Kipper – a Grid bridge to Identity Federation

A. K. Kiryanov[1,2,a], O. Keeble[3], A. Manzi[3]

[1] Petersburg Nuclear Physics Institute, Orlova Roscha 1, Gatchina, Russia

[2] NRC "Kurchatov Institute", Akademika Kurchatova pl. 1, Moscow, Russia

[3] European Organization for Nuclear Research CERN, Geneva, Switzerland

E-mail: [a] globus@pnpi.nw.ru

Identity Federation (IdF, aka Federated Identity) is the means of interlinking people's electronic identities stored across multiple distinct identity management systems. This technology has gained momentum in the last several years and is becoming popular in academic organisations involved in international collaborations.

One example of such a federation is eduGAIN, which interconnects European educational and re-search organisations, and enables trustworthy exchange of identity-related information.

In this work we will show an integrated Web-oriented solution code-named "Kipper" with a goal of providing access to WLCG resources using a user's IdF credentials from their home institute with no need for user-acquired X.509 certificates.

Kipper achieves "X.509-free" access to Grid resources with the help of two additional services: STS and IOTA CA. STS allows credential translation from the SAML2 format used by Identity Federation to the VOMS-enabled X.509 used by most of the Grid, and the IOTA CA is responsible for automatic issuing of short-lived X.509 certificates.

Kipper comes with a JavaScript API considerably simplifying development of rich and convenient "X.509-free" Web-interfaces to Grid resources, and also encouraging adoption of IOTA-class CAs among WLCG sites.

We will describe a working prototype of IdF support in the WebFTS interface to the FTS3 data transfer engine, enabled by integration of multiple services: WebFTS, CERN SSO (member of eduGAIN), CERN IOTA CA, STS, and VOMS.

Keywords: Federated Identity, Grid, X.509, Kipper

This work was funded in part by the Russian Ministry of Education and Science under contract №14.Z50.31.0024

1. Introduction

Identity Federation (IdF, aka Federated Identity) is the means of interlinking people's electronic identities stored across multiple distinct identity management systems. This technology has gained momentum in the last several years and is becoming popular in academic organisations involved in international collaborations.

Identity Federations like eduGAIN [eduGAIN] allow convenient and transparent user authentication without the burden of re-registration at each and every organisation he or she has collaboration with. This mostly applies to accessing sensitive or restricted information on web pages, personal mail, mailing lists, etc.

Historically Grid systems had their own methods for user authentication based on asymmetric cryptography and Public Key Infrastructure (PKI) [PKI]. Such trust infrastructures were normally coordinated by organisations like EUGridPMA [EUGridPMA] or OSG [OSG].

A usual Grid workflow includes obtaining of an X.509 personal digital certificate, registration in at least one Virtual Organisation (VO) and subsequent delegation of rights to Grid resources by means of X.509 proxy certificates. While looking somewhat complicated, for end-users accessing the Grid from UNIX-like console this was quite convenient, but with the emergence of more and more sophisticated and rich Web-interfaces this has become an impediment.

One of the main problems is that Web browsers only have a very basic support for X.509 certificates. They can handle neither custom certificate extensions (so-called attribute certificates holding VO membership data) nor proxy certificates necessary for delegation. Moreover browsers have a very limited security API for Web applications that makes it almost impossible to access browser's internal key store.

These two systems (IdF and PKI) in fact serve the same purpose: user authentication. It would be convenient to somehow link one to the other to avoid duplication and extra work on user's side. In this paper we will describe the Web-oriented software solution named Kipper, which enables translation of IdF credentials into X.509 ones and allows Web applications to talk directly to Grid resources.

2. Motivation

Today's distributed systems for scientific communities rely heavily on Web interfaces. One of the main reasons for this is interoperability and universality of Web-based user interface, as many users have different operating systems on their portable, home and office devices.

A web interface can be opened on any device, including a laptop computer, tablet or smartphone. Checking the status of the whole system or just your computing job cannot be easier. It would have been even more appealing if one could actually initiate Grid-related task (data transfer or job submission) right from the browser.

While such Web interfaces already exist, there's usually almost no connection between the user that accesses the interface and the Grid user identity that is seen by the Grid resources. The server side of such interfaces normally identifies itself to the Grid with a so-called robot proxy certificate which makes individual users indistinguishable from the resource point of view.

In order to avoid this obscurity and increase both traceability of resource usage and the level of control of the Grid resources over their users we need a mechanism that would convert Web-oriented user credentials to the Grid-oriented ones.

3. Necessary bits

For the credential translation to work we need the following parts:

1. Signed credentials coming from the IdF;

2. A service that would create an X.509 certificate based on IdF credentials;
3. A trusted Certification Authority that will sign such a certificate;
4. Optionally a VOMS service that will handle user mapping in a framework of a specific VO;
5. Kipper – a library with an API that will make interaction between aforementioned parts easy and convenient for a developer.

User interaction with IdF normally happens through a local Single Sign-On (SSO) service. When a user accesses a Web page that requests authentication, he is redirected to the standard SSO log-in page that might look like the one on fig. 1. Both SSO itself and the look of its log-in page are specific to the organisation and used software, but their functions are the same: provide a user with uniform way of authenticating himself. In the CERN case the SSO service is based on Microsoft Active Directory Federation Services (ADFS) while some other organisations rely on Linux-based Shibboleth software.

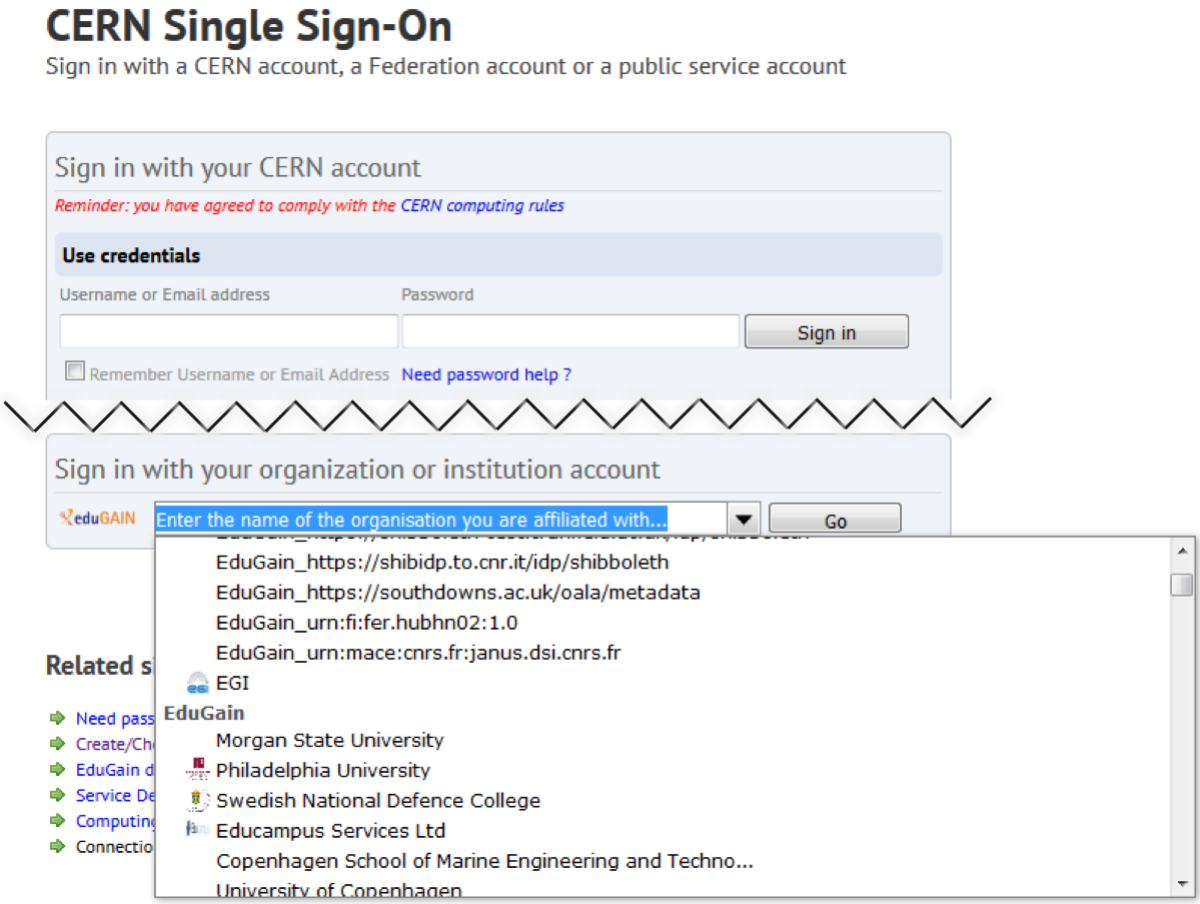

Fig. 1. CERN SSO log-in page.

Each Web server that is integrated with SSO must be properly registered and needs one extra component: an SSO plug-in, which handles interaction with SSO and provides standard authorization hooks to the Web server. For Apache on Linux there are two possible solutions both of which are natively supported by Kipper:

1. Shibboleth – widespread, supports all possible standards and can also be used as SSO server software. This is the first solution that was supported at CERN. The only shortcoming of Shibboleth is a complex XML-based configuration.

2. Mellon (aka mod_auth_mellon) – lightweight and pure SAML2 Service Provider with simple configuration. Older versions had problems with ADFS but all of them were solved with help of CERN IdF team.

IdF credentials are usually serialized in a form of SAML2 Assertion. It is a signed XML document which lists user properties like name, organisation, email address, e-groups, etc. This list is used internally by the Web server to authorize users based on property values and is not exposed to the user by default. A typical SSO login workflow is shown on fig. 2.

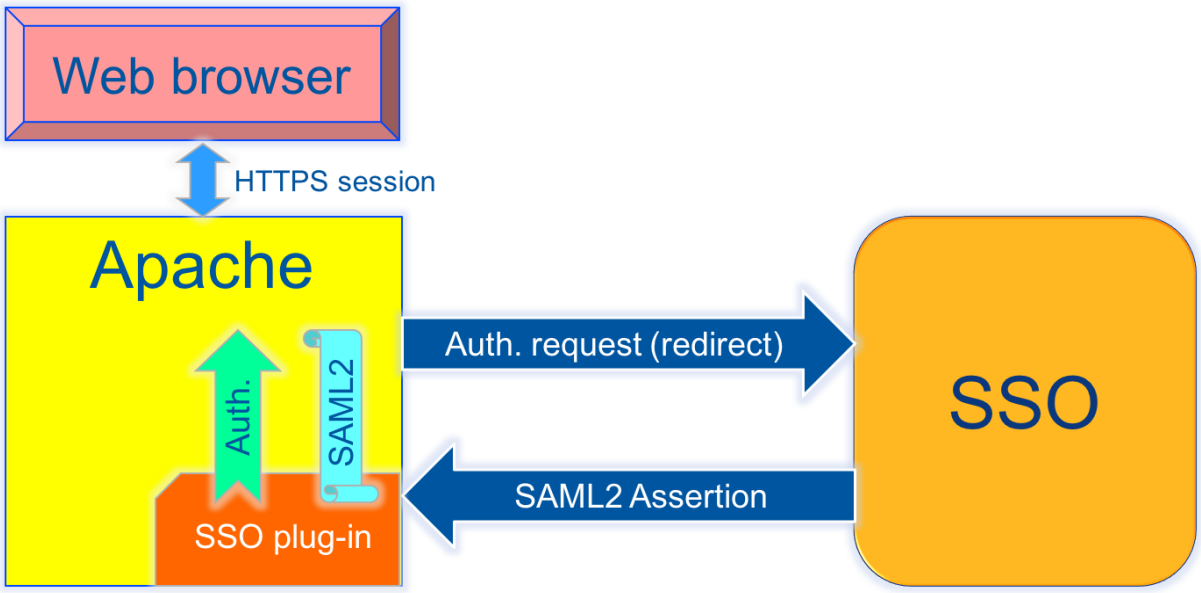

Fig. 2. SSO log-in process.

While such a workflow is sufficient for most use cases, in our case IdF credentials must be passed from the Web server to the user, or more precisely into the user's browser context. There's no security compromise in this as it is perfectly fine for a user to look at his own credentials. After that, two more extra services come into play: STS and IOTA CA.

3.1. STS

STS stands for Security Token Service [STS]. It is an implementation of the WS-Trust [WS-Trust] OASIS standard. This service is responsible for consuming SAML2-based credentials and producing X.509-based ones. It was first developed in the framework of the EMI project, and was then extended at CERN to support some extra features, including VOMS DN mapping and the CERN IOTA CA client.

In our workflow a user sends his IdF credentials to STS and receives X.509 credentials in return. From security point of view STS is a critical service because it is responsible for verifying the validity of IdF credentials and asking IOTA CA to sign an X.509 certificate derived from them. As an optional extra step, the STS can also ask a configured VOMS server to add extensions to the newly signed certificate.

The current implementation of STS allows configuration of only one Web application and VOMS server, which makes it necessary to deploy a dedicated instance of STS for every Web application and every VO it supports.

3.2. IOTA CA

An IOTA (Identifier-Only Trust Assurance) CA [IOTA] is a special profile of Certification Authority that signs short-living (days) certificates in response to requests from STS. In contrast to "normal" CAs that usually sign user certificates once a year and require in-person verification of user's authenticity, IOTA CA relies exclusively on information coming from an IdF.

One important requirement that arises from such a usage scenario is the uniqueness of the user identification string (Certificate Subject or Distinguished Name). Usually CAs have full control over DNs of signed certificates and can ensure that DNs are unique and never reused. However, this is not as easy with Identity Federation as this requirement cannot be strictly enforced.

For eduGAIN it was decided that the eduPersonPrincipalName attribute should be considered unique among its members. Identity Providers that fail to secure uniqueness of this attribute should not be enabled in SSO.

A document containing all the details for the new IOTA CA at CERN has been prepared in 2015 by the CERN IdF team. This document went through the review process of EUGridPMA and was accepted. As a result, the CERN IOTA CA has been included in IGTF Trusted Anchor Distribution since version 1.72 under a special profile, which makes it globally available at WLCG sites.

4. Implementation details

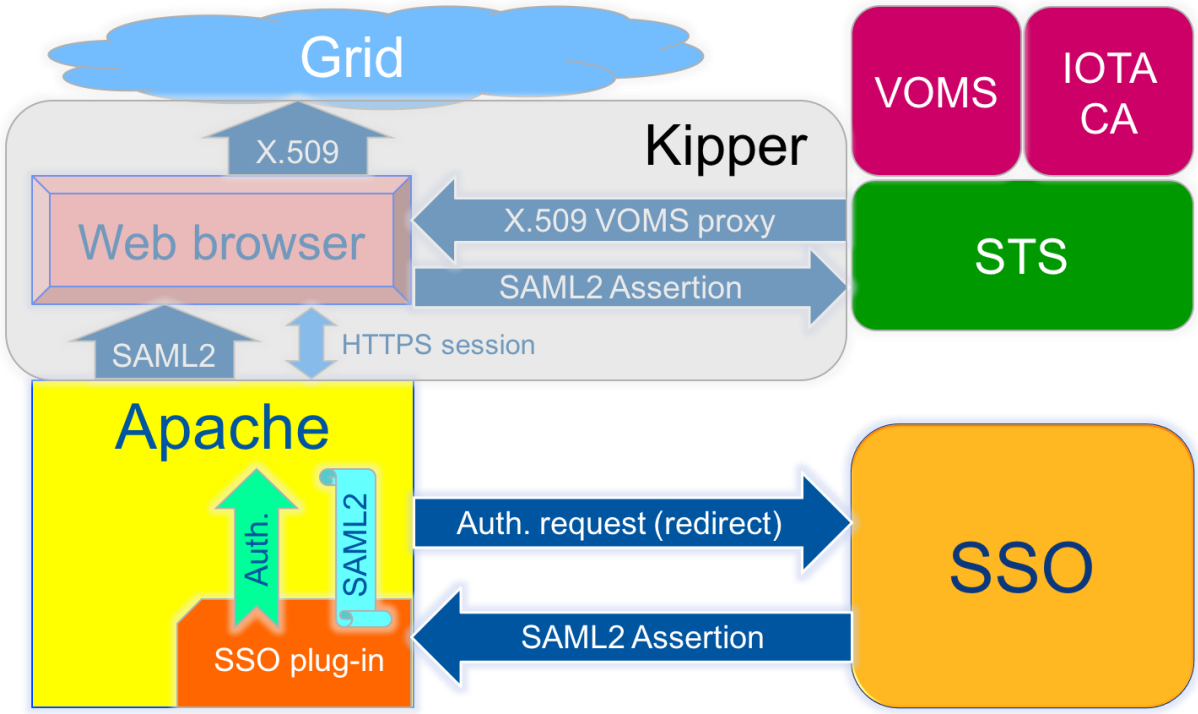

Fig. 3. Kipper workflow.

Kipper is implemented as an open-source JavaScript library that can be used directly from the Web application. Kipper handles all the basic steps in credential translation including the retrieval and

parsing of SAML2 Assertion from the Web server and interaction with STS, providing a user with an RFC-compliant X.509 proxy certificate with VOMS extensions right in the browser context (fig. 3).

Kipper and STS sources are publicly available in CERN Git repository [STS] along with documentation and configuration examples. So far Kipper can only work with the Apache web server on Linux platform (both 32-bit and 64-bit) but this may change in the future as most of the code is platform-independent.

If a user already has a "normal" Grid certificate, his IOTA DN could be mapped to the same user record in VOMS, however Grid middleware will not be aware of this mapping and will interpret different DNs as different users. This is not always a problem as in many cases access control is based on VOMS group and role rather than on user DN.

5. Use cases

Web applications that could benefit from Kipper include all sorts of portals that need to talk directly to Grid resources: data and workload management interfaces as well as various monitoring tools. Kipper key features include:

- Clear distinction between users without the need for catch-all robot proxy certificates;
- Full SSO integration, no need to maintain App-specific user database;
- VOMS support, possibility to limit the users to the members of a specific VO.

Kipper was initially developed as part of the WebFTS [WebFTS] interface to FTS3 data transfer engine. After successful proof of concept it was decided to segregate Kipper from WebFTS and make it available to other projects.

There's an ongoing integration of ATLAS PanDA monitor with SSO which will then allow exploiting Kipper to transparently access job and monitoring log files stored on Grid storage elements.

CERN is developing a pilot portal that will allow eduGAIN members that are also members of LHC VOs to get a proxy certificate out of their eduGAIN credentials.

6. Acknowledgements

The authors thank for valuable technical assistance the following people:
Henri Mikkonen
Romain Wartel
Emmanuel Ormancey

References

eduGAIN [Electronic resource]: http://services.geant.net/edugain/Pages/Home.aspx
EUGridPMA [Electronic resource]: https://www.eugridpma.org/
IOTA [Electronic resource]: https://www.igtf.net/ap/iota/
OSG [Electronic resource]: http://www.opensciencegrid.org/
PKI [Electronic resource]: http://www.net-security-training.co.uk/what-is-a-public-key-infrastructure/
STS [Electronic resource]: https://gitlab.cern.ch/sts
WebFTS [Electronic resource]: https://webfts.cern.ch/
WS-Trust [Electronic resource]: https://docs.oasis-open.org/ws-sx/ws-trust/v1.4/errata01/os/ws-trust-1.4-errata01-os-complete.html