SAML Profile for Privacy-enhanced Federated Identity Management

Rainer Hörbe, Identinetics GmbH

Abstract

This profile for the SAML WebSSO use case specifies an enhancement that allows users to limit their observability by IdPs and APs. It is based on the general Model for Privacy-enhanced Federated Identity Management[1], which describes a 3-tier model resembling an enhanced hub-and-spoke federation model. It includes the SAML WebSSO and SLO profiles, and adds messaging capabilities.

Document Version 1.2
Date: 6 May 2015

1 Interfaces

This section describes the extensions to a SAML WebSSO use case as specified in the PE-FIM model referenced above.

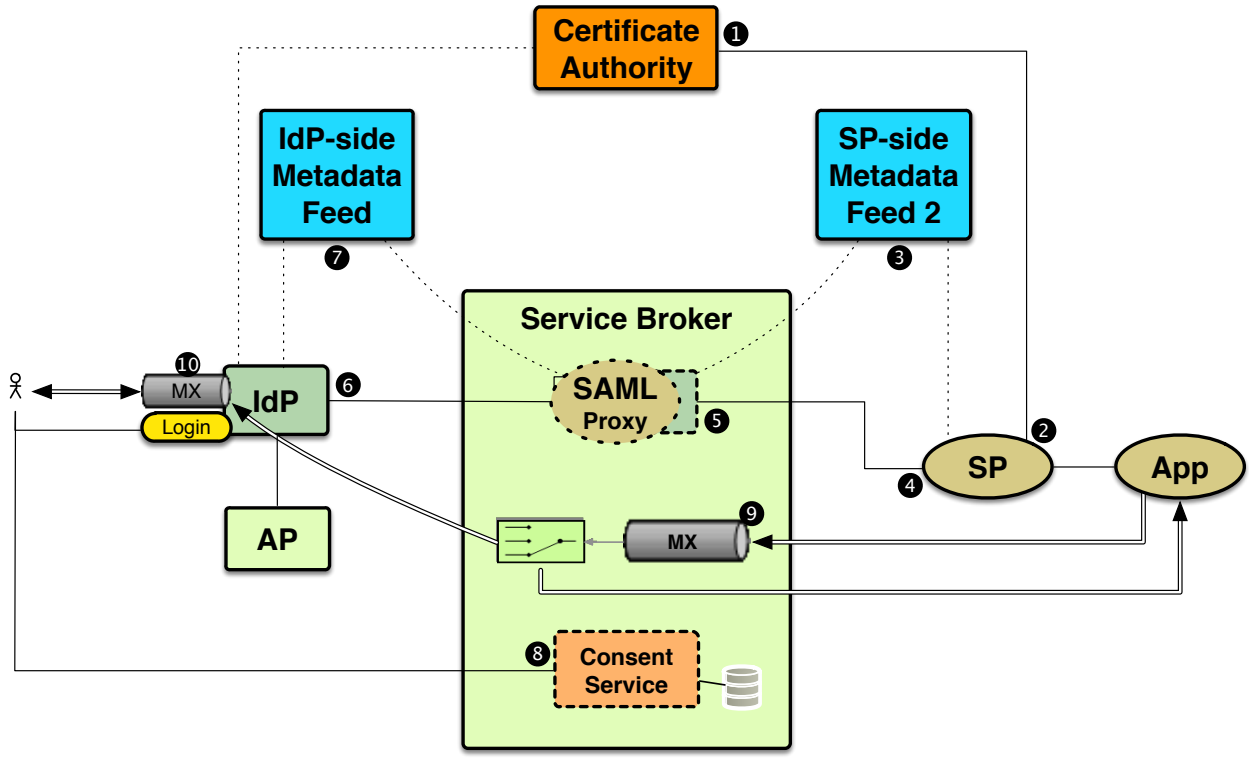

Fig. 1 High-level layout of the PEFIM model (WebSSO use case)

The following overview describes the interfaces that deviate from standard SAML WebSSO profiles, corresponding to the numbered references in Fig. 1.

1. The CA provides an interface for pseudonymous short-term certificates. An SP may obtain encryption certificates that assert that the SP is a federation member in good standing. Certificate serial numbers must be well randomized to diffuse any relationships to SPs that obtain certificates in blocks. CSRs may be authenticated using either a secure channel or signed messages.
2. An SP must implement the complementary interface to (1). Each authentication request must use a unique encryption key certified by the CA. For efficiency, signing multiple CSRs in batch-mode using CMS-signatures is recommended.
3. The metadata feed describes entities according to SAML2MetaIOP [2] with an SP-side view. That is, that IdPs are represented by proxies in the SB. `<EntityDescriptor>`

elements must not include encryption keys in the <SPSSODescriptor> element, because these keys are for one-time use only and therefore submitted in the authentication request.

4. An SP needs
    a. to send an <AuthnRequest> with a new one-time certificate in the <Extensions> element, as specified in 2.2 and
    b. understand the <Response> as specified in 2.3.
5. SB/IdP proxy for the SAML WebSSO profile. The SB proxies <AuthnRequest> and <Response> elements as follows:

	IdP	SB	SP
AuthnRequest		Rewrite AuthnRequest, filtering SP-identifying attributes (destination, audience and issuer)	Issue AuthnRequest to IdP proxy.
Response	Target Proxy-SP. Provide TID_1 in NameID. Encrypt attribute assertion.	Rewrite Response. Create TID_2 from TID_1.	Decrypt attributes.
Response (optional)	Aggregate encrypted attributes assertions	Pass thru	Decrypt attributes. Delete encryption key.

SB/IdP proxy for the SLO profile.

6. An IdP or AP using and validating the SP's encryption key contained in <pefim:SPCertEnc> MUST search the certificate using the public key and MUST NOT use the subject name or serial number. The encryption key MUST be verified using X.509 path validation.
7. Like SP-side metadata, but for IdP-side.
8. The SB implements a pseudonymous consent service. It allows users to grant, review and revoke consent. It may only operate on attribute names, not values to protect pseudonymity. Consent data is stored using TID_2 and SP-entityID as keys.
9. SB/Message Broker: MTA rewriting (a) TLD_2 to TLD_1 addresses or (b) TLD_2(SP1) to TLD_2(SP2) addresses and in the reverse direction.
10. IdP/Message Broker: MTA rewriting (a) TLD_1 to email addresses and in the reverse direction.

Note 1: This model requires SP-first authentication flows, as the IdP must not know about the principal's registered services. Service discovery may be implemented by an SB.
Note 2: The typical size of encryption certificates in PEM-format will be around 2k. Internet Explorer limits the URL length to 2083 characters, hence POST-binding is recommended to convey the AuthnRequest.
Note 3: Message content must not contain PII, because the SB or IDP could link this up with other data and violate the unobservability requirement. Solutions are end-to-end encryption or sending links to authenitcated contents.

2 Data Structures

2.1 Namespaces

Prefix	XML Namespace	Comments
ds:	http://www.w3.org/2000/09/xmldsig#	This namespace is defined in the XML Signature Syntax and Processing specification [XMLSig] and its governing schema [XMLSig-XSD].
pefim:	urn:net:eustix:names:tc:PEFIM:0.0:assertion	Namespace for elements introduced by this spec.
saml:	urn:oasis:names:tc:SAML:2.0:assertion	This is the SAML V2.0 assertion namespace, defined in a schema [SAML-XSD]. The prefix is generally elided in mentions of SAML assertion-related elements in text.
samlp:	urn:oasis:names:tc:SAML:2.0:protocol	This is the SAML V2.0 protocol namespace, defined in a schema [SAMLP-XSD]. The prefix is generally elided in mentions of XML protocol-related elements in text.

2.2 AuthnRequest

`<AuthnRequest>` elements must contain the encryption certificate used to encrypt the assertion with the attribute statement. The encryption key is represented within a `<ds:KeyInfo>` element. Its XPath is:
/samlp:AuthnRequest/samlp:Extensions/pefim:SPCertEnc/ds:KeyInfo/
ds:X509Data/ds:X509Certificate.

2.3 Response

A <Response> contains a single assertion that has following properties:

- It MUST have a subject, containing the Targeted ID in a `<NameID>` element;
- It MUST have an authentication statement;
- It MAY have (and usually will have) one or more `<EncryptedAssertion>` element contained in the `<advice>` element, each itself containing an attribute statement. Multiple `<EncryptedAssertion>` elements are useful to aggregate

attributes from multiple attribute providers without any other party than the SP reading them in clear.

- The encrypted assertion issued by the IDP must be self-contained with respect to XML namespaces[1].

An encrypted assertion MUST NOT contain a subject (because the TID_1 in nameID MUST NOT be revealed to the SP)

[1] The SB is creating a new <Response> element with its own QNAMES and does not know about namespace definitions in the encrypted part. Including all namespace definitions from the IDP or even having pre-defined QNAMES do not seem viable alternatives for the implementation.

2.4 Sample Instances

2.4.1 Authentication Request

(Non-normative – always implement according to specification – do not copy examples)

```
<samlp:AuthnRequest xmlns:samlp="urn:oasis:names:tc:SAML:2.0:protocol"
                    AssertionConsumerServiceURL="https://echo.kuk.portalverbund.at/SAML2/POST"
                    Destination="https://idp5.test.portalverbund.gv.at/idp/profile/SAML2/Redirect/SSO"
                    ID="_d11257d39d92042c860f5e8ee147a160" IssueInstant="2014-02-07T11:30:31Z"
                    ProtocolBinding="urn:oasis:names:tc:SAML:2.0:bindings:HTTP-POST" Version="2.0"
                    xmlns:ds="http://www.w3.org/2000/09/xmldsig#"
                    xmlns:pefim="urn:net:eustix:names:tc:PEFIM:0.0:assertion">
    <saml:Issuer xmlns:saml="urn:oasis:names:tc:SAML:2.0:assertion">
        https://echo.kuk.portalverbund.at/sp.xml
    </saml:Issuer>
    <samlp:Extensions>
        <pefim:SPCertEnc xmlns:pefim="urn:net:eustix:names:tc:PEFIM:0.0:assertion">
            <ds:KeyInfo>
                <ds:X509Data>
                    <ds:X509Certificate>
                        MIIC8jCCAlugAwIBAgIJAJHg2V5J31I8MA0GCSqGSIb3DQEBBQUAMFoxCzAJBgNV
                        BAYTAlNFMQ0wCwYDVQQHEwRVbWVhMRgwFgYDVQQKEw9VbWVhIFVuaXZlcnNpdHkx
                        EDAOBgNVBAsTB0lUIFVuaXQxEDAOBgNVBAMTB1Rlc3QgU1AwHhcNMDkxMDI2MTMz
                        MTE1WhcNMTAxMDI2MTMzMTE1WjBaMQswCQYDVQQGEwJTRTENMAsGA1UEBxMEVW1l
                        YTEYMBYGA1UEChMPVW1lYSBVbml2ZXJzaXR5MRAwDgYDVQQLEwdJVCBVbml0MRAw
                        DgYDVQQDEwdUZXN0IFNQMIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQDkJWP7
                        bwOxtH+E15VTaulNzVQ/0cSbM5G7abqeqSNSs0l0veHr6/ROgW96ZeQ57fzVy2MC
                        FiQRw2fzBs0n7leEmDJyVVtBTavYlhAVXDNa3stgvh43qCfLx+clUlOvtnsoMiiR
                        mo7qf0BoPKTj7c0uLKpDpEbAHQT4OF1HRYVxMwIDAQABo4G/MIG8MB0GA1UdDgQW
                        BBQ7RgbMJFDGRBu9o3tDQDuSoBy7JjCBjAYDVR0jBIGEMIGBgBQ7RgbMJFDGRBu9
                        o3tDQDuSoBy7JqFepFwwWjELMAkGA1UEBhMCU0UxDTALBgNVBAcTBFVtZWExGDAW
                        BgNVBAoTD1VtZWEgVW5pdmVyc2l0eTEQMA4GA1UECxMHSVQgVW5pdDEQMA4GA1UE
                        AxMHVGVzdCBTUIIJAJHg2V5J31I8MAwGA1UdEwQFMAMBAf8wDQYJKoZIhvcNAQEF
                        BQADgYEAMuRwwXRnsiyWzmRikpwinnhTmbooKm5TINPE7A7gSQ7l0RxioQePPhZO
                        zkM27NnHTrCe2rBVg0EGz7QTd1JIwLPvgoj4VTi/fSha/tXrYUaqc9AqU1kWI4WN
                        +vffBGQO9mo+6CffuFTZYeOhzP/2stAPwCTU4kxEoiy0KpZMANI=
                    </ds:X509Certificate>
                </ds:X509Data>
            </ds:KeyInfo>
        </pefim:SPCertEnc>
    </samlp:Extensions>
</samlp:AuthnRequest>
```

2.4.2 Response

t.b.d.

3 Glossary

Attribute Assertion

An <EncryptedAssertion> element containing an <AttributeStatement>.

Targeted Identifier (Targeted ID)

A persistent, non-reassigned, privacy-preserving identifier for a principal shared between a pair of IdPs and SPs. An IdP uses the appropriate value of this attribute when communicating with a particular SP (or SP affiliation), and does not reveal that value to any other service provider except in limited circumstances. Many similar definitions can be found for EduPersonTargetedID[2].
Synonym: Persistent ID
NOTE: This concept is extended for the PE-FIM model by decomposing the Targeted ID into TID_1 and TID_2.

Targeted Identifier 1 (TID_1)

A targeted ID between an IdP or AP and an SB.

Targeted Identifier 2 (TID_2)

A targeted ID between an SB and an SP.

[2] e.g. SWITCH AAI attributes: http://www.switch.ch/it/aai/support/documents/attributes/

References

[1] R. Hoerbe. (2014, A Model for Privacy-enhanced Federated Identity Management. *arXiv preprint arXiv:1401.4726*. Available: http://arxiv.org/abs/1401.4726

[2] OASIS, "SAML V2.0 Metadata Interoperability Profile Version 1.0," ed, 2009.

Change History

Version	Date	Changes
1.0	17. January 2014	Initial version
1.1	5. March 2015	Encrypted assertion issued by the IDP must be self-contained with respect to XML namespaces.
1.2	6. May 2015	Corrected type with samlp:Extensions being in Singular