. Following these guidelines will ensure interoperability and ease of deployment.

2.4.3.1 *SAML*

The SAML 2.0 standard itself does not specify specific methods for handling MFA; the type of authentication performed by the IdP is beyond the scope of the standard. Consequently, the standard also does not contain any profiles for step-up authentication. The Shibboleth SAML implementation has some basic support for MFA in its IdP component. A community-created extension created specifically for the – now deprecated – version 2 of this IdP supports the Initiative for Open Authentication HMAC-based One Time Password algorithm (OATH-HOTP) second-factor standard [RFC4226]. This extension allows a user to use a second factor besides his username/password credentials.

Because the Task 2 team focuses on the definition, analysis and formal description of different MFA use-cases, implementing these has to be done in the next phase of the GÉANT project. Since version 3 of the Shibboleth IdP, limited support for MFA is built in. The IdP v3 implements a multi-context broker (MCB), which allows users to use different authentication contexts. An authentication context describes the authentication method and additional information that is relevant to the authentication process. Using the MCB, the IdP can define multiple authentication contexts, such as password-based authentication or multi-factor authentication, and prompt the user with or for the appropriate context. The appropriate context is determined based on the context requested by the SP, the context the user selects, and the context hierarchy (respectively). The context hierarchy specifies which context automatically satisfies another context. This is used when the IdP of an already authenticated user has to determine whether to ask the user for authentication again if the context from the SP differs from the context the user's session was initially started with. This can be configured by an IdP administrator. The *authn/standard* bean defines a list referencing an authentication context class with ID `http://id.example.org/standard`. The second *authn/strong* bean lists authentication context classes `http://id.example.org/standard` and `http://id.example.org/strong`. This can be used to define that, if a user is already authenticated using (strong) MFA, she does not have to re-authenticate if the SP is requesting (standard) password authentication. However, if the user is only authenticated via password and the SP is requesting MFA, the user has to provide the MFA authentication.

Using these features of the Shibboleth IdP v3, the use-cases User-1, IdP-1 and SP-1 can be implemented. A step-up of the authentication level as described in use-case SP-2 is also supported. The SP would request an authentication using MFA from the IdP, the IdP would determine – based on the context hierarchy – that the existing password authentication is not sufficient and prompt the user for authentication using MFA.

2.4.4 Conclusions and Recommendations

Two major issues prevent widespread adoption of multi-factor authentication in the NREN and R&E community: 1) vendor lock-in by single-vendor solutions, and 2) lack of support for step-up authentication in standards and implementations. Because some standards already exist on the SP side to prevent vendor lock-in, the IdP side has to be further investigated. However, the improvement in security provided by MFA is needed for keeping service and (sensitive) data safe. Therefore, this Task recommends that activities in the field of MFA, and to promote generic solutions and implementations to stimulate adoption by the R&E field and to prevent vendor lock-in, be continued in the next phase of the GÉANT project.

2.5 Standardisation Work

Members of the Task have participated in the IETF (e.g. for Certificate Transparency), the OpenID Foundation (FedLab OpenID Connect certification testing), and the Kantara working groups (user-managed access).

2.6 Dissemination

Table 2.5 below shows the main dissemination activities JRA3 T2 has undertaken during GN4-1 to present, raise awareness of and encourage discussion about its progress and achievements. In addition to giving presentations at various conferences, one out-of-the-ordinary dissemination activity in which JRA3 T2 has been involved is training: two members of the Task have given two two-day courses on OAuth2, OIDC and UMA for GÉANT project participants and two two-day courses for InCommon members.

Type	Title	Event	Location	Date	Link
Presentation	eduKEEP	Internet2 TechExhange	Cleveland, OH	07-10-2015	http://meetings.internet2.edu/media/medialibrary/2015/10/16/20151007-Kremers-eduKEEP.pdf
Workshop	OIDC federation	Internet Identity Workshop	Mountainview, CA	27-10-2015	–
Workshop	OIDC test tool RP&OP	Internet Identity Workshop	Mountainview, CA	28-10-2015	–
Presentation	Catlfish – An Implementation of Certificate Transparency in GN3+	TNC2015	Porto, Pt	17-06-2015	https://tnc15.terena.org/core/presentation/205
Course (x2)	OJOU (OAuth2/JW*/OpenID Connect/UMA) course	GÉANT audience	Utrecht, NL	23&24-09-2015 09&10-02-2016	Presentation material: https://github.com/rohe/ojou_course.git
Course (x2)	OJOU course	InCommon audience	Denver, CO	22&23-02-2016 24&25-02-2016	Course material: https://github.com/rohe/openid_course

Table 2.5: Dissemination activities

3 Conclusions and Recommendations

To summarise the conclusions and recommendations for the areas of work undertaken by JRA3 T2 during GN4-1:

- The Certificate Transparency team has created an independent, stable, open source implementation of CT, named Catlfish, currently being used for a log that holds around 6 million certificates. The work on Gossip significantly enhances CT goals by providing a useful means of preventing potential partitioning attacks on the CT system. The architecture of Catlfish was designed with a multi-organisation Internet-wide deployment environment such as GÉANT in mind.

  The CT team recommends that work on Certificate Transparency and Gossip be sustained, by continuing its development and starting a pilot and production deployment in the next phase of the GÉANT project.

- The idea of using software statements for OpenID Connect federations – offering a more dynamic solution to managing federations than SAML, and filling a gap in current authorisation and authentication standards – has been presented and discussed by this work item's team at a number of workshops and conferences. The development of software implementations using this approach has shown that the basic idea is feasible.

  Proposed and recommended future work is to start testing OpenID Connect federations in a production setting parallel to existing SAML2 federations.

- FedLab provides an invaluable test service for the different entities and protocols that exist in the federation space. The complete redesign of the test suite ensures that new protocols and tests can be added as efficiently as possible. The development of a test editor for SAML2 facilitates adding tests created by different people to the total test pool available. The OIDC identity provider test tool has now been in production for about a year and is the official test tool used for OIDC certification. During that period, 23 installations have been certified for at least one profile, with the total number of tested instances somewhere between 100 and 150.

  It is the Task's recommendation that support and development of FedLab be continued, particularly to develop an OIDC service provider-instance testing tool and to enable the user-managed access test tool to support three-entity scenarios, either in the next phase of the GÉANT project or in other suitable environments such as REFEDS.

- Two major issues prevent widespread adoption of multi-factor authentication in the NREN and R&E community: 1) vendor lock-in by single-vendor solutions, and 2) lack of support for step-up authentication in standards and implementations. However, the additional security offered by MFA is indisputable.

This Task recommends that activities in the field of MFA, and to promote generic solutions and implementations to stimulate adoption by the R&E field and to prevent vendor lock-in, be continued in the next phase of the GÉANT project.

- Task members have made significant contributions to standardisation work with regard to Certificate Transparency (IETF), OIDC (the OpenID Foundation) and user-managed access (Kantara working groups).
- JRA3 T2 has carried out several high-profile dissemination activities during GN4-1, including at TNC2015 and Internet2 TechExchange, and the delivery of two two-day courses for GÉANT project participants and InCommon members.

  It is recommended that the level of activity be maintained, to continue raising awareness of the advances and enhancements in trust and identity technologies being made available to the R&E community, and to elicit feedback.

References

[Catlfish]	https://www.ct.nordu.net/
[Certification]	http://openid.net/certification/
[CT]	https://www.certificate-transparency.org/
[CT-WhatIs]	https://www.certificate-transparency.org/what-is-ct
[ECCS]	https://technical.edugain.org/eccs/
[FedLab]	http://fed-lab.org/
[Gossip]	https://datatracker.ietf.org/doc/draft-ietf-trans-gossip/?include_text=1
[Metadata]	https://docs.oasis-open.org/security/saml/v2.0/saml-metadata-2.0-os.pdf
[OIDC-CReg]	https://openid.net/specs/openid-connect-registration-1_0.html
[OIDC-Fed]	https://github.com/its-dirg/oidc-fed
[RFC4226]	https://www.ietf.org/rfc/rfc4226.txt
[RFC6962]	https://tools.ietf.org/html/rfc6962
[SATOSA]	https://github.com/its-dirg/SATOSA
[Trans]	https://datatracker.ietf.org/wg/trans/documents/

Glossary

API	Application Programming Interface
AS	Authorisation Server
C	Client
CA	Certificate Authority
CT	Certificate Transparency
ECCS	eduGAIN Connectivity Check Service
HMAC	keyed-Hash Message Authentication Code
HOTP	HMAC-based One Time Password algorithm
HTTP	HyperText Transfer Protocol
HTTPS	HTTP Secure
IdP	Identity Provider
IEC	International Electrotechnical Commission
IETF	Internet Engineering Task Force
ISO	International Organisation for Standardisation
JRA	Joint Research Activity
JRA3	GN4-1 Joint Research Activity 3 Trust and Identity Research
JSON	JavaScript Object Notation
JW*	JSON Web Encryption, JSON Web Signing and JSON Web Tokens
KTH	Royal Institute of Technology in Stockholm
MCB	Multi-Context Broker
MFA	Multi-Factor Authentication
NREN	National Research and Education Network
OATH	Initiative for Open Authentication
OIDC	OpenID Connect
OP	OpenID Provider (synonymous with IdP)
OTP	One-Time Password
PIN	Personal Identification Number
PKI	Public Key Infrastructure
R&E	Research and Education
REFEDS	Research and Education Federations
RP	Relying Party (synonymous with SP)
RPC	Remote Procedure Call(s)
RS	Resource Server
SAML2	Security Assertion Markup Language version 2

SATOSA	A configurable proxy for translating between different authentication protocols such as SAML2, OpenID Connect and OAuth2.
SCT	Signed Certificate Timestamp
SP	Service Provider
SAML	Security Assertion Markup Language
SOAP	Simple Object Access Protocol (an XML-based RPC mechanism)
SSL	Secure Sockets Layer
T	Task
T2	GN4-1 JRA3 Task 2 Trust and Identity Technologies
TLS	Transport Layer Security
UMA	User-Managed Access
XML	Extensible Markup Language