University of Twente
Enschede - The Netherlands

Alfa & Ariss
Network Security Solutions

On A-Select and Federated Identity Management Systems

Joost Reede

August 4, 2007

Master's Thesis

Information Systems Chair
Computer Science Department
University of Twente

This thesis is supervised by:

dr. M.U. Reichert, University of Twente
A. Wombacher, University of Twente
R. de Vaal, Alfa & Ariss

Abstract

In the current web oriented information technology world, an increasingly large number of services take place via the web. These web services often require the user to identify himself in order to receive personalized information or services. The result is an increase in the amount of authentication procedures a user must deal with. One of the methods to cope with this, is using federated identity management.

Federated identity management provides the framework for the organizations to implement methods that reduce the frequency in which users must authenticate. Single sign-on is a prominent method, and is often used in federations. With single sign-on, a user needs only one authentication procedure, in order to be allowed access to resources from all federation members. One of the initiatives that allows for single sign-on in a federation environment is Alfa & Ariss' core product "A-Select", which is also the target system of this thesis.

Alfa & Ariss are keen on extending their collective knowledge on the subject of federated identity management and in particular on the effect of a broader adoption of federated identity management in A-Select. Gaining this knowledge is necessary to extend A-Select in such a way that it can cope with future challenges of working in federated environments.

The required information is mainly divided into two fields which are, in short, the basics of federations and the problems that are typically occur in federated identity management environments. I have been researching the topics by literature survey and working out a concrete scenario to pinpoint the actual issues discussed in the literature.

The basics of federated identity management are worked out in fine detail. The illustrated problems are illustrated extensively and where possible, solved up to a certain extend. In general, I am convinced that this document can be a source of knowledge that helps the development of A-Select a great deal.

Preface

This thesis is written to present findings of my research on federative identity management, conducted roughly from September 2006 until April 2007. In this period I have worked at Alfa & Ariss trying to find out what consequences the use of this specific form of identity management would inflict on their core product, A-Select. The A-Select project started in the mid-90's and is currently used for the authentication procedures of a large amount of web services, mostly provided by organizations in the Dutch public sector. The information gathered during my research project helps Alfa & Ariss develop A-Select further.

A lot of people have helped me in some way during the course of the project, for which I am very grateful. I will note a few of them. First of all I would like to thank my parents, sister, cat and all my friends for their mental support. Also, this thesis wouldn't be here without the help of my ever-jolly colleagues at Alfa & Ariss and their infinite A-Select wisdom. I especially would like to thank Remco for his in-depth reviewing of the thesis and Ali for giving me the opportunity to do this project. Last but not least, I thank the people at the University. In particular Andreas, for his devotion in supervising the process and providing me with invaluable feedback, and Manfred, who has the burden of assessing my work.

Contents

List of Figures

Chapter 1

Introduction

1.1 Background

In our current society where information means everything, web services play an increasingly important role in providing this information. For the user, to find his way around, he has to work his way through a jungle of web services. Also, more and more of these services require that the identity of the user is known, in order to present a personalized web page. The page may contain privacy sensitive information or simply a convenient selection of information that the user is interested in. Because these web services require knowledge of the user's identity, user credentials are needed to get access to the content. So, for each service the user wants to subscribe to, he needs to fill in a form that, most of the time, contains the same information. This can be very annoying for the user. Even more important is the fact that the user loses track of all the passwords associated with the services. This results in the user trying to find ways to manage his passwords. Since human beings do not have the mental capability to remember an endless amount of passwords [20], they use alternatives. Passwords are written down on pieces of paper, one password is used for multiple services, easy to remember (and therefore computationally weak) passwords are chosen, etc. All these alternatives have in common that they severely lower the level of safety that the identity of the user possesses.

In order to keep the identity of the user safe and still allow access to all the different services, the identity of the user must be unified instead of shattered among countless web services. A Federated Identity Management (FIM)

system provides the means to share the user's identity between two or more trusted parties. By sharing this information, the principle of ***Single sign-on*** (***SSO***) can be applied. A Single sign-on service only requires the user to authenticate once, but lets the user access an arbitrary number of identity dependent services.

Enabling communication between the services is not a straightforward process. If two services want to share their user databases it is very likely that the form the user information is stored in is not 1-to-1 compatible. A standard to transfer user information is required. One of these standards is developed by the OASIS group and is named Security Assertion Markup Language (SAML) [35]. It is the most widely adopted standard providing this connectivity.

The focus system in this thesis is A-Select, developed by Alfa & Ariss. Alfa & Ariss was founded in 1999 and specializes in developing authentication solutions. A-Select [2] is one of the core products of Alfa & Ariss, and can be used to build an authentication system. Its flexibility allows it to be a fundamental part in a federation. Therefore, a federation is also the application of A-Select that is the context of this thesis. A more elaborate technical survey on A-Select can be found in paragraph 2.6.1.

1.2 Motivation

In the near future A-Select will likely be used in more and larger Federated Identity Management systems. In order to be sure that A-Select is ready for such large scale systems, the consequences for Federated Identity Management systems in general and A-Select in particular need to be studied. Together with the company we have determined two main fields of which the knowledge level within the company should raise, and where more research should be conducted on: The basics of federated identity management and issues that apply to A-Select in the federated identity management context. These two research fields have been split into three goals, one for the former, two for the latter.

1.3 Goals

This thesis reports on how the main goal is reached. This main goal is to *Identify problems to occur in large federated identity management systems based on A-Select*. The main goal is split up into three sub goals:

- **Goal 1** - Identify the principles and foundations of Federated Identity Management;
- **Goal 2** - Identify potential issues of federations, with special attention to scalability;
- **Goal 3** - Find out which of the issues exist in A-Select and give solutions on how to solve them.

The process is divided into three phases. Each of the three sub goals is the main goal of the corresponding phase.

1.4 Approach

The first part of the thesis consists of a literature survey of which the results are used to form a conceptual model of a federation and to define sets of use cases and models (goal 1). In the second part, I take a look at what issues are presented in the literature. Since the goal is to find the issues of large federations, the issues are also examined with a significant expansion of scale in mind (goal 2). Finally, the issues found in the second part are addressed in the third. The findings are adapted to the A-Select scenario. This results in a series of recommendations and considerations that can be taken into account by Alfa & Ariss in the future (goal 3).

1.5 Thesis structure

The remainder of this document is organized to match the three sub goals mentioned above. Chapter two elaborates on the process to identify main principles of federated identity management and presents results thereof. It distincts several layers of abstraction within the general principles and discusses the fundamental pieces of each of those.

Chapter three illustrates some of the problems apparent in the field of federated identity management. These issues are distilled from literature and filtered for their importance to the company. Where applicable, it describes which of the core abstraction layers or concepts, mentioned in chapter two, are affected by the issue. If possible, a solution to the problem is also presented.

In chapter four, the issues discussed in chapter three are applied to a use case scenario based on a system of A-Select enabled components. In this use case, the principles as discussed in chapter two are made concrete. This enables for a good overview of A-Select functions within the context set by the principles discussed earlier. The second part of this chapter gives an indication of what issues of chapter three are applicable to A-Select and at what point in the scenario they become apparent. Also, possible solutions are given. If such a solution is not known or applicable, ideas are given on the general direction where a exact solution can be found.

The thesis is concluded in chapter five. In this chapter a summary of the problems is given and discussed. Also, the solutions from chapter four are summarized and casted into recommendations to the company. Finally, I present a list of open, but important points that still need to be covered, in a future work section.

Chapter 2

Federation fundamentals

This chapter elaborates on the concept of federations and federated identity management, with an overview of the common terms and principles in the field. Point of focus in this chapter is a conceptual model of federations. The model will clarify the basic concepts and the relations between them. Preceding the next chapter, the used federation models and use cases are described in detail at the end of this chapter.

2.1 Conceptual model

The conceptual model illustrates core components of a federated identity management system and how these concepts relate. This model provides the most abstract of overviews of a federation and forms the foundation on which the next sections are built. The model is constructed with the use of literature [19, 21, 34, 35, 43] and terminology used by the leading parties in the development of standards used for federated identity management [10, 11, 12].

The model has evolved from an incoherent and far too large collection of definitions to what it is now. The first step was to find similarities among standards and other reviews and eliminating specific non-fundamental elements. The set of concepts distilled from these standards appeared far too extensive to work with. Attempts to construct a conceptual model containing all concepts resulted in a model that was for too complicated. In order to avoid that, the set of concepts to work with was striped to the core concepts. Building a model out of this set gained the satisfactory result; A clear overview of the core of fed-

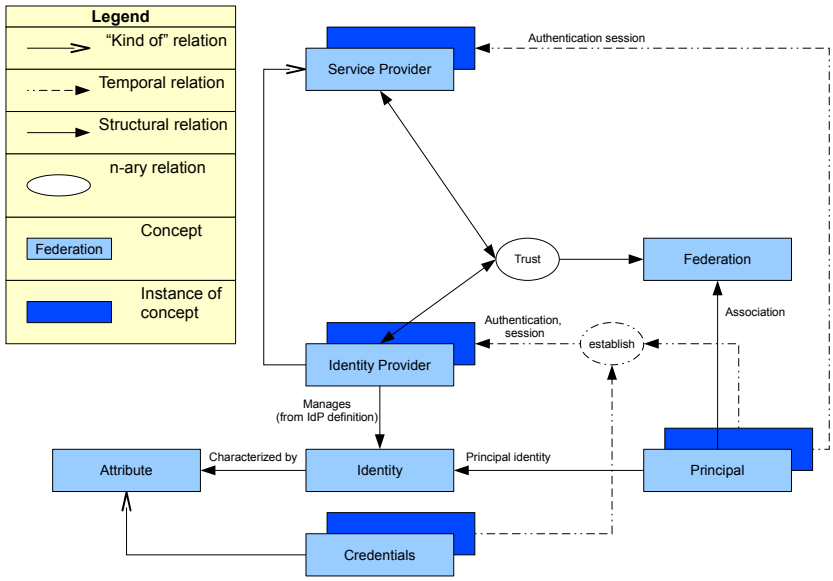

Figure 2.1: Federation conceptual model

erated identity management (figure 2.1). Another problem that occurred was the classification of certain entities that were neither clearly a concept nor a relation. One of these was the "credentials", that is ultimately defined as being a concept, but was also considered a relation between identity provider and principal.

A distinction must be made between the abstract concepts and their instances in the model, as can be seen in the model (represented by the light and darker boxes respectively). This same distinction is apparent regarding the structural and temporal part of the model, which is also present in the figure. Although both kind of relations are a substantial part of a federation, the structural relations (that also includes the kind-of relation) differ from the temporal relation in that they connect abstract concepts and are present as long as the federation exists. The temporal relations connect instances of concepts and exist only during the lifetime of a principal relation with the federation.

2.2 Concepts

Now we take a closer look at the separate concepts and determine their role within a federation. The terms that are displayed in ***italics*** can be found in the glossary at the end of the thesis. The definitions in the glossary are mostly extracted from the specific initiative glossaries [10, 11, 12].

2.2.1 Federation

The ***Federation*** forms the conceptual center of the model. Technically, it consists of a number of service providers that have mutual trust relations and principals that are associated to it. Concretely, this implies a network of service providers and identity providers that have a direct or indirect trust relation, depending on how the federation is modeled (see section 2.3) and the position of the provider in that model.

It must be noted that, besides the general definition of the word there is also the act of federation. This refers to the act where a party enters into a trust relation with a second party.

Related to that is the term ***Identity Federation***. It is used when a principal's identity, stored at an identity provider, is in some way associated with another identity of that principal, stored at a different identity provider. Methods to perform the act of identity federation are discussed extensively in section 2.5.

2.2.2 Principal

Principal is a general term for a party that benefits from being able to authenticate itself. This can be a tangible individual, also known as a user, or an abstract process such as a web service. A principal is said to be associated to a federation when it is part of a domain and has its identity stored at the domain's identity provider. Although a principal refers mostly to a concrete user, we keep referring to them as principals, since it is the most general term. In the rest of this thesis, when the term "user" is applied, this refers to a principal that can only be a real person.

2.2.3 Provider

The term "Provider" is used to annotate service providers and identity providers. The ***Service Provider*** (***SP***) is an entity in the federation that is able to provide

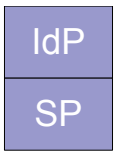

Figure 2.2: A composite domain, consisting of a service provider and an identity provider

services to requesting principals. The ***Identity Provider*** (***IdP***) is a kind of Service Provider. Identity providers store identity information of the principals and are able to verify identity information originating from other providers with the intention to authenticate a principal or federate a principal's remote with his local identity.

A clear distinction between the two must be made. In order to do that, an ***Administrative Domain*** providing functionality of both providers is modeled as having both, although it could well be that it is in fact only a service provider with a user database attached. This kind of domain is displayed as in figure 2.2.

2.2.4 Identity

The ***Identity*** is often mapped to a Principal (***Principal Identity***) and uniquely describes a principal by its characteristics. In the context of federated identity management, an identity is often a set of attributes that is stored at an identity provider. When multiple identities of a principal are stored at different identity providers, the act of identity federation can be conducted by ***Identity Mapping***, also known as ***Account Linkage***.

2.2.5 Attribute

An ***Attribute*** is a characteristic of the identity of a principal. Mainly, attributes are stored at the identity provider, in the form of a value that is identified by its key. This format and the key name are crucial in order to guarantee portability in the communication protocol and therefore compatibility among providers. An identifier is a special kind of attribute that, alone or in a set, identify an identity and thus a principal. A similar kind of attribute is the ***Credentials***. These are sent by the principal to prove its identity (the act of ***Authentication***)

to the identity provider and establish a trust relation called a ***Authentication Session***. Although credentials are identifiers, not all identifiers are credentials. Credentials do not have to be stored at the identity provider, but must allow the identity provider to authenticate a principal through some predetermined process.

2.3 Trust Relations

As discussed in the last section and as can be seen in the conceptual model, federations are formed using mutual trust relations between service and identity providers. Multiple providers can together form a trust network. The Liberty Alliance have identified what models can be applied when defining trust relations [28]. These give a good idea of how to classify trust models as used by the different initiatives. The models indicate how a small number (2 or 3) parties relate, but hybrid lay-outs can be formed to model more extensive trust networks. Liberty distinguish between business trust in the form of a contract between the parties and authentication trust that is enforced using cryptographic key certificates.

2.3.1 Direct and Indirect Trust

The concept of ***Trust*** is a classification of the relations between providers in a federation that indicates that one provider is willing to rely on the other providers for handling its principal identities. When a provider $\mathcal{B}$ has trust relations with both provider $\mathcal{A}$ and provider $\mathcal{C}$ and $\mathcal{A}$ and $\mathcal{C}$ do not have arranged a direct trust relation, then $\mathcal{A}$ and $\mathcal{C}$ are said to have an indirect trust relation, in which $\mathcal{B}$ is the intermediary (see figure 2.3).

2.3.2 Trust Models

Liberty classify models depending on the status of the business trust relation. Within a model the authentication trust relation can be either direct or indirect. Opposite to the business trust relation, the authentication trust relation must be present.

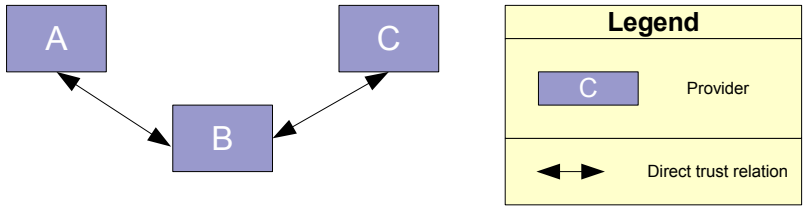

Figure 2.3: An indirect trust relation between A and C

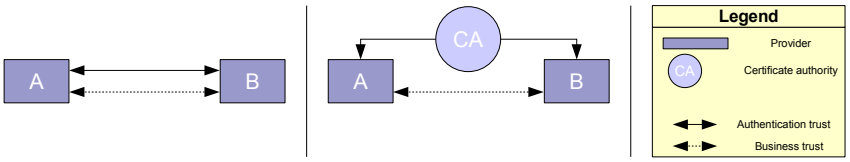

Figure 2.4: Pairwise business trust relation with direct (left) and indirect (right) authentication trust

Pairwise trust model

The pairwise trust model illustrates a direct business trust relation. This model ensures strong trust business wise, but is very limited when it comes to scaling. Although possible, authentication trust relations are not likely to be indirect, since new parties need to establish business agreements with federation members (see figure 2.4).

Brokered trust model

When business trust is established indirectly, the brokered trust model applies. Indirect trust is negotiated via an intermediary that acts as the trusted party on behalf of the local service provider. This provider can either decide to indicate what remote providers to trust or leave that decision up to the intermediary, thereby creating a more dynamic trust network. In both ways this model is more dynamic than a pairwise trust model, since far less business agreements need to be formed (see figure 2.5).

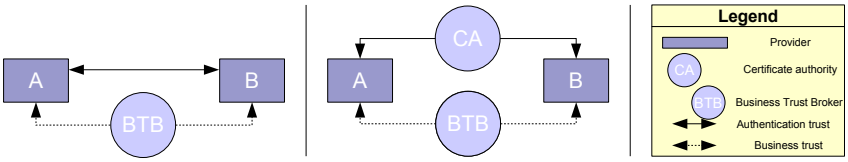

Figure 2.5: Brokered business trust relation with direct (left) and indirect (right) authentication trust

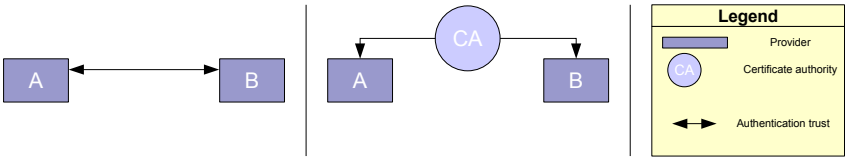

Figure 2.6: Community business trust relation with direct (left) and indirect (right) authentication trust

Community trust model

In the community trust model, there are no business agreements between providers, either direct or indirect. Instead of by business agreements, trust is established purely by technical means, such as by using trust infrastructures such as Public Key Infrastructure (PKI) (see figure 2.6).

2.4 Standards and evolvement

Now that the conceptual model and trust models are clear, we can take a look at the practical side of the matter; The standards that are present in federated identity management and how they evolved over time.

2.4.1 Introduction

Multiple parties in the field of identity management recognized the necessity of federating identities across multiple domains. The fact that this implied the use of standards resulted in the formation of consortiums that aimed at forming these standards: The Liberty Alliance created the Identity Federation Framework (ID-FF); Microsoft, IBM and others worked on the Web Services

(WS-*) specifications and the Organization for the Advancement of Structured Information Standards (OASIS, [32]) supervised the development of the Security Assertion Markup Language (SAML) standard. Another party involved with development of Federated Identity Management Systems is Shibboleth, which not only extended SAML, but also provides a complete implementation.

2.4.2 SAML

The Security Assertion Mark-up Language (SAML) [35] is a framework, based on eXtensible Markup Language (XML, [44]) that enables the exchange of user information. Supervised by OASIS, version 1.0 of the SAML specification was adopted as standard in November 2002. This version provided support for only basic functionality, such as Web Single Sign-on. Since then, SAML has evolved to a flexible and extensive framework. SAML version 2.0, which is the focus framework in this thesis, is built up out of several components that are designed in a way that allows a flexible framework, independent on the foundation on which it is used.

Assertions

An assertion is a piece of information that contains a statement about a subject. Usually, an assertion is sent from the ***Asserting Party*** to the ***Relying Party*** containing information of a subject of the asserting party. There are three kinds of statements that can be made using an assertion:

- *Authentication statements* indicate that is principal is successfully authenticated and at what time that happened,
- *Attribute statements* contain specific attributes of a principal and
- *Authorization decision statements* define actions a principal is allowed to do.

Assertions are composed using XML. SAML has a specific XML schema that defines the valid structure and contents of an assertion.

Protocols

The SAML protocols define what sequence must be followed in order to form a proper SAML conversation. There are several standard protocols:

- *Authentication request protocol* - A principal can request assertions carrying authentication statements and attribute statements;
- *Single logout protocol* - Protocol for activation of sign-off procedures at all providers;
- *Assertion query and request protocol* - Defines queries with which assertions can be requested, either by ID, subject or statement type;
- *Artifact resolution protocol* - Defines a mechanism for the use of artifacts. Artifacts are small pieces of data referencing to the actual protocol message. After receiving an artifact this protocol describes how to contact the original message creator for obtaining the message;
- *Name identifier management protocol* - Allows a requester to change the value or format of the name identifier or to end the association of a name identifier;
- *Name identifier mapping protocol* - Allows mapping of two name identifiers.

Again, SAML specifies a XML schema for all protocols.

Bindings

SAML bindings define how protocol messages can be carried using technology of the transport layer, such as HTTP. SAML v2.0 defines the following bindings:

- *HTTP Redirect binding* uses HTTP redirect messages,
- *HTTP POST binding* defines how to transport messages using POST and HTML forms,
- *HTTP Artifact binding* specifies the transport of artifacts over HTTP,
- *SAML SOAP binding* is used for defining how to carry messages using SOAP 1.1 over HTTP,
- *Reverse SOAP (PAOS) binding* specifies a multi-stage protocol that allows an ordinary HTTP client to communicate using SOAP and
- *SAML URI binding* defines how to retrieve an assertion using a universal resource locator (or URI).

Profiles

SAML profiles define a series of use cases and indicate what bindings, protocols and assertions are required to employ the use case. The SAML v2.0 profiles include:

- *Web browser SSO profile* defines single sign-on (SSO) with ordinary web browsers using the authentication request protocol, SAML response messages and assertions and HTTP redirect, HTTP POST and HTTP artifact bindings;
- *Enhanced Client and Proxy (ECP) profile* also defines SSO, for the use with specific clients using the PAOS and SOAP bindings;
- *Identity provider discovery profile* allows service providers to find out what identity providers a user has visited before;
- *Single logout profile* defines how to combine the Single logout protocol with SOAP, HTTP redirect, HTTP POST and HTTP Artifact bindings;
- *Assertion query/request profile* combines the Query and request protocol with synchronous bindings such as SOAP;
- *Artifact resolution profile* specifies how to combine the Artifact resolution protocol with synchronous bindings such as SOAP;
- *Name identifier management profile* defines the combination of the Name identifier management protocol with the SOAP, HTTP redirect, HTTP POST and HTTP artifact bindings;
- *Name identifier mapping profile* combines the Name identifier mapping protocol with a synchronous binding such as SOAP.

2.4.3 Liberty Alliance

The Liberty Alliance [26] is an organization involved in the development of standards for identity management. The goal of this consortium is to provide a set of standards to form a holistic specification to especially federated identity management. These standards are divided into three modules providing different functionality, namely ID-FF, ID-WSF and ID-SIS [27].

ID-FF

The Liberty Identity Federation Framework (ID-FF) is a set of protocols, schema and profiles that can be used to implement identity federation by supplying features such as account linkage and single sign-on. It is developed, based on the SAML v1.1 standard, with flexibility in mind and is therefore suitable for heterogeneous platforms and all kinds of systems, so it will integrate well with systems and specifications that are used. The single sign-off feature from ID-FF is used in SAML 2.0.

ID-WSF

The Liberty Identity Web Services Framework (ID-WSF) provides another set of protocols, schema and profiles for interoperability that is built on top op the ID-FF. ID-WSF can be used for discovery, consumption and creation of identity web services and provides features for enhanced and personalized identity services and attribute sharing.

ID-SIS

The Liberty Identity Service Interface Specification (ID-SIS) is built on top of the ID-WSF and contains a collection of specifications for services providing specific functionality, so Liberty enabled organizations are able to exchange these services. Examples of the possibilities are a calendar or a contact book. The Liberty Alliance have at this point released two specifications, the ID Personal profile (ID-SIS-PP) and the ID Employee profile (ID-SIS-EP).

2.4.4 WS-*

WS-* is the name of a collection of standards providing a framework for security for web services. It was originally developed by a group of corporations lead by Microsoft and IBM. It contains specifications for all kinds of security purposes, of which WS-Security and WS-Federation are of most interest to this document.

WS-Security

Web Services Security (WS-Security) specifies how security tokens can be attached to ordinary SOAP messages and how to add signature and encryption

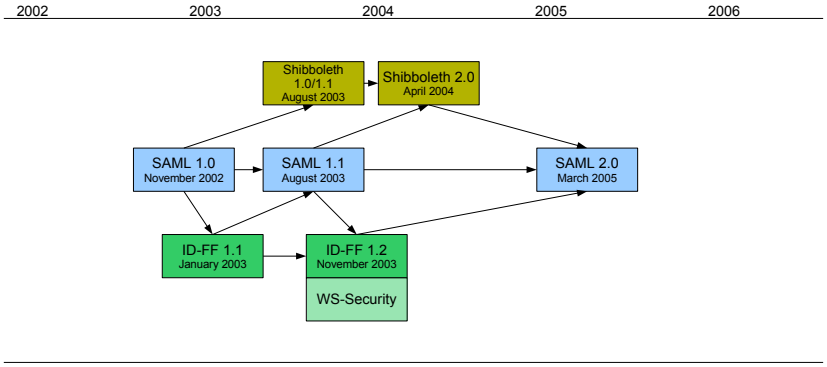

Figure 2.7: Time line illustrating the convergence of the different standards

headers to the message. WS-Security is used in the Liberty ID-FF 1.2 standards and is now a part of the OASIS specifications.

The WS-Security specifications provide a secure foundation on which SAML messages can be built. WS-Security can be used in combination with SAML assertions to embed security tokens and to ensure message integrity and confidentiality.

WS-Federation

WS-Federation specifications define how to share identity information in a cross-***Realm*** environment. It distincts two major profiles: The ***Active Requestor*** profile uses a fully SOAP compliant client and ***Passive Requestor*** profile only needs a HTTP compliant client such as any modern browser. WS-Federation invokes WS-Trust, which is also part of WS-*, for specification of security token exchange and issuance.

2.4.5 Time line

In the early days of federated identity management, there were multiple initiatives available to accommodate such a system. Every initiative focused on a different section of the market. SAML (v1.0, v1.1) aimed at the business-to-business world, Liberty focused on the business-to-customer market and Shibboleth (see paragraph 2.6.2) was most active in the educational sector. Since provider interaction is a key-factor in a federation, a single standard is

preferred over a system working with providers that use different standards and need to use adapting middleware to be able to interoperate. With this knowledge in mind, OASIS, Liberty Alliance and Shibboleth started working on SAML 2.0, that would converge their standards into a generally supported one.

SAML version 2.0 is based on version 1.1 and is compiled with a great deal of participation from the Shibboleth project and the Liberty Alliance's ID-FF 1.2 (see fig. 2.7).

2.5 Identity Federation Use Cases

In this section we elaborate on the identity federation use cases that are most important to the rest of this document. These scenarios are part of the SAML 2.0 standard mentioned above and are described in the SAML technical overview [35]. All use cases assume a user want to access a resource at a service provider that is not his home identity provider and he decides to federate his local identity with the one of his home identity provider. In the case of a persistent federation, the user needs to authenticate only once to access data on both providers, after the identity federation process. Federation using transient pseudonyms has other advantages, as is discussed later.

2.5.1 Out-of-band

Out-of-band identity federation uses a procedure and communication procedure that is not part of the SAML standard. Therefore, it was the only kind of identity federation possible in SAML 1.0 and 1.1. Providers that federate their stored identities must agree on the method used for federating and the mapping of the identities. The method can be either in batch form or a real time synchronization.

2.5.2 Persistent Pseudonym

After a successful authentication at the identity provider, the service provider receives an authentication assertion. Added to that assertion is a reference to the user's identity at the identity provider (the pseudonym), that is only linked to the user profile at the identity provider, thereby not disclosing the identifier of the user. The user is asked whether or not the identities must be federated

and if so, a record is stored at the service provider that holds a reference to the local identity, the identity provider and the pseudonym gained from the identity provider. The identity provider creates the same kind of record. Using this information, the trust relation between identity provider and service provider enables the user to use both identities with a single log in procedure.

2.5.3 Transient Pseudonym

The transient pseudonym is similar to the persistent pseudonym, but differs in the fact that the pseudonym provided by the identity provider is not in any way linked to an identity of the service provider. The implications of that are that the user does not have to authenticate at the service provider and therefore does not have to identify himself. The user remains anonymous to the service provider and only uses the trust relation between the service and identity provider and authentication at the latter in order to gain access to information at the service provider. Using the transient pseudonyms, it is possible to have a service provider that does not have any identities stored, saving a lot of maintenance load.

2.6 Initiatives

The trust models, standards and use cases discussed above form an abstract foundation on which federated identity management is built. What remains are concrete implementations of these principles. There are several initiatives providing implementations of federated identity management systems, of which we will discuss a few. The focus is on the SURFnet federation and Shibboleth, of which the former is the perspective from which the following chapters are defined. Although the analysis of the issues (chapter 3) is broadly oriented, the discussion of the practical application of the issues applies only to the SURFnet.

2.6.1 SURFnet Federation

The SURF foundation [40] manages the SURFnet [41], a technologically advanced physical network connecting Dutch educational organizations mutually and to the Internet. The foundation also develops all kinds of services to customers to optimize the use of the network. One of these services is the SURFnet Federation (SNF). The SNF is the most common form of federation

that A-Select makes possible. A-Select facilitates authentication mechanisms in the SURFnet, thereby providing it with federation capabilities. The A-Select server must be configured to run in so-called "cross A-Select" mode. A cross A-Select set-up needs some elaboration.

First of all, there are some different naming conventions used in cross A-Select. The use cases described in the SAML documentation and above use the terms "service provider" for the provider where the principal wants to access a resource and that needs an authentication assertion, that in A-Select is called a "Ticket Granting Ticket" (TGT). Cross A-Select uses the term "Local A-Select server" for that entity. The identity provider from SAML, that is used to authenticate the user, is named the "Remote A-Select server".

Cross A-Select can be used for remote authentication, in the same way the web browser/single sign on SAML profile functions. The local A-Select server defines all remote servers and imports their public key certificates in the local keystore. The remote A-Select server defines the local servers in the configuration and imports their respective public key certificates.

Cross A-Select can also be used as a proxy between a group of local and a group of remote servers. The certificates and administration is managed by the proxy on behalf of the local and remote servers. To the local servers, the proxy acts as a remote server and to the remote servers it is a local server. This centralizes the management of server administration, which makes management less complicated.

Supported use cases

At the moment of writing, there is no support for active user account linkage. This implies that identity federation using persistent pseudonyms is not possible. Out-of-band identity federation would theoretically be possible, although there is no system that provides such functionality. On the other hand, the use of the transient pseudonym is not only possible, but is also the main use of the SURFnet Federation.

Supported trust relations

Trust in a cross A-Select environment is established with the use of Public Key Infrastructure (PKI). Plain cross A-Select relies on direct trust only. As indicated, certificates need to be exchanged mutually, thereby forming a direct trust

relation between two A-Select servers. Cross A-Select in combination with a proxy, in comparison, is a good example of indirect trust. Servers exchange their certificates with the proxy that becomes the authority which the server trust. A local server that requests an authentication assertion from a remote server, now requests this assertion via the proxy. Since the proxy trusts the remote server and the local server trusts the proxy, an indirect trust relation exists.

In the case of SURFnet, business trust is established only via brokered trust. New parties only form agreements with the SURF Foundation that manages the so-called Root Identity Provider (RootIdP), that functions as the intermediary between two providers that want to exchange identity information.

Incorporated standards

Communication in A-Select is based on HTTP with URL encoding. An alternative is to use Simple Object Access Protocol (SOAP, [38]), but this is not conventional since this channel is not fully SOAP compliant. This also applies to compatibility with the SAML standard. Although a pure A-Select system uses its own communication protocol, there is an adapting module available to enable the use of the SAML 1.1 protocol, especially for communicating with Shibboleth-based networks.

Supported use cases

At the moment of writing, there is no support for active user account linkage. This implies that identity federation using persistent pseudonyms is not possible. Out-of-band identity federation would theoretically be possible, although there is no system that provides such functionality. On the other hand, the use of the transient pseudonym is not only possible, but is also the main use of the SURFnet Federation.

2.6.2 Shibboleth

The Shibboleth Project [37] is an initiative by the Internet2 community [18] aiming at providing a Single Sign-on and attribute exchange solution for use in the Higher Education sector. Shibboleth is based on the OpenSAML [33] implementation which is itself based on SAML v1.1. Shibboleth's federated

identity system framework's most distinctive component is the Where Are You From (WAYF) server that is used to determine the home identity provider of the user that needs to be authenticated or of which some specific attributes need to be retrieved.

As is the case in A-Select, naming conventions differ from the Liberty/SAML standard. Although many documents on Shibboleth use Liberty/SAML vocabulary, there originally were other terms. The service provider from the SAML example, discussed in section 2.6.1, is named the "Target". The identity provider is often referred to as the "Origin".

Supported use cases

Since SAML 1.1 is the foundations underneath Shibboleth, and SAML 1.1 does not support identity federation within the specification, this also applies to Shibboleth. As is with SAML 1.1, out-of-band identity federation is theoretically possible, using some external communication application.

Supported trust models

The Shibboleth project did not define any trust models. This is because Shibboleth is a direct descendant of SAML 1.1 and therefore supports SAML's trust models, that were originally defined by the Liberty Alliance.

Incorporated standards

Shibboleth is built on top of SAML 1.1, with extensions added to it. The attribute profile that was originally part of Shibboleth, is adopted by the OASIS and is now part of SAML version 2.0.

Chapter 3

Federation issues

3.1 Introduction

In the previous chapter the basic concepts, trust models, use cases and standards for federated identity management are illustrated. This chapter discusses common problems that occur in the world of federated identity management. The rest of this chapter is organized as follows: First, the methods that are used to identify the issues are discussed and the boundaries to which the contexts of the issues are confined are defined. Next, the issues are elaborated on in section 3.2 and the chapter is closed with conclusions that are drawn from the issue discussions in section 3.3.

3.1.1 Methodology

The issues discussed in this chapter are extracted from a wide array of literature on identity management. Some of these issues are related directly to federated identity management. For others, the original context was identity management in general and these need some more elaboration and reflection on identity federation management, in order to make the transition to that context. The result of the literature survey was an extensive list of issues. A filter was applied to the list in order to filter out any issues that are not relevant for the topic. The remaining issues were categorized. A second round of filtering was applied, in which the company indicated the relevance they assigned to each issue. The issues in the categories were then merged to form a single issue per category, resulting in the list of issues discussed below. These issues

therefore form a good overview of the problems that are both relevant to the topic and of interest to the company.

3.1.2 Boundaries

In the last section a filtering process was mentioned. Issues mainly were filtered out because their context was not within the boundaries of this research. The main characteristics for every issue regarded important for the thesis are:

- SAML - If an issue is dependent on the information transport layer of the network, it must be concerning SAML, either 1.x or 2.0. However, solutions for the issues can originate from any standard since it is the general idea that counts, not the implementation;

- Zero footprint - If an issue relates to the interaction between principal and network, only naive clients are taken into account, most likely ordinary web browsers (also known as passive requestors, in WS-*);

- Issue not applicable - As indicated earlier, many issues originate from the field of identity management. Some of these issues are also applicable to federated identity management. However, some issues cannot be adapted to the federated identity management context and these are therefore skipped.

Now that the boundaries are determined, it is clear in what context the issues must be placed.

3.2 Issues

3.2.1 Issue 1: User control

Any system that is developed to be used by an arbitrary set of users should encourage people to use it. For identity management systems, this encouragement should come from establishing trust in the system among the user group. People tend to trust the system when they feel that they are in control of the information managed by it [4]. Therefore, an identity management system should have user control as one of the core requirements.

Problem

Reflecting this issue onto federated identity management systems, there are some specific differences compared to regular identity management. First of all, there is the distinction between federations that enforce user privacy rules by policy, such as Liberty enabled systems, and federations that do not [30]. Next, it is important what environment the federation is in. User control is much easier to apply in a homogeneous environment than in a heterogeneous one.

Known solutions

Applying user control policies should guarantee that each provider in the federation puts users in control. Especially in a heterogeneous federation, the enforcement of a policy is difficult to maintain and verify, so extensive testing of new parties is necessary. The paper of GailJoon Ahn and John Lam [1] indicates four key concepts that need to be taken into account, when devising policies regarding user control and consent:

- *Notice* - Users should receive prior notice of the information practices;
- *Choice* - Users have a choice to specify what information will be used and the purpose for which the information is collected;
- *Access* - Users should be able to access and modify their personal information if necessary and when needed;
- *Security* - Users should be assured that the organizational system is capable of securing their personal information.

Also, the paper identifies different scenarios for notifying and communicating with users about their preferences regarding their attributes. These give a good all round plan for ensuring user control and could therefore serve as a good foundation in federation policies on that topic.

3.2.2 Issue 2: Identity provider discovery

Web services that allow delegated authentication, as is the case with service providers in federated identity management, must be able to guide a user to his home identity provider, thereby knowing to what identity provider the user

is directed to and what is the nature of the identity provider. This knowledge is crucial to the web service in order to perform the right steps for communication and trust negotiation.

Problem

In a static federation where every identity provider is known by default, a list of possible principal's identity providers is always present and can always be consulted. If, however, this is not the case and the federation is of a dynamic nature and new identity providers are arbitrarily connected, sending the principal to his home identity provider is not so straightforward. The list that the service provider can present to the principal is most likely to be incomplete. The users of the new identity provider are not able to choose their home identity provider, and therefore cannot be authenticated. A failed authentication results in the principal not being able to access the resource of the service provider.

Known solutions

This problem is recognized as being notoriously difficult to solve [19]. SAML v2.0 defines a special profile for identity provider discovery. This profile uses a domain cookie to store preferred identity providers of the principal. There are two drawbacks to this approach. First of all, it requires the use of cookies that can be read by identity provider that is not the original issuer of the cookie information. Usually, this is prohibited by default in any modern user agent, and thus requires the user to lower the security level of the agent [24]. This is not recommended and should therefore not be a necessary step in the identity provider discovery solution. Without lowering the security level, this approach can only be used within one administrative domain, and this is impossible in larger federations. A second drawback is that the cookie only stores preferred identity providers. If a user has not defined his preferred identity provider yet, then it will not be known to the service provider and cannot be selected by the user.

3.2.3 Issue 3: Minimal information disclosure

In every identity management system, there is always the possibility of identity theft. Of course, the chances of such an event occurring must be reduced to the

minimum. Also, one must not forget to reduce the impact of an identity theft event as well. Reducing impact can be accomplished by storing the absolute minimum of the required information. This implies not only to store the least amount of information, but also only for the time it is needed and no longer. For example, a service provider that hosts information for different age groups, does not need to know the birthday of users. Storing only the age category is sufficient in this case [4].

Problem

Identity theft prevention in a stand-alone identity management system is a difficult problem, and this problem is even more severe in the case of federated identity management. In federated identity management, many identity providers can store identity information of one principal. So, identity theft can occur via multiple channels. The impact is even higher when identities of the principal are federated. In the worst case, the trespasser can access information on all identity providers the principal has stored information on and has federated identities with. In general, federating identities tends to enlarge both the chances and impact of identity theft, unless serious measures are taken. Note that, although important, the focus is not on the technical point of view of reducing the chances of identity theft. We assume that identity theft is an inevitable event that has a certain chance of occurrence at a single point in the federation. We focus on reducing the chance and impact of identity theft in a federation as a whole.

Known solutions

There are several measures that can be taken in order to reduce chance and impact of identity theft. SAML v2.0 introduces a new dynamic way of federating identities by using transient pseudonyms, which guarantees anonymity of principals. Identities are linked only temporally and without mutual knowledge of the other party's identity information (see also paragraph 2.5.3). This reduces the impact of an identity theft event. Besides this technical solution, there is also the necessity of defining policies. Policies should define and clarify how identity providers should cope with identity information stored locally or received from remote identity providers. Because of the open and dynamic nature of a federation, the policy system used must also be dynamic

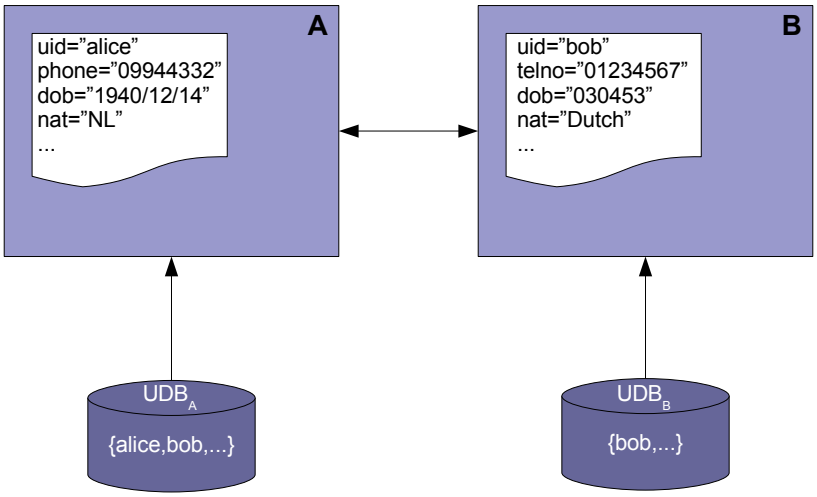

Figure 3.1: Two identity providers with attribute format discrepancy

and provide a fine granularity. The approach of Bertino, Bhargav-Spantzel and Squicciarini [3] defines five categories of policies, to be implemented by identity providers and principals that cover the full spectrum of information flow within a federation. The authors also propose an assertion language that can be used to define flexible and dynamic policies, necessary for information management in federations. A different approach is that of Cassasa Mont, Pearson and Bramhall [30]. This approach allows for very fine granularity policy definition, since policies can be attached to each individual message. Although the approach is focused on regular principal-service provider communication, it is applicable to federation systems as well.

3.2.4 Issue 4: Enforcement of attribute mappings

A user's identity that is stored at an identity provider is composed of attributes. These attributes hold a single property of a principal. Usually, these attributes have a key=value format. In the case a principal wants his identities federated, this can cause problems [13, 14, 15, 36, 39, 42].

Problem

Suppose user "bob" has an identity at his identity provider $\mathcal{B}$, containing all kinds of identity related information. He also has a lightweight account at identity provider $\mathcal{A}$, consisting of some other arbitrary attributes, and wants to federate his identities, so $\mathcal{A}$ can use his identity information in a session directly. As indicated by identity information of "alice" in figure 3.1, $\mathcal{A}$ is aware of the same identity attributes as $\mathcal{B}$, but the information differs in format. This implies that $\mathcal{A}$ cannot directly access information of $\mathcal{B}$, because of discrepancy in the attribute mapping. In the field of databases, this problem is well known, and is often referred to as the data integration problem [25]. Extensive research has been performed in this field as well. However, within the context of this thesis the focus will be on attribute mapping problems and their solutions. Different kinds of mapping problems are([14]):

Attribute name differs
Often the mapping fails on retrieving a certain attribute, because this attribute is not known by the name the requesting identity provider uses. This is illustrated by the "phone" attribute of $\mathcal{A}$. This is in fact the same attribute as the "telno" attribute of $\mathcal{B}$, but has a different name. This fact complicates the phone number attribute query of $\mathcal{A}$ to $\mathcal{B}$ and vice versa.

Attribute syntax differs
Another problem lies in the fact that the names are the same, but the syntax of the values is not. This can be seen in the example, regarding the "dob" attribute. This attribute, representing the date of birth of the principal, has different notations at the different identity providers and the attribute value must therefore be processed before it can be used by $\mathcal{B}$.

Attribute semantics differ
Similar to the problem of syntactical differences, is the problem of semantical differences. In contrast, the values of the attributes are semantically different. This is illustrated by the "nat" attribute, that indicates the nationality of the principal. Although the values of the attribute represent the same country, the values can not be translated without a list of countries' nationality names (for "Dutch") and codes (for "NL") to translate the values.

Composed attribute differences
The mapping problems described above illustrate what issues can arise when a 1-on-1 mapping is possible. However, there is also the possibility that an attribute is split up at one identity provider and is composite at the other. For example, the birthday attribute could also be split up into separate attributes for day, month and year of birth.

Known solutions

It is clear that this issue puts quite a severe strain on federated identity management, considering the amount of times this is referenced in the literature. Often it is only referenced to as apparent, without presentation of a solution. However, there are some papers that present solutions varying in level of detail.

Attribute layout transformation [14]
The first solution is to create an adapting interface between providers that is capable of translating attributes from any and to any schema. Hommel and Reiser [14] propose such a solution. It involves an attribute converter, present at any identity provider. The attribute schema need not be changed, since all conversions take place in the stand-alone converter. Conversions are carried out using XSL Transformations [6], that is useful because of its flexibility and powerful XML transformation possibilities. The last addition proposed in the paper is the centrally placed Federation Schema Correlation Service, that is in charge of conversion rule distribution and is therefore capable of providing every attribute provider with appropriate rules for transforming their attribute assertions into the requested form for any other party to communicate with.

Standard schema [13]
A different approach is the use of standard schema throughout the federation [13]. The Liberty Alliance have defined two profiles that can be applied when developing providers that need or distribute personal (ID-SIS-PP, [23]) or employee (ID-SIS-EP, [22]) information. Profiles are in development for contact book, geolocation and principal presence services. Besides using pre-defined schema such as in the ID-SIS profiles, it is also possible to define schema for a federation. All providers should comply to this schema and therefore it is vital to also define proper and clear policies for existing and new parties to

maintain.

Combined solution
It is possible that the two approaches that are discussed are combined into one. For example, it is hard to match semantically different attributes, since additional information is needed for that. It could well be, that a clear definition of such attributes is stated in the policies and the decision on how to format the other attributes is left to the individual providers in order to interfere with their internal processes as little as possible. However, the border between technical and legal enforcement is a thin and vague one and should be properly defined.

3.2.5 Issue 5: Provider identity and trust

As already indicated in the discussion on User Control, the trust of the user in the system is vitally important. Trusting a system means trusting the individual components, which in this case are the providers, user clients and the connections between them. Therefore, the providers need to properly identify themselves to users to gain their trust and the providers need to establish mutual business trust and authentication trust relations with each other in order to enhance trust of the users in the network as well [36]. In this case, identity and trust are intertwined concepts. Determining the true identity of the other component is a significant part in the process of establishing a trust relation between two components. If a trust relation exists, this implicitly means that the identities of the trust relation participants are known.

Problem

Concretely, the problem is how to distribute trust in a federated identity management system. Especially in a large federation consisting of many indirect trust relations, it may not be desirable to add another provider using only manual certificate exchanges for authentication trust distribution. Although it is possible to exchange certificates with only one provider and rely on indirect trust for secure communication with the rest of the providers, this generates an increase in overhead interaction, when the network grows. Besides overhead growth, the communication also becomes more fragile when depending on an increasing amount of parties. Both the safety and the reliability of the communication channel suffer from this increase.

Known solutions

For providers to be able to verify the identity of other providers, and establish trust relations with them, an administration is necessary. In order to keep this information available for all providers, it needs to be a central administration. Such an approach is presented in [9]. This paper introduces the concept of a notary server in federated identity management. This central authority is able to distribute trust to other federation members.

The notary server authority allows only identity providers of good reputation. Identity providers can rely on the fact that registered identity providers are to be trusted.

3.2.6 Issue 6: Single sign-off

One of the most prominent uses of federated identity management is the Single sign-on feature. Specifications for that use case are well evolved and extensively discussed in many papers and technical documents. A less popular subject, that is in fact equally important, is ***Single Sign-off***. The process of Single sign-off has the exact opposite result of the Single sign-on process, and causes all current sessions the principal has with the different identity and service providers to be closed.

The single sign-on feature is managed within a single sign-on session. This session is responsible for the distribution of the authentication assertions as soon as a requesting party requires it. This is where the single sign-on session differs from the ordinary session, that can only maintain authentication for and grant access to a single web service.

Problem

From a federated identity management perspective, a complicated problem appears. Independent of whether the single sign-off process is initialized by the service provider or the identity provider, the identity provider is responsible for sending sign off messages to all service providers. In the case of a service provider initialized single sign-off process, the service provider delegates the responsibility of fulfilling the task to the identity provider. Although the process for performing a single sign out action is well defined in the Liberty profile [5] (later adopted in SAML [16]), the profile does not indicate how

to maintain the necessary administration of involved service providers. This administration is necessary for determining what sessions need to be closed. Sessions are closed by sending a logout request from identity provider to service provider using `HTTP GET` messages in combination with redirects via the principal, or directly using SOAP messages.

Known solutions

WS-Federation [21] is the only of the major specifications that mentions the administration inconveniences. It proposes the use of a loose framework, based on messaging as defined in WS-Eventing [7]. WS-Eventing defines a messages protocol that can be used to subscribe to an information service. The information service uses the resulting administration to send information messages to all subscribers. The federated sign-off functionality as described in the WS-Federation specification indicates the use of two different message flows to accomplish global sign out in a federation. The sequential method relies on every individual provider that maintains a local session of the principal to pass the single sign-off message through to the next provider. The last provider then sends a message back to the initializing identity provider. This method is considered very fragile; If one link in the chain fails to re-send the sign-off message, the rest of the chain is left unnoticed. A second approach is to notify the individual providers in parallel. This approach is much more stable and is therefore the mandatory implementation in systems that comply to WS-Federation. In order to receive sign out messages, the providers must subscribe to the originating identity provider. Upon receiving a sign-off message, the provider is indicated to clear its session information. Although this sequence provides better results, the protocol does not provide a fully predictable outcome. As said, the protocol is quite loose, and relies on one-way sign-off messages, requiring no reply. Furthermore, the sign-off messages act as a hint and providers are not obliged to clean up the session.

3.3 Conclusion

Although the choice for using federated identity management can ease identity management administration, there are some serious issues to consider. Some of those issues relate directly to “plain” identity management and intensify in the

federated identity environment. Others are new and originate from federated identity management. Although this list of issues is not complete, it presents a good overview of issues that may occur in any federated identity management system.

Chapter 4

Application to A-Select

4.1 Introduction

In the previous chapter, some issues are discussed that can arise in a federated identity management system. Now that these issues are identified and, if possible, accompanied by a solution, the next step is to see whether these issues are also apparent in the target system. The target system we are interested in is the SURFnet Federation.

This chapter is organized as follows: First, the methodology that is used to find the issues is explained in section 4.2. In section 4.3, a complete scenario is described that will be used to search for and pinpoint the problems. Section 4.4 handles the issues and elaborates on every issue individually in order to establish its presence and its place in the model.

4.2 Methodology

The issues that are discussed in the last chapter vary in the level of abstraction of the federated identity model they can occur. Therefore, to find out if and where these issues occur within the SURFnet Federation, all abstract levels (mostly discussed in chapter 2) need to have a concrete example to work with. Next, a concrete example of a simple federation is presented that consists of the most basic components. Of every issue, the place in the example is identified so it is possible to give a good indication of the implications the issue carries for the SURFnet federation. In the next chapter, the conclusions of

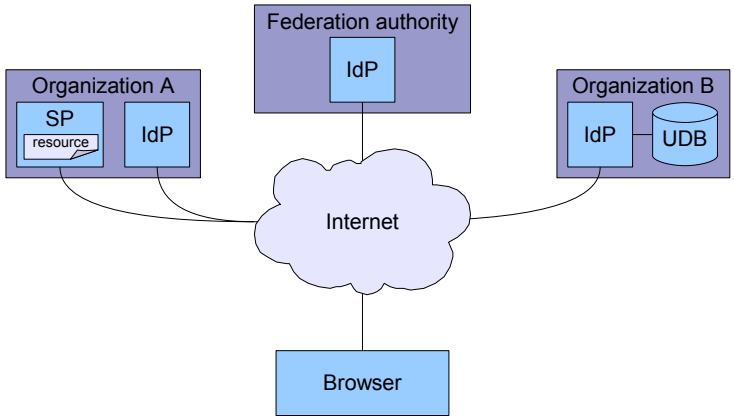

Figure 4.1: Architecture used in the example

the analyses are discussed and recommendations are drawn.

4.3 The scenario

4.3.1 Situation sketch

Suppose a principal, Bob, is trying to access a resource that is hosted and maintained by Organization A. Organization A is connected to the SURFnet Federation, and so is Organization B. Organization B is the home identity provider of our principal Bob. The SURFnet Federation also has a root identity provider that is known and trusted by all organizations in the federation. It could well be that there are more organizations connected to the SURFnet federation. These are not important in the example and are therefore not discussed. Let us take a look at the technical specifications of the individual components (see figure 4.1):

Organization A

Organization A is an administrative domain that hosts a service provider, containing the resource, and an identity provider.

Service Provider

The service provider contains resources, among which is the resource Bob wants to access. The resources are supplied only to those who are able to present a

valid authentication assertion. Validity of the assertion is guaranteed if the assertion is given out by a trusted party, in this case organization A or Organization B.

The service provider can be any web service platform. This access control is managed by a component called the A-Select filter. This filter negotiates with the local identity provider in order to determine whether a certain user is allowed to access the requested resource. The A-Select filter can be seen as an extension of the identity provider and will not be discussed further.

Identity Provider

The identity provider is an A-Select server set to communicate with the root identity provider, using the "Cross A-Select" mode. This enables the Identity Provider to act as a service provider requesting authentication assertions, thereby delegating the act of user authentication to the root identity provider (see also the extensive description in paragraph 2.6.1). Cross mode also implies this communication mode is bi-directional, thereby letting the root identity provider act as a service provider requesting authentication assertions from the local identity provider. The local identity provider can have a local user database, but this is not necessarily the case in this example. Since this would not be used, the user database is omitted in this example. As we will see later, the identity provider of organization A does not have a substantial role in the scenario, since the example principal is not part of organization A.

This identity provider is known in the SURFnet context as a client-only identity provider. A client-only identity provider cannot authenticate users itself, but is able to send authentication requests and receive authentication assertions.

Organization B

Organization B is an administrative domain and is the home organization of Bob. Organization B can also host a service provider although that is not the case in this example.

Identity Provider

The identity provider of organization B is the actual home identity provider of the principal. In the context of the SURFnet federation, this identity provider is required to be a "full" identity provider. A full identity provider can serve

as the asserting party (remote identity provider in A-Select terms) for other identity providers to receive authentication assertions from.

User database
The user database is an A-Select component that forms a meta database containing specific A-Select properties and linking to accounts in the backend systems that are used to actually authenticate the user. Multiple backends are supported (e.g. LDAP, RADIUS or a plain relational database). The backend is not regarded in this example. We assume that all information necessary can be consulted by the identity provider, regardless of where that information originates from.

Federation authority

The federation authority is the administrative center of the federation. Possible new members subscribe at the authority and are audited before being allowed admittance to the federation.

Identity Provider
The identity provider at the federation authority, which is also referred to as the root identity provider, serves as a proxy for authentication requests in the same way as the Where Are You From server operates in the Shibboleth situation. Requesting identity providers, the relying parties, query the root identity provider. The authentication request is then sent through to the user's home identity provider, the asserting party, that is capable of authenticating the principal.

Browser

The principal in the scenario uses an ordinary modern HTTP compliant client, from now referred to as the browser, for executing the steps to perform the communication with the different providers.

Disregarded parts

In the example, only the most trivial of parts are taken into account. Components that are not used are omitted for the sake of simplicity. Such parts include:

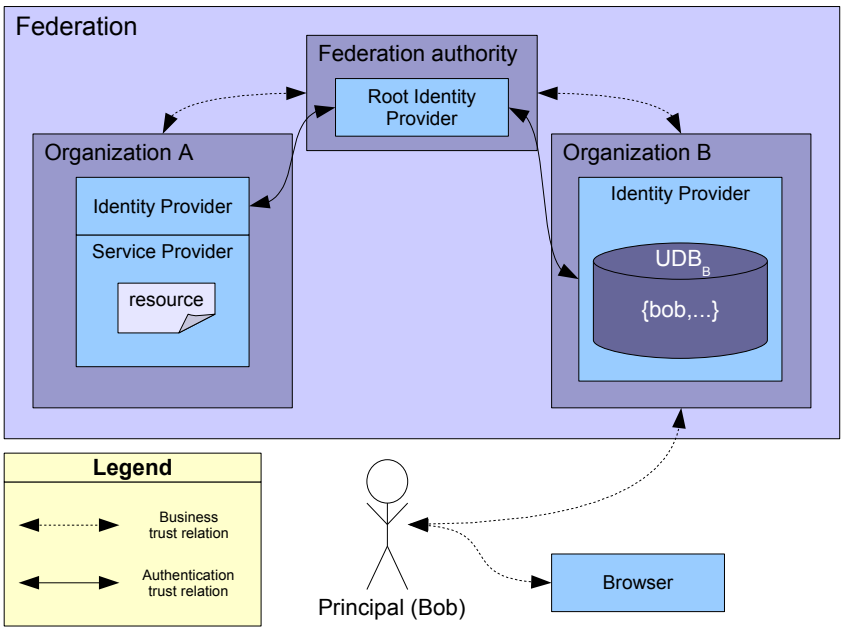

Figure 4.2: The trust architecture used in the example

- A load balancer, that forms the entry point of every administrative domain and distributes the network traffic among every physical identity provider, presenting them as one logical identity provider to the outside world;
- The Authentication Service Provider (AuthSP), a component used by the asserting party to delegate the actual authentication process.

4.3.2 Trust architecture

Figure 4.2 illustrates which trust relations exist in this scenario and what the nature is of these relations. Organization A and organization B both have a direct relation to the federation authority and mutually have an indirect brokered trust relation and no direct relation. The federation authority is responsible for distributing both business trust and authentication trust. As explained in section 2.3, these roles are known as the Business Trust Broker and Certificate Authority respectively.

Business trust between the organizations and federation authority is established by policy enforcement. Legal agreements between the parties on what

policies should be obeyed allow for mutual trust. Without the signing of these contracts, the organization cannot be a part of the SURFnet federation.

There has also been an exchange of certificates between the identity provider of organization A and the root identity provider. These certificates are used to sign the assertions. This exchange has also taken place between the root identity provider and the identity provider of organization B, for the same purpose. The authentication assertions consist of an session identification combined with a proof. The relying party receives the authentication assertion, that was send via the user, and verifies this by directly querying the assertion party, by presenting the session id and proof. Signing makes sure that the assertions can not be forged or altered.

Furthermore, the principal has formed an agreement with organization B, that makes that organization B trusts him and adds his information to the local administration. Most probably, Bob has signed documents stating that he will not abuse his gained privilege to be authenticated by organization B. Therefore, Bob and organization B have a business trust relation.

The principal has a business trust relation with the browser in a less strict way than illustrated in the discussion on trust models. He did not sign any formal agreement with the browser, but trusts the browser and its input devices in the sense that no other (unauthorized) principals can eavesdrop in order to obtain Bob's credentials.

4.3.3 Standards

The SURFnet Federation uses A-Select for its authentication actions. A-Select servers communicate mutually using URL encoded parameters. A-Select was not specifically built for use in federation-like environments, since this concept was not common at the time. However, structural design and techniques that are used often in federation are also used in the SURFnet. For instance, the cross A-select mode allows for Single sign-on scenarios that are comparable to the Web-browser SP initiated SSO profile. So, although A-Select does not comply to the standards on federated identity management mentioned before, the overall structure can be characterized as such; A-Select's flexibility allows it to be used in the federation environment. Furthermore, A-Select is currently capable of producing SAML v1.1 compatible communication messages in order for it to communicate with Shibboleth environments. This is not the case

in this scenario, since no Shibboleth enabled identity providers are part of the federation.

4.3.4 Use case

The scenario is based on the transient pseudonym use case of SAML. The transient pseudonym use case is the only of the SAML use cases that is supported in the current version of A-Select. This is due to the fact that linking accounts is not possible at the moment, and this functionality is necessary for the out-of-band and persistent pseudonym use cases. The transient pseudonym use case can be used to authenticate principals without having to manage their data locally, which is the case in A-Select.

4.3.5 Scenario process

Assumptions

There are some assumptions that apply to the illustrated process.

- For the sake of simplicity, the authentication step is not regarded. In A-Select, the act of authentication is performed by a component called the Authentication Service Provider (AuthSP). The A-Select server communicates with a AuthSP in order to authenticate a principal. There are several different AuthSP's available for different authentication methods, but the exact method of authentication is not relevant within the context of this scenario. It is therefore assumed that there is an arbitrary AuthSP present that is capable of authenticating the principal;
- The principal does not have any existing single sign-on sessions with any Service or Identity provider, but is willing to start them. As we will see later, because the principal does not have any SSO sessions yet, he needs to perform the full authentication process;
- but does have a user account at the Identity provider of organization B;
- The credentials that the principal provides are valid and allow the principal to authenticate properly;
- The user name that Bob uses to identify himself directly relates Bob to his home organization. This can be achieved by using a naming convention

within the federation, for instance 'username@organization'. Bob's user name is 'bob@organizationb'.

- The A-Select authentication procedure usually consists of two major steps: The authentication request, which is described below and the credentials verification step. The latter is not discussed here, since it is only used to verify that the authentication assertion is received properly by the relying party. We assume that that is always the case, so any call to verify the credentials receives a positive answer. Because of the homogeneous nature of the SURFnet federation, we can also assume that this method of verification is valid and understood by both the asserting and relying party.

Scenario steps

The steps taken in the scenario are illustrated in the sequence diagram of figure 4.3. A step by step description of the scenario process follows now:

1. Bob uses his browser to access the resource at the service provider of organization A;
2. The service provider receives the request and verifies the status of the session;
3. The session does not have the authentication tokens required for the disclosure of the resource Bob requested. The service provider detects this and redirects Bob to the local identity provider, also part of organization A;
4. The local identity provider needs to know and requests the user name of Bob;
5. Bob enters his user name and domain and sends the information back to the identity provider;
6. The local identity provider recognizes that the user 'bob@organizationb' is a concatenation of a user name and a domain name and determines that this user cannot be authenticated locally;
7. The local identity provider redirects Bob to the root identity provider, in order to delegate the authentication of Bob;

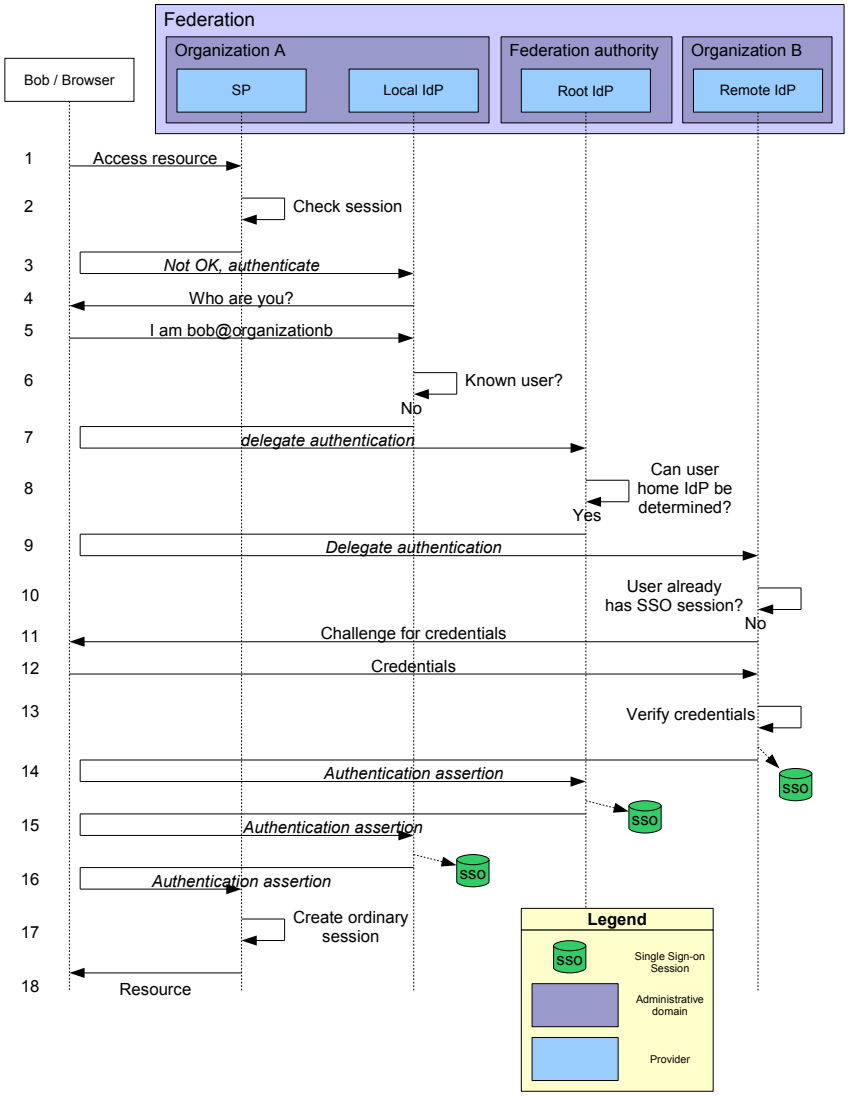

Figure 4.3: Sequence diagram of single sign-on procedure in the SURFnet Federation

8. The root identity provider recognizes that user 'bob@organizationb' is a user of organization B;
9. Bob is redirected to his home identity provider, that of organization B, in order to delegate the authentication process. The root identity provider remembers which identity provider Bob came from;
10. The identity provider of organization B determines if Bob already has a Single Sign On session with this identity provider;
11. Because Bob does not have an established SSO session, Bob is required to authenticate. Bob's user name is already known, so Bob is now challenged for his credentials to prove his identity;
12. Bob sends his credentials to his home identity provider;
13. The credentials are verified with the user database and Bob's identity is confirmed;
14. Bob's home identity provider creates an authentication assertion that can be used to prove Bob's identity at any of the other identity providers, and stores Bob's single sign-on session in the session database. With this assertion, Bob is redirected back to the root identity provider;
15. The root identity provider recalls what identity provider originally requested the authentication of Bob and redirects Bob back, with the authentication assertion still attached. The root identity provider also adds Bob to the session database to enable Bob's single sign-on session;
16. The local identity provider creates a single sign-on session and redirects Bob back to the service provider;
17. This step is equal to step 2, except for the fact that the user possesses a valid authentication assertion and so a session is created;
18. The service provider is now able to verify Bob's identity and thereby decides to supply Bob with his requested resource.

Alternate path

The procedure indicated above is used most often, but not the only possible scenario. An often used alternative is by using a list of identity providers.

At step 5 Bob indicates that organization A is not his home organization, but does not supply his user name yet. Step 6 is skipped, since there is no user name known. Step 7 is the same, except from the fact that no user name is provided. In step 8, the root identity provider can not determine the home identity provider of the user, and therefore sends a list of identity providers to Bob. Bob selects his home identity provider, so that that identity provider is known to the root identity provider and Bob is redirected to his home identity provider (Step 9). In step 11, Bob is not only asked for his credentials, but also required to submit his user name, since it is not known at this point. Steps 13 to 18 are left unchanged.

4.4 Issue analysis

4.4.1 User Control

As indicated in chapter 3, full user control in a federated identity management is hard to maintain. In a homogeneous environment such as the SURFnet, making sure that user control is taken into account by every federation member needs to be enforced by policies. Fortunately, there is a central federation authority that is in charge of auditing possible federation members before exchanging authentication trust tokens, such as certificates.

The federation authority is capable of enforcing user control requirements of the federation members through one of the following means:

- The federation authority can audit currently available policies on how users can access their own information. Since it is the first time the possible federation member becomes part of an open system such as a federation, it most probably does not specify policies on this subject. In this case the next option can be applied;
- The SURFnet federation authority can specify its own policies on how user control requirements need to be enforced. Whenever an organization wants to be a part of the federation, it is obliged to apply the policy as offered by the authority.

The current version of A-Select however is not yet capable of enabling full user control. This was not the intention of the product as it is, since it was developed to be an authentication mechanism, not an identity management

system. In a federated identity management context as the SURFnet federation, A-Select is used in an identity management system. In the SURFnet federation, A-Select has grown beyond being a simple authentication mechanism, because it is also responsible for releasing user attributes and managing policies for that. This is the point where user control becomes important. Therefore, it would be a wise idea to make A-Select progress to the next step in identity management by integrating a great deal of user control.

If we take the four concepts that were depicted in section 3.2.1, this results in the following concrete changes to A-Select:

- *Notice* - A-Select already provides some information on what happens. However, it is up to the application of the requesting party to inform the user on what attributes are requested. An user control improvement would be to indicate what attributes will be exchanged *before* authentication, so that the user can decide to cancel the procedure;

- *Choice* - A next step is to improve this information message to a module that allows the user to select attributes that he does or does not want to exchange;

- *Access* - Following this, a module that allows the user to change his attributes, or at least some of those, can be added. Of course, the user must not be able to change identifiers or other attributes that are necessary for the authentication procedure itself;

- *Security* - This is a point that would not require many changes. A-Select has proved itself secure, so users do not have to worry about security problems.

4.4.2 Identity provider discovery

This issue is discussed in paragraph 3.2.2. In this paragraph an answer is presented to the questions if and how this problem manifests itself in the SURFnet federation example.

The question is whether the dynamic connection of new identity providers could form a problem, regarding the list of identity providers a user is presented with when the home identity provider of the user needs to be determined for authentication. In the case of SURFnet, the great advantage is that

there is only one authority that is responsible for distributing both business and authentication trust (see paragraph 4.3.2). It has full control of the workflows that come with federated identity management, since all go through the root identity provider. The root identity provider thereby acts as a kind of router. Parties that want to connect to the SURFnet, therefore contact the root authority and sign contracts to comply to the policies as stated by the SURF foundation (business trust). Only then, the certificates are exchanged (authentication trust) and the party is connected to the SURFnet. With this information in place, a list of identity providers can be generated easily.

Now, when a principal needs a delegated authentication from a different SURFnet member, and the home identity provider cannot be determined using a cookie or extracted from the user name, the root identity provider can present the user with the identity provider list from which the user can choose his home identity provider, as is the scenario sketched in the paragraph on the alternate path for authentication. See paragraph 4.3.5 on the alternate path for remote authentication in the SURFnet.

So, in the SURFnet federation the problem of identity provider discovery does not really exist. Therefore, the use of standardizing tools such as the SAML Identity Provider Discovery Profile [16] does not have to be taken into account for as long as the federation authority remains the leading party in managing the federation. When this position is lost, the problems may become apparent. This may happen when the SURFnet becomes a part of a far wider network of educational facilities of which it is only a local authority. At that point, not all other identity providers may be known to the root identity provider.

4.4.3 Minimal information disclosure

A federated identity management system should aim at reducing the chance and impact of an identity theft event. As indicated in paragraph 3.2.3, using federated identity management may result in an increased chance and impact in the case of a security breach. This is due to the increased amount of possible entry points and references that exist between identity providers. Reflecting this onto the SURFnet Federation example, it is essential to take a look at the relations between identity providers, both from the architectural and the trust perspective.

Policy-wise measures

The minimal information disclosure problem in a federation can be reduced using attribute release policies. In A-Select these policies are possible, although their functionality is very limited. Depending on the name of the attribute, the attribute is either released (if found) or not released. Policies can be specified per provider. It is possible to indicate a wildcard attribute name, thereby automatically releasing all attributes from the indicated backend.

It would be a good idea to thoroughly revise the attribute release policy technique used in A-Select. Defining policies more fine grained, for example a policy for every user or user group, would be a good start. Also, an indication of the time a certain attribute may be retained is useful. Furthermore, the user should be able to indicate what attributes are released to what provider. This was also discussed in the paragraph on user control. This is a big step, but a necessary one, since the user is the only one who can properly determine what should be done with his identity information.

Technical measures

As indicated earlier, the transient pseudonym identifier federation approach which is part of the SAML v2.0 approach provides a fully anonymous authentication service. Since the SURFnet Federation only uses this for federating identities, it is capable of providing authentication for users, that are left anonymously to the service provider. With the anonymous identifier, attributes can be gathered from the identity provider. So, even if the user is anonymous, an attribute release policy must be present to manage what attributes can be distributed to what service provider. Special attention in specifying these attribute release policies must be paid to the disclosure of identifying information. If, within this transient pseudonym scenario, identifying attributes can be acquired, the anonymity advantage of it is lost.

The positive fact of supporting only the transient pseudonym method is that there is still only one point of entry in the system, thereby reducing the chances of a security issue.

4.4.4 Enforcement of attribute mappings

There are two components in the SURFnet federation that determine the way the attributes are exchanged. First there is the A-Select attribute release policy. This is a part of A-Select that can be used to define what attributes may be disclosed to what party. That second party may be a web application (service provider) or an identity provider, in the case of a cross A-Select configuration in a federation. Of course, the second situation is the most interesting in this context. The problem is that, apart from the attribute naming, there is no way this component can change the attribute. So, the attribute's syntax and semantics are determined solely by the backend from which the attribute is collected. The SURFnet foundation anticipated on this problem by specifying a list of attributes, including their naming, syntax and semantics. Unfortunately, this list consists only of two attributes, being the user ID and the organization of the user. These attributes are considered to be essential in order to let the federation function properly. Although true, the short list of attributes is of course insufficient for enabling the federation with attribute exchange capabilities. It may therefore be interesting to take a look at how to overcome this problem, although it may be partial.

Standard schema

A first and easy solution would be to oblige members of the federation to present their attributes compliant to a standard schema. Examples of these kinds of schema are the defined standards by EduPerson [8], InCommon [17] and ID-SIS-* [22, 23]. These schema are far more extended compared to the schema consisting of the two existing attributes defined by the SURF foundation.

Automatic schema mapping

A different approach is to solve the problem technically, by transforming the attribute layout. As can be read in paragraph 3.2.4, a mediator component is necessary to transform the attributes from the supplier's lay-out to the lay-out used at the requester side. The level of complexity of the mediator component determines whether it is possible to perform translations of attribute names, attribute syntaxes or even attribute semantics.

In the case of the SURFnet federation, the implementation of this approach is interesting to discuss. Because attribute exchange also takes place via the root identity provider, that manages the mediator component, each local identity provider need not worry about 3rd party attribute lay-outs. It can delegate transformations to the root identity provider that, for example with the entrance of the organization in the federation, receives attribute information as it is used by the identity provider. The federation authority then adds the appropriate XSLT sheets to the mediator component and the attributes can be exchanged.

An implication is however that, as noticed in [14], the central administration of attribute transformer XSLT sheets requires a large amount of sheets, to map every schema to all other schema ($n \times (n-1)$, where n is the amount of identity providers). So, especially if the federation becomes extensive and many identity providers are added, it would be wise to think about more efficient storage of mappings.

4.4.5 Provider identity and trust

The problem of provider identity and trust distribution is in essence purely a trust distribution issue (see paragraph 3.2.5). Therefore we can divide it into two parts: Distribution of business trust and distribution of authentication trust. In our example, it is apparent that these two forms of trust distribution are elaborated to different extends. Because of the presence of a central authority that takes care of validating potential federation members, the distribution of business trust is clearly very well taken care of. Before these potential members are connected to the SURFnet, they are obliged to sign contracts that cover a wide array of legal aspects. So, federation members can safely presume that the parties the authentication assertions are exchanged with have met the legal requirements of the federation.

The authentication trust distribution is a more complicated process, since it requires the exchange of certificates between the root identity provider and the new identity provider. This action must be performed manually, so it is a great advantage that there is only one central authority to exchange certificates with. If the SURFnet federation were fully decentralized, an automatic certificate exchange process was far more urgent. Therefore, a Notary server that can be used in for automatic distribution of certificates as was proposed in [9] and

was discussed in paragraph 3.2.5, is unnecessarily complex for implementation in the SURFnet federation.

Manual certificate exchange will suffice for the time being. However, there are features that the use of a Notary server will bring that may be interesting in the future:

- The notary server approach allows the identity provider to blind the assertion. This makes the assertion impossible to extract information from by the root identity provider;
- Using the notary server will introduce a high level of accountability, because all assertion signatures are stored. There is currently no traceability function in the SURFnet root identity provider, thus actions cannot be accounted for.

To draw a brief conclusion, there are some interesting features a notary server will enhance the SURFnet federation with. The downside is that it incorporates a great deal of cryptographic steps, making it a far more complicated system, compared to the current SURFnet.

4.4.6 Single sign-off

The Single sign-off scenario is well defined in the WS-Federation specifications [21]. However, the specification does not mention how to make sure that all sessions are cleaned properly and how to inform the user in the case of a success or failure of the single sign-off process. The scenario as sketched is quite loose and voluntary. This is also due to the fact that parties relying on sign-off messages must register at the identity provider using WS-Eventing protocols.

The current version of the SURFnet federation does not support Single Sign-off at all, so a solution is desirable. The solution must be easy to apply and should not change the way the local single sign-on sessions are managed. However, it should be so flexible that it is capable of stopping single sign-on sessions throughout the federation.

When we take a look at the SURFnet Federation, and in particular the simple use case depicted in paragraph 4.3, it could well be that three single sign-on sessions are created. The actual number depends on how the individual identity providers are configured. Each of the three identity providers can initialize a single sign-on session for itself. Each of the single sign-on sessions have a

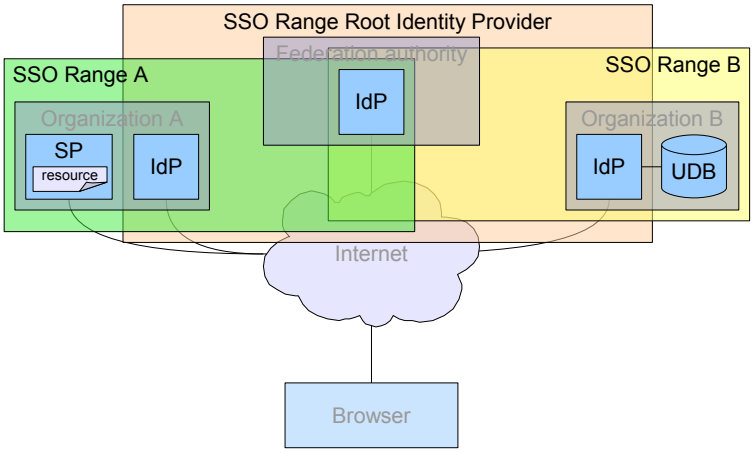

Figure 4.4: The ranges of the individual single sign-on sessions

range in which they can influence authentication requests of the relying parties. These parties can be service providers, such as the service provider of organization A, or can be other identity providers. The possible ranges of the single sign-on sessions are illustrated in figure 4.4.

What is desired is a solution to clear all single sign-on sessions, without having to change much in the existing federation. Fortunately, the root identity provider is a component in the SURFnet federation that can be altered severely without frustrating the individual organizations' federations. Furthermore, there is the advantage of the identity providers being the passing hatch for all authentication requests. This enables the identity providers to built an administration that can be used later, when the single sign-on sessions must be cleared.

Using this administration, there are two basic models that can be applied when implementing a solution to the Single Sign-off problem; By using redirects and by using direct IdP-to-IdP messaging. Both have advantages and disadvantages that will be discussed next. In both cases, it is assumed that sessions are maintained at the identity provider, as opposed to client-side. This implies that if a principal's identifiers are removed from the session administration, he is signed off, although the cookie identifying the principal may still exist. If the principal tries to access a resource for which a session is needed, the identifier in the cookie will not match a record in the session administration

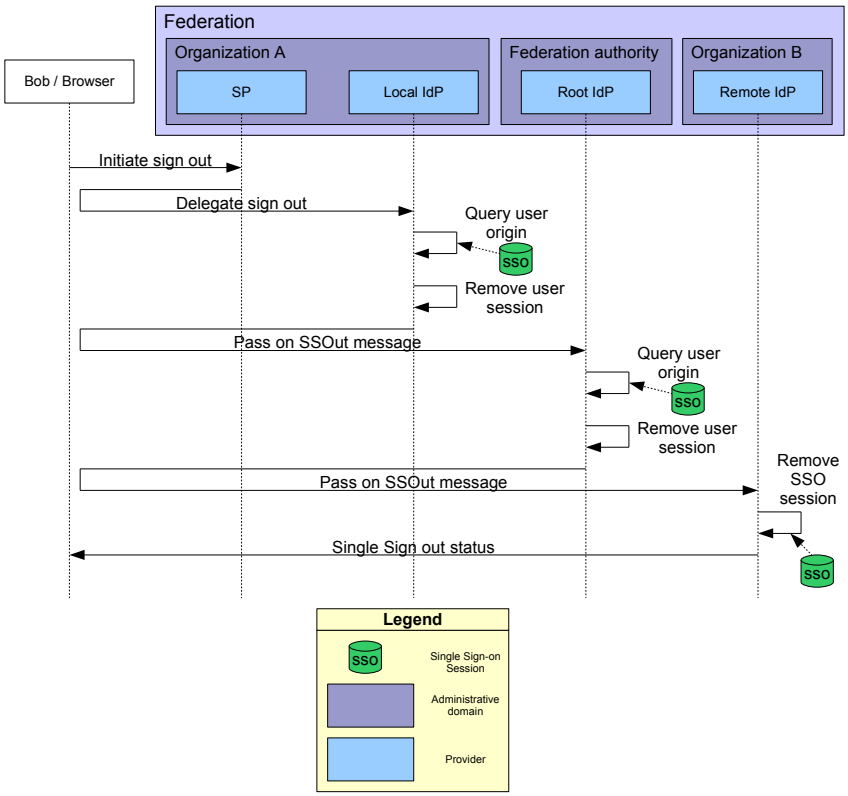

Figure 4.5: Single sign-off procedure by using browser redirects

and will therefore not grant the principal access. The principal can be notified in the case of such an event.

The first scenario is depicted in figure 4.5. Although highly discouraged in the WS-Federation standard [21] because of its fragility (see also paragraph 3.2.6), it has some positive points as well. First of all the user is in full control of the procedure and can also be informed immediately and adequately in case a failure in the process occurs. Furthermore, because of its front end approach, standard requests can be used. Although special methods must be implemented, that generate a redirect to the next identity provider, it does not require the use of special messaging protocols.

The second solution to the problem can be found in backend messaging. At first sight, this appears similar to the approach discussed in the WS-Federation standard. However, there is a substantial difference. Since a accurate administration of single sign-on sessions is already present, the WS-Eventing approach

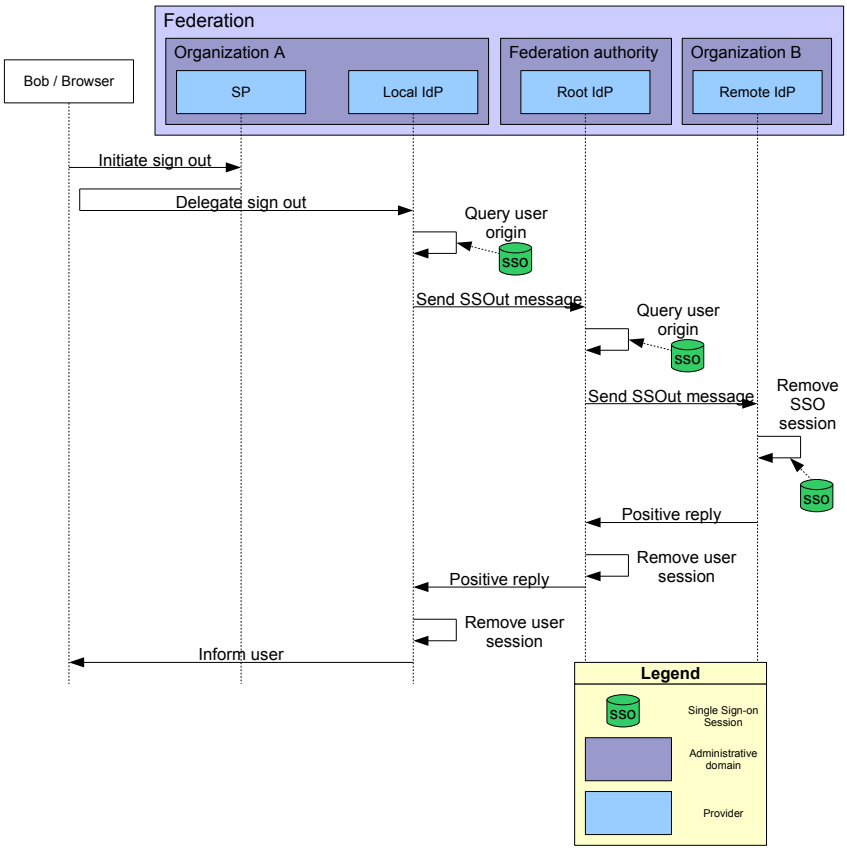

Figure 4.6: Single sign-off procedure by using IdP-to-IdP messaging

is no longer necessary. The identity provider is aware of the tickets and to which providers they are granted. The messaging scheme is therefore quite straightforward. The scenario is illustrated in figure 4.6. The disadvantage in comparison to the redirect approach is that in the case of a failure somewhere beyond the initializing identity provider, it cannot determine the exact point of failure. Therefore, the identity provider cannot guarantee that the procedure is finished successfully. Therefore, it could be that the user is left with sessions that he is not aware of, which is a security hazard in some cases. Furthermore, this approach requires the necessary additions in message handling, since the backend messages that indicate a logout notification must be processed. In contrast to the messages used for the frontend approach, handling these messages is not implemented by default.

It is clear that both solutions have advantages as well as disadvantages. Further research to the more subtle differences is necessary to find a solution that is solid, easy to monitor and easy to implement.

Chapter 5

Conclusion

The final chapter of the thesis discusses the conclusions and recommendations. These conclusions are founded in chapters 2, 3 and 4. The basis for the recommendations can be found mostly in chapter 4.

5.1 Discussion

To conclude, let us recall the main goal that was formed the basis for the thesis. It states: "*Identify problems to occur in large federated identity management systems based on A-Select*". As said, it is divided into three sub-goals. Each of these sub-goals are discussed in the following three paragraphs. Because the third sub-goal completes the trajectory towards the main goal, the main conclusions can be found in that paragraph; 5.1.3.

5.1.1 Establishing the principles and foundations

The thesis establishes the following: In chapter 2 the principles of federated identity management are discussed. The purpose of this research is to make clear statements on federations and its underlying foundations. The information described in this chapter forms a basis for the rest of the thesis. This fact illustrates that the chapter contains the right subjects and depth. The chapter provides a broad and yet deep enough overview of the concept which will serve Alfa & Ariss well and provides enough insight to the reader.

Specifically the conceptual model and the glossary, which was a required part, may be used in the future to reach a higher quality for products (A-Select)

and documentation respectively. The products and other business solutions may benefit from the conceptual model supported by the other parts illustrated in this chapter. These form a reminder of the structure of a federated identity management system, thereby indicating the role of the individual components and how they relate. The glossary will provide a consistent vocabulary, that was needed for the composition of clear documentation accompanying the software.

5.1.2 The issues of federations

Chapter 3 discusses some of the problems that are currently to be expected in a federated identity management system. The issues that are depicted are chosen in corporation with Alfa & Ariss following an extensive literature research. These issues define what Alfa & Ariss have experienced in the past and are expected to run into.

For many of the issues, solutions are added. The solutions can be applied to A-Select, and keep in mind the scale, potential other models and use cases that A-Select may be working with in the future, in the federated identity management world. This helps the concrete application of the issues to A-Select, of which results are discussed next.

5.1.3 The results considering A-Select

Finally, the application of the issues to a concrete scenario gives an insight into whether these issues are to be expected and if so, what can be done about it to solve them or reduce the impact. The occurrence of the problems are found by looking at a concrete example of A-Select, the federation identity provider, in a federation. The knowledge gained in developing this SURFnet Federation helps producing a thorough investigation.

There are some prominent conclusions that can be drawn from the application:

- The use of a central authority solves a great amount of potential problems. Especially the trust distribution related problems can be easily avoided by linking business trust establishment and authentication trust establishment processes, by coupling legal contract signing and certificate exchange. Also, the central authority controls the channels between

federation members, thereby placing restrictions on data exchange. The reduction of what data can be exchanged and how this data should be formatted, reduces the chances of unnecessary information disclosure and data format discrepancies severely. It is very clear that without this central authority many more problems are to be expected;

- The SURFnet federation is a very closed environment. Potential members are obliged to sign contracts and use specific software components to connect to the SURFnet federation. This latter requirement makes the SURFnet federation a homogeneous federation. Many of the interconnectivity problems apparent in heterogeneous federations do not occur in the homogeneous variants. So, besides the use of a central authority, the fact that the SURFnet federation is homogeneous and legally restricted reduces the amount of problems as well;

- The adaptation of A-Select to work in a federation has made it an identity management system. A-Select was originally built for the purpose of authentication only. Now it is applied in a federation system, as an organization wide identity provider as well as a root identity provider playing the role of central authority. A-Select is now responsible for releasing attributes, making it a true identity management system. This transition is not yet recognized as such, although it should be.

5.2 Recommendations

There are some specific recommendations following directly from the aforementioned conclusions. Since most of the problems are less significant in the SURFnet situation, the recommendations should mostly be considered when placing A-Select in a typical heterogeneous federation. This can also imply a very large federation of which the SURFnet forms a small part.

- The users for whom identity information is managed using A-Select should be given the privileges and possibilities to alter their information. This is not a simple addition to A-Select, so a transition with multiple steps, as discussed in section 4.4.1 is recommended;

- Related to the previous point is the fact that finer grained attribute release policies must be applicable. Also, users should be able to indicate what

information they want to exchange with other federation parties. So, attribute release policies should be defined partly by the identity provider and partly by the user, thereby keeping in mind that the user is not able to technically limit the federation functionality (e.g. The user should not be able to prohibit the exchange of a pseudonym identifier);

- In order to overcome compatibility issues when interchanging information with partners in heterogeneous networks, it is wise to make use of other already established standards. For example SAML 2.0 would be a good start. The use of the techniques for attribute conversion can be used for communication with specific parties, but the fact that this method costs more time to implement makes this the less desired solution;

- In order to ensure a certain degree of user privacy, the use of policies should be enforced. Organizations that want to use A-Select in a federated environment can be informed on the use of policies that can be agreed on together with other federation members. In centrally organized federations such as the SURFnet federation, the central authority has the power to make these agreements mandatory.

5.3 Future work

Since the research area of this thesis is quite large, there are some open items that can be investigated further in the future. Some of these topics form minor extensions, others can resolve in extensive projects, stretching over several months.

- The Involvement of distrust. In the research, there has been a large deal of attention to the topic of trust and trust relations. However, there is no mentioning of distrust. Concretely, this means that it is unclear how to handle a distrust relation in federated identity management systems, for instance when two providers are connected indirectly, but are not willing to exchange information. This is not an easy subject and requires more work;

- Although not a specifically valid future research topic when you look at federated identity management, the transition towards user centricity can be very interesting for A-Select. Many of the issues discussed in this

document have a relation with the enlargement of user interaction in the identity management system. This can ultimately lead to a user centric identity management system. A possible transition towards User centricity for A-Select takes a severe amount of research work to be done, before the product can be changed. A-Select has hardly left the status of authentication mechanism, so user centric identity management is quite far off;

- More research may also be needed on heterogeneous federations. Of course this depends on the amount of interest Alfa & Ariss have on making A-Select more flexible in order to be compatible with other federation supporting systems. Further research may be conducted on what standards to implement and how typical protocols and models can be fit within the A-Select product.

Bibliography

[1] Gail-Joon Ahn and John Lam. Managing privacy preferences for federated identity management. In *DIM '05: Proceedings of the 2005 workshop on Digital identity management*, pages 28–36, New York, NY, USA, 2005. ACM Press.

[2] The A-Select Authentication System. http://www.a-select.org.

[3] Elisa Bertino, Abhilasha Bhargav-Spantzel, and Anna Cinzia Squicciarini. Policy languages for digital identity management in federation systems. In *POLICY '06: Proceedings of the Seventh IEEE International Workshop on Policies for Distributed Systems and Networks (POLICY'06)*, pages 54–66, Washington, DC, USA, 2006. IEEE Computer Society.

[4] Kim Cameron. The laws of identity. http://msdn.microsoft.com/library/en-us/dnwebsrv/html/lawsofidentity.asp, last visited 11 Oct 2006.

[5] Scott Cantor, John Kemp, and Darryl Champagne. Liberty ID-FF Bindings and Profiles Specification, version 1.2-errata-v2.0. Technical report, 2004.

[6] James Clark. XSL Transformations (XSLT). W3C Recommendation, http://www.w3.org/TR/xslt.

[7] Alan Geller (Editor). Web Services Eventing (WS-Eventing). Specification, http://download.boulder.ibm.com/ibmdl/pub/software/dw/specs/ws-eventing/WS-Eventing.pdf, 2004.

[8] Eduperson object class. Overview of specifications, http://middleware.internet2.edu/dir/schema/, last visited 2 May 2007.

[9] Michael T. Goodrich, Roberto Tamassia, and Danfeng Yao. Notarized federated identity management for web services. In *DBSec*, pages 133–147, 2006.

[10] Hugo Haas and Allen Brown. Web services glossary. W3C Working Group Note, http://www.w3.org/TR/ws-gloss, last visited 29 Nov 2006.

[11] Jeff Hodges. Liberty technical glossary v2.0-05. Liberty specification.

[12] Jeff Hodges, Rob Philpott, and Eve Maler. Glossary for the oasis Security Assertion Markup Language (SAML) v2.0. OASIS Standard, 2006.

[13] Wolfgang Hommel. An architecture for privacy-aware inter-domain identity management. In *16th IFIP/IEEE Distributed Systems: Operations and Management (DSOM 2005)*, pages 49–60, 2005.

[14] Wolfgang Hommel and Helmut Reiser. Federated identity management in business-to-business outsourcing. In *Proceedings of the 12th Annual Workshop of HP OpenView University Association (HPOVUA 2005)*, pages 81–93, 2005.

[15] Wolfgang Hommel and Helmut Reiser. Federated identity management: Shortcomings of existing standards. In *9th IFIP/IEEE International Symposium on Integrated Network Managment (IM 2005)*, 2005.

[16] John Hughes, Scott Cantor, Jeff Hodges, Frederick Hirsch, Prateek Mishra, Rob Philpott, and Eve Maler. Profiles for the OASIS Security Assertion Markup Language (SAML), v2.0. Technical report, OASIS, 2005.

[17] Incommon Attribute Table. Technical Specification, http://www.incommonfederation.org/attributes.cfm, last visited 2 May 2007.

[18] Internet2. http://www.internet2.edu.

[19] Internet2. The Shibboleth Architecture. Non-normative Technical Overview, 2005.

[20] J. Jeff, Y. Alan, B. Ross, and A. Alasdair. The memorability and security of passwords – some empirical results. Technical report, Computer Laboratory, University of Cambridge, 2000.

[21] Chris Kaler and Anthony Nadalin (editors). Web Services Federation Language (WS-Federation), version 1.1. Specification, 2006.

[22] Sampo Kellomäki and Rob Lockhart. Liberty ID-SIS Employee Profile Service Specification (ID-SIS-EP). Liberty specification, http://www.projectliberty.org/liberty/content/download/1031/7155/file/liberty-idsis-ep-v1.1.pdf.

[23] Sampo Kellomäki and Rob Lockhart. Liberty ID-SIS Personal Profile Service Specification (ID-SIS-PP). Liberty specification, http://www.projectliberty.org/liberty/content/download/1028/7146/file/liberty-idsis-pp-v1.1.pdf.

[24] Charles Knouse. SAML Implementation Guidelines. OASIS Specification, 2004.

[25] Maurizio Lenzerini. Data integration: a theoretical perspective. In *PODS '02: Proceedings of the twenty-first ACM SIGMOD-SIGACT-SIGART symposium on Principles of database systems*, pages 233–246, New York, NY, USA, 2002. ACM Press.

[26] Liberty Alliance. http://www.projectliberty.org.

[27] The Liberty Alliance. Introduction to the Liberty Alliance identity architecture. Paper, 2003.

[28] John Linn. Trust Models Guidelines. OASIS Non-normative specification, 2004.

[29] Teruko Miyata, Yuzo Koga, Paul Madsen, Shin ichi Adachi, Yoshitsugu Tsuchiya, Yasuhisa Sakamoto, and Kenji Takahashi. A survey on identity management protocols and standards. *IEICE TRANSACTIONS on Information and Systems Vol.E89-D No.1*, pages 112–123, 2006.

[30] Marco Casassa Mont, Siani Pearson, and Pete Bramhall. Towards accountable management of identity and privacy: Sticky policies and enforceable tracing services. HP Technical Report (HPL-2003-49), 2003.

[31] NMI-EDIT glossary. Website, http://www.nmi-edit.org/glossary/index.cfm, last visited 29 Nov 2006.

[32] Organization for the Advancement of Structured Information Standards (OASIS). http://www.oasis-open.org.

[33] OpenSAML. http://www.opensaml.org.

[34] Birgit Pfitzmann and Michael Waidner. Federated Identity-Management Protocols where user authentication protocols may go. In *Security Protocols Workshop*, pages 153–174, 2003.

[35] Nick Ragouzis, John Hughes, Rob Philpott, and Eve Maler. Security Assertion Markup Language (SAML) v2.0 technical overview. OASIS non-normative overview, 2006.

[36] Audun Jøsang, John Fabre, Brian Hay, James Dalziel, and Simon Pope. Trust requirements in identity management. In *ACSW Frontiers '05: Proceedings of the 2005 Australasian workshop on Grid computing and e-research*, pages 99–108, Darlinghurst, Australia, Australia, 2005. Australian Computer Society, Inc.

[37] Shibboleth Project. http://shibboleth.internet2.edu.

[38] Simple Object Access Protocol (SOAP). W3C recommendation, http://www.w3.org/TR/soap/.

[39] Abhilasha B. Spantzel, Anna C. Squicciarini, and Elisa Bertino. Integrating federated digital identity management and trust negotiation. Technical report, Purdue University, 2005. CERIAS TR 2005-46.

[40] SURF Foundation. http://www.surf.nl, (page in Dutch, English selectable).

[41] SURFnet. http://www.surfnet.nl, (page in Dutch, English selectable).

[42] Erik Vullings, Markus Buchhorn, and James Dalziel. Secure federated access to grid applications using SAML/XACML. In *Australian Partnership for Advanced Computing (APAC) Conference and Exhibition on Advanced Computing, Grid Applications and eResearch*, 2005.

[43] Thomas Wason. Liberty ID-FF architecture overview. Liberty non-normative description, 2006.

[44] Extensible Markup Language (XML). http://www.w3.org/XML.

Glossary

Account Linkage

"*A method of relating accounts at two different providers that represent the same principal so that the providers can communicate about the principal. Account linkage can be established through the sharing of attributes or through identity federation*"[12].

Active Requestor

"*An active requestor is an application (possibly a Web browser) that is capable of issuing Web services messages such as those described in WS-Security and WS-Trust*"[10].

Administrative Domain

"*An environment or context that is defined by some combination of one or more administrative policies, Internet Domain Name registrations, civil legal entities (for example, individuals, corporations, or other formally organized entities), plus a collection of hosts, network devices and the interconnecting networks (and possibly other traits), plus (often various) network services and applications running upon them. An administrative domain may contain or define one or more security domains. An administrative domain may encompass a single site or multiple sites. The traits defining an administrative domain may, and in many cases will, evolve over time. Administrative domains may interact and enter into agreements for providing and/or consuming services across administrative domain boundaries*"[12].

Asserting Party

"*Formally, the administrative domain that hosts one or more SAML authorities. Informally, an instance of a SAML authority*"[12].

Attribute

"*A distinct characteristic of an object (in SAML, of a subject). An objects attributes are said to describe it. Attributes are often specified in terms of physical traits, such as size, shape, weight, and color, etc., for real-world objects. Objects in cyberspace might have attributes describing size, type of encoding, network address, and so on. Attributes are often represented as pairs of "attribute name" and "attribute value(s)", e.g. "foo" has the value "bar", "count" has the value 1, "gizmo"' has the values "frob" and "2", etc. Often, these are referred to as "attribute value pairs". Note that Identifiers are essentially "distinguished attributes""*[12].

Authentication

"*Authentication is the process of confirming a system entitys asserted identity with a specified, or understood, level of confidence*"[11].

Authentication Session

"*The period of time starting after A has authenticated B and until A stops trusting Bs identity assertion and requires re-authentication. Also known simply as a session, it is the state between a successful login and a successful logout by a Principal*"[11].

Credentials

1: "*Data that is transferred or presented to establish either a claimed identity or the authorizations of a system entity*"[11]. 2: "*An electronic identifier and corresponding personal secret associated with an electronic identity. An identity credential typically is issued to the person who is the subject of the information to enable that person to gain access to applications or other resources that need to control such access*"[31].

Federation

1: "*A federation is an association of organizations that come together to exchange information as appropriate about their users and resources in order to enable collaborations and transactions*"[31]. 2: "*The act of establishing a relationship between two entities*"[11].

Identity

"The essence of an entity. One's identity is often described by one's characteristics, among which may be any number of identifiers. A Principal may wield one or more identities"[11].

Identity Federation

"Creating associations between a given system entitys identifiers or accounts"[11].

Identity Mapping

"Identity Mapping is a method of creating relationships between identity properties. Some Identity Providers may make use of identity mapping"[10].

Identity Provider

"A kind of service provider that creates, maintains, and manages identity information for principals and provides principal authentication to other service providers within a federation, such as with web browser profiles"[12].

IdP

see Identity Provider

Passive Requestor

"A passive requestor is an HTTP browser capable of broadly supported HTTP"[10].

Principal

"A principal is a system entity whose identity can be authenticated. In Liberty usage, the term Principal is often synonymous with natural person or user. A Principal's identity may be federated. Examples of Principals include individual users, groups of individuals, organizational entities e.g. corporations, or a component of the Liberty architecture"[11].

Principal Identity

"An identity being wielded by a Principal, or that is mapped to a Principal in some fashion"[11].

Realm

"A realm or domain represents a single unit of security administration or trust"[10].

Relying Party

"*The recipient of a message that relies on a request message and associated assertions to determine whether to provide a requested service*"[11].

Service Provider

"*A role donned by a system entity where the system entity provides services to principals or other system entities*"[12].

Single Sign-off

"*Inverse function of Single Sign-on, provides a synchronized session logout functionality across all sessions that were authenticated by a particular identity provider*"[29].

Single sign-on

"*From a Principal's perspective, single sign-on encompasses the capability to authenticate with some system entity in the Liberty context, an Identity Provider and have that authentication honored by other system entities, termed Service Providers in the Liberty context. Note that upon authenticating with an Identity Provider, the Identity Provider typically establishes and maintains some notion of local session state between itself and the Principal's user agent. Service Providers may also maintain their own distinct local session state with a Principal's user agent*"[11].

SP

see Service Provider

SSO

see Single Sign On

Trust

"*Trust is the characteristic that one entity is willing to rely upon a second entity to execute a set of actions and/or to make set of assertions about a set of subjects and/or scopes*"[10].