Identity Management for Converged Networks

A Lucent/Sun white paper

February 2006

Editor:

Arnaud Sahuguet, Bell Labs research

Contributors:

Fulup Ar Foll, Sun Microsystems

Rick Hull, Bell Labs research

Todd Morgan, Lucent Technologies

Pat Patterson, Sun Microsystems

Satish Ramamoorthy, Lucent Technologies

Nourredine Rouibah, Sun Microsystems

Arnaud Sahuguet, Bell Labs research

Lauren Wood, Sun Microsystems

1 INTRODUCTION

Federated identity management is a recent, vast and promising area. It has received a lot of attention from various industries, including the telecom industry. Some aspects have already been standardized as part of forums such as 3GPP, OASIS, OMA and Liberty Alliance.

In the context of converged networks, the overall goal of this effort is to "identity-enable" selected telecom network elements and allow converged operators to gain the benefits of standards-based identity and data federation technologies.

By empowering themselves to play a key role in federated identity management, network operators reduce the likelihood of becoming "dumb pipes" and create around them a rich ecosystem of services and business relationships.

In this white paper, we will:

- provide a high-level view of the identity management ecosystem (using Liberty Alliance as the archetypical example)
- identify the value-add for end-users and the benefits & opportunities for the providers of services
- explain some of requirements for the role of identity provider (IdP)
- propose a reference architecture for identity management
- list the key differentiators of the Lucent-Sun product combinations
- present proof-of-concept use cases to validate with a service provider.

We also include a glossary and a list of relevant resources.

Note to the reader: even though this white paper will be, as far as possible, technology agnostic, we will use some terminology and examples borrowed from the Liberty Alliance specification, because it is the most mature technology to address federated identity management in its full scope.

Table of Contents

2 THE IDENTITY ECOSYSTEM

In this section we describe the various players and roles of the identity management "game". We refer the reader to the glossary for some detailed definition of the various terms we will use.

2.1 Overview

A federated identity management ecosystem consists of 3 roles:

- user
- service provider (a.k.a. SP)
- and identity provider (a.k.a. IdP)

A **user** corresponds to a human end-user, a device or a service using or consuming another service.

A **service provider** is often a website providing goods or services, e.g. e-commerce web site, gaming, content, weather forecast, SMS web service, etc.

An **identity provider** is an authority for authentication. It is responsible for delivering authentication assertions regarding a given user that can be presented to service providers for authentication. The authentication itself does not necessarily have to be performed by the IdP and can be delegated to other network components (e.g., RADIUS, LDAP, etc.). The authentication usually does not reveal the user's real identity but rather uses a pseudonym, to preserve the user's privacy.

This implies that the IdP has ways to authenticate users and that there exists a trust relationship between the IdP and the SP. The grouping of entities sharing trust relationships is called a Circle of Trust.

The interaction between the various entities is represented in the figure below.

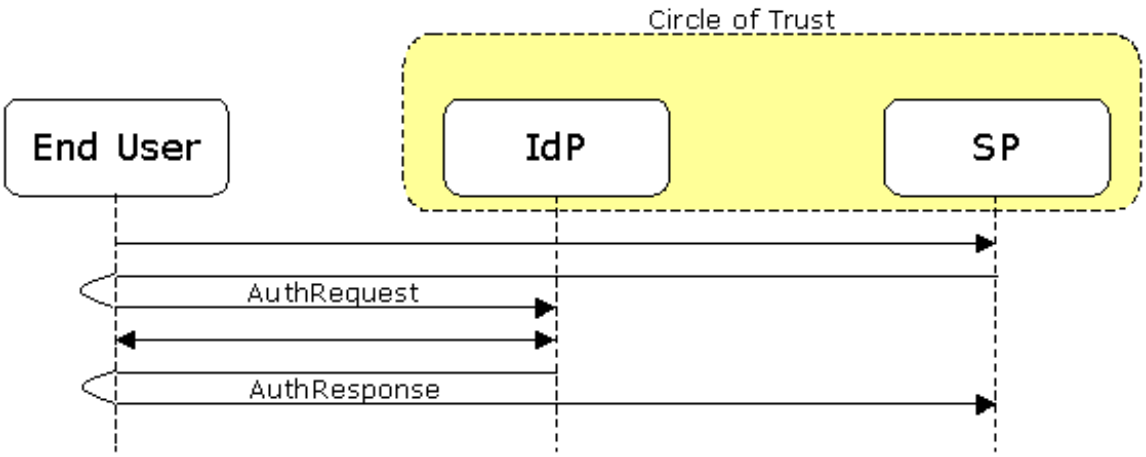

Figure 2.1-a: Liberty typical interaction

Note that Liberty specifications only standardize the request for authentication (<AuthRequest>) and the response (<AuthResponse>). The details of the authentication per se are left to the implementers.

This leaves a lot of freedom and room for differentiation such as:

- methods of authentication (e.g. password, challenge-based, SIM card, etc.)
- choice of identities or personas for a given SP
- policies for generation of pseudonyms (fresh, reuse, etc.)

We provide in the appendix a short glossary with the key terms and concepts that are useful in understanding this new ecosystem.

2.2 Value-add for end-users

Identity management will provide the following benefits to end-users.

2.2.1 Single sign-on (SSO)

Thanks to SSO and identity federation, end users only need to authenticate once to their preferred identity provider and can present claims about this authentication to any service provider inside the circle of trust. There is no need to manage separate usernames and passwords for each service any more.

Moreover, with the presence of network components dedicated to authentication and identity management (IdPs), users can be offered a large choice of authentication methods (e.g. password, challenge-based, voice authentication, SIM card, etc.).

Service providers too have some flexibility in terms of which authentication mechanism should be used. All service providers are not equal and may require different levels of security, e.g. login / password (for mail access), one time password (for VPN access), two factor authentication (for financial transactions), etc.

2.2.2 Anonymity/pseudonymity

SSO – as defined by Liberty – relies on pseudonyms (i.e. opaque identifiers) to guarantee minimum disclosure of identity information. Third parties can access user information or invoke services for a given user with no knowledge of the user's real identity. The use of pseudonyms during web service interaction also minimizes the risk of collusion and tracking.

2.2.3 User profile roaming across services

In the federation model, user data can be shared among SPs and IdPs. End users can provision their data once and let data consumers access it: "enter once, share everywhere". This is very convenient when the data changes.

For instance, a change of address implies a change of billing information for credit cards. Without data roaming, the user would have to manually change the information for all the web services that require this information. Having the data in one place makes the change much easier and also guarantees consistency.

2.2.4 Privacy conscious data sharing

The federation model makes it possible for the end user to control how her data can be shared: when, with whom, for what purposes, etc.

Note that the Liberty architecture offers mechanisms (through the Interaction Service; see Use case #3: interaction service) to interact with the user to get her agreement when disclosure of information is needed.

2.3 Benefits for providers of services

2.3.1 Outsourced authentication

By relying on IdPs for authentication, SPs only need to support a limited set of authentication protocols, leaving all the hard work to this IdP. In some cases, SPs do not even need to manage user accounts.

By accepting users coming from multiple IdPs, service providers can extend their reach to a much wider audience. This is particularly useful in the context of inter-enterprise applications or cross operator services.

Pushing it to the extreme, some services may never actually know the real identity of their users. Having them authenticated by a trusted IdP is often enough.

2.3.2 Outsourced data management

As with authentication, the use of Liberty specifications makes it possible for SPs to rely on a third party to store user related data. This makes it possible for SPs to focus on their core businesses.

2.3.3 Access to richer user information through data sharing

Thanks to Liberty specification data sharing capabilities, SPs can access user data managed by other services to offer a more personalized service. Typical examples include location-based services (where the location is provided by the network operator playing the role of the IdP) or game download where the user device details (color, resolution, etc.) are available through the network operator.

2.3.4 Lower development/deployment costs

By offering ways to outsource authentication and data management, the use of Liberty specifications makes the cost of building and deploying a new service much lower. The service provider can really focus on the service it offers.

2.4 Opportunities for identity providers

2.4.1 Authentication related revenues

Liberty specifications introduce the role of the identity provider, which is responsible for authenticating end users and generating authentication assertions (including anonymous authentication) for service providers. This role requires a trust and security infrastructure and the related services will be offered for a fee.

Vanilla authentication will probably be bundled as part of the end-to-end service. Stronger authentication (e.g. two-factor) may be charged on a per authentication basis. For instance an operator hosting a third party service may want to increase its revenue sharing agreement from 30% to 40% because it now also offers integrated authentication.

Other funding models (e.g. based on advertising) may also be applicable, given the unique position of the IdP: the IdP knows about the service providers visited by the end users. Obviously such services will need to preserve the user's privacy and respect their preferences.

Other revenues and models will emerge when the authentication is being used not only by end users but also by web services.

2.4.2 Data sharing related revenues

Liberty introduces Data Service Template for identity services that offer access to user profile information.[1]

Several data services have already been standardized: personal profile, employee profile, location, presence, calendar, etc. DST also provides extensibility so new data services can be defined for use within a circle of trust.

Information owned by the end user might be free. Information coming from network elements (e.g. presence, location, phone status) will probably be charged for.

2.4.3 Targeted advertising (for IdP)

As mentioned above, identity providers are in a unique position because they become aware of all the service providers visited by the end user. Based on this information, they could – with the agreement of the user – offer some targeted advertising (e.g. via SMS).

[1] Note that this is not the only way to disseminate user's data: attribute assertions embedded in an authentication assertion or SAML attribute queries can be used as well.

3 BEING AN IDENTITY PROVIDER

The IdP is clearly a keystone role in the identity architecture. As mentioned in the previous section, the IdP role can offer some interesting opportunities. In this section we look at the requirements for playing such a role. We also speculate on some of the obvious (and not so obvious) candidates to play such a role.

3.1 Requirements to be an IdP

It is not clear yet who will play the role of an IdP. We can however enumerate the key requirements for this role.

3.1.1 Reputation

Since the IdP is responsible for authenticating users and issuing statements about them, end users will pick their IdP based on the reputation of the organization hosting/running the service. The IdP will post some privacy guidelines or offer a service agreement.

3.1.2 Large set of business agreements

The value of an IdP is proportional to the number of entities in its circle of trust, i.e. the number of business agreements between the IdP and various SPs. The more, the merrier. If the user wants to access a service not part of her IdP circle of trust, she will have either to authenticate directly to the SP or use a second IdP. Having an IdP with a very large and comprehensive circle of trust guarantees the full benefits of the single sign on experience.

3.1.3 Early point of access

The sooner the user accesses the IdP, the better. Ideally, the IdP should be the first point of access the user touches when accessing the network. This way, the single sign on can be achieved very early on with maximum benefits (no more need to authenticate). Reaching the IdP very early also creates some opportunities to (re)use network authentication to authenticate the user to the IdP itself.

3.1.4 Good user experience

Finally, the IdP has to offer a compelling user experience. Authentication has to be as simple and intuitive as possible. This implies supporting a large range of authentication methods to accommodate the user.

3.2 IdP candidates

3.2.1 Web portals

It is pretty clear that the Googles and Yahoo!s of the world will try to become IdPs.

They already are their own IdPs for their own ecosystems: Google with GMail, GTalk, Orkut; Yahoo! with Yahoo! 360 or Yahoo! stores.

As web portals, they are often the first point of access for end users. This puts them in a unique position to let the user authenticate once.

Moreover, this makes a lot of sense to them from a marketing point of view. By delivering authentication credentials to SPs, they know what sites a user is visiting. And they can easily use this information to target their advertising accordingly. This is good for end users (who get ads for products / services they potentially care about) and it is good for advertisers as well.

3.2.2 Access providers

Access providers (e.g. cable, mobile, DSL, Wifi, WiMax, etc.) are also good candidates to become IdPs. Since they provide network access to end users, they are in a unique position (even better than the web portals) to offer authentication services. In some cases, the authentication used to access the network can be (re)used: this offers some unique opportunities for zero-sign-on experience (i.e. where device authentication to the network can be leveraged to authenticate the user using the device, with no need for explicit login from the user herself).

Access providers also have a special relationship with their end-users through billing, enabling users to buy services from the SP. These services may be bought on an anonymous, pseudonymous, or identified basis, depending on the service and the user preference.

3.2.3 Trusted institutions

In Europe, governments have decided to use Liberty Alliance specifications for many e-government services. Government themselves often do not want to be IdPs. Rather they will delegate this role to trusted institutions such as banks, insurance companies and utility companies.

Many universities have already adopted identity management for their own student / faculty / scholar population with Shibboleth.

Banks, credit card companies (Visa, Master-Card, American Express), insurance companies, and Internet security companies (e.g. Verisign) are also good candidates to become IdPs.

3.2.4 Industry specific aggregators

Some specific industries may prefer to establish trust relationships with industry specific IdPs, rather than mainstream or traditional IdPs.

Adult content or online gambling are some industries that come to mind. Governments might also require dedicated IdPs, for contract bidding or security reasons.

3.3 Comparing the IdP candidates

The relative advantages of the various "candidates" are summarized in the table below.

	Reputation	Biz Agreements	Early Access	User experience
Web portals	■	■■■	■■■	■■■
Access providers	■■■■	■■■	■■■■	■■
Institutions	■■■	■■■	■■	■
Industry specific	■■■	■■■■	■	■■■

Note: for institution and industry-specific IdP candidates, the reputation and business agreements are "in the context of the institution or the industry".

4 SHARING USER DATA

With SSO, end users can access services with no need for local authentication, as long as they have authenticated to their IdP.

A natural extension of "*authenticate once, be authenticated everywhere*" is "*enter data once, share it everywhere*", i.e. data federation.

Data sharing can be achieved using 3 different mechanisms:

- attribute statement embedded in authentication assertion
- attribute statement in response to SAML attribute query
- Liberty Data Service Template

4.1 Data sharing via attribute embedded in authentication assertions

An authentication assertion is issued by the IdP in response to an <AuthRequest> from the SP.

The assertion contains various components such as subject description, expiration time, authentication context and digital signature of the issuer.

The IdP can also include zero or more <AttributeStatement> with some <Attribute> and <AttributeValue> elements.

Note that such data sharing suffers from the following limitations:

- data is not query-able. There is an implicit prior agreement between the SP and the IdP in terms of what data needs to be embedded. If these requirements change, so does the agreement.
- embedded attribute assertions are not suitable for highly dynamic data. Given the time to live of an authentication assertion (e.g. 20 minutes), real time data such as presence or phone status will become stale very quickly.

4.2 Data sharing via SAML attribute queries

SAML 2.0 offers a clean way to query data about a given user. The data consumer can simply craft a <AttributeQuery> with <Subject> and <AttributeDesignator> elements to represent the user identity and the list of attributes to query.

In the current version of the specification, an attribute designator is a URI representing an attribute name. In a proposed extension, XPath could be used.

The SAML attribute query protocol is for query only. Updates and publish/subscribe mechanisms are not supported. The protocol is schema-less in the sense that the data consumer must know which attribute to ask for.

Note also that SAML attribute queries rely on an implicit schema that must be agreed upon by the two parties (data consumer and data producer).

4.3 Data sharing via Liberty ID-WSF

Given the importance of data services (i.e. identity services offering access to user profile information) and the critical aspect related to user privacy and confidentiality, Liberty has defined a dedicated framework for data sharing.

The framework consists of 3 pieces:

- Liberty Data Service Template (DST), a set of interfaces for identity services managing user profile data
- Liberty Discovery Service (Disco), a set of interfaces to advertise and lookup resources related to user profile data
- Liberty SIS (service interface specification), a set of schemas and interfaces for actual data services (e.g. personal profile, employee profile, location, presence, etc.)

A data service manages some resources for a given user. A data service offers the DST interfaces, mainly generic query, update, and subscribe operations; SIS are the specific information templates carried within or referred to by DST operations. A data service advertises its resources to the discovery server. A data consumer queries the discovery server for resources about a given user. The Discovery Service returns zero or more resource descriptions pointing to data sources. The data consumer then queries the corresponding data sources.

The Discovery Service can be used as a policy enforcement point to provide access control to data resources. For instance a user may restrict who can have access to her presence and location information. In the figure below, an application (data consumer) needs to retrieve presence and location information for a given user. The application sends a request to the Discovery Service, which returns a list of resources corresponding to the data requested. The application contacts each resource (i.e. an instance of a Liberty Data Service Template) to retrieve the corresponding data. The Discovery Service can enforce access control by simply returning an empty resource.

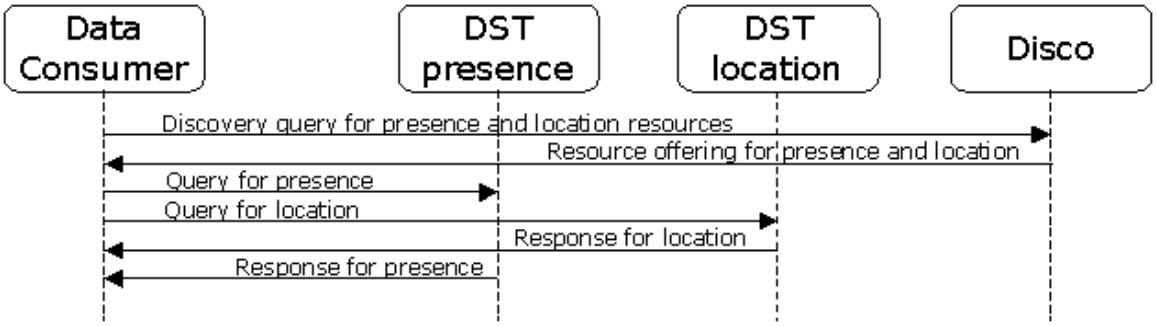

Figure 4.3-a: Discovery Service

4.4 Which sharing method to pick?

Which method to pick depends on the nature of the data to be shared and the mode of consumption. The pros and cons of the various approaches are summarized in the table below.

	Embedded	**SAML query**	**ID WSF**
Request type	Push	Pull	Push & Pull
Query	No	Yes, using attribute name	Yes, using XPath
Support for pub/sub	No	No	Yes
Schema for data	No	No	Yes
Complexity of deployment	Low	Low	High
Interoperability	Low	Low	High
Best for	Static data	Any data	Any data

Figure 4.4-a: data sharing alternatives.

5 MAKING SENSE OF IDENTITY MANAGEMENT

Identity management and related technologies are technically interesting but will only be deployed because they make sense, i.e. they help service providers achieve their two main goals: reduce cost and increase revenue.

5.1 Reducing costs

- Rationalization of subscriber management via a unified subscriber datastore
- Unified provisioning of identity data
- Unified service deployment through unified authentication architecture
- Cheaper deployment of services via 3rd party development, with access to authentication and data sharing done through standardized interfaces.

5.2 Increasing revenue

- Sharing (for a fee) network user data exported through standardized interfaces
- Partnering (with revenue sharing) with 3rd party made easy through standardized interfaces for authentication and data sharing
- Leveraging network authentication for single sign-on (e.g. zero sign-on)
- Leveraging other asserts, such as payment solutions

5.3 Reference architecture to achieve these two goals

A proposed reference architecture to address these two goals is presented in Figure 1.

The arrows in red represent Liberty compliant messages:

- between the user and the SP
- between the user and the Identity Federation component for the SSO protocol (authentication request initiated by the SP)
- between the SP and the Data Federation component, for either a SAML attribute query or a request via the DST.

An extra benefit of such a reference architecture is that it will reduce risks such as identity theft and security breaches.

Notes:

- The link between Identity Federation and Data Federation represents the case where information is shared via attribute assertions embedded in authentication assertions.
- The details of the authentication between the user and the IdP are not standardized by Liberty.
- The details of the provisioning (e.g. how data about a given user ends up in a given datastore) are not standardized by Liberty.

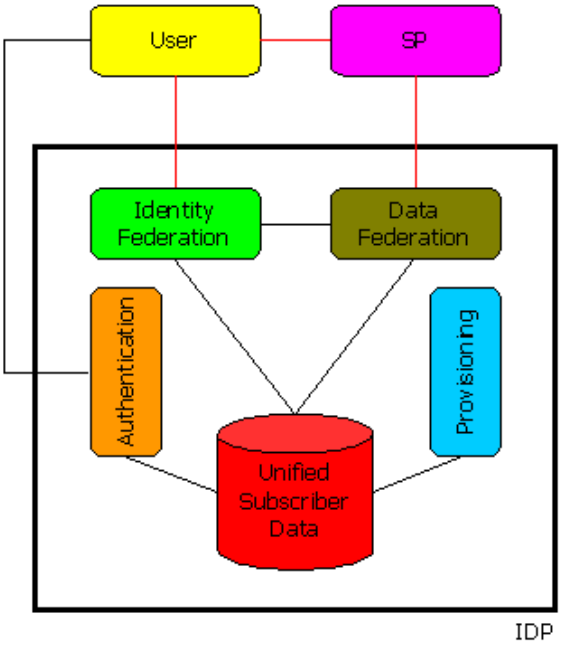

Figure 5.3-a: Reference architecture for identity management.

6 LUCENT/SUN KEY DIFFERENTIATORS

6.1 Key technologies from Lucent

6.1.1 Lucent Datagrid™

Datagrid is a telecom-targeted data integration capability that provides a client application with a global view and a single point of access to subscriber-related data. Integrated access to a wealth of subscriber data (e.g., IMS profile data, location, presence, multimedia mail server) provides substantial time-to-market advantages.

The client application interfaces to Datagrid using one of the supported protocols (e.g., LDAP) and performs queries and updates against the global view with near real-time performance.

Datagrid and data sources may reside in the same trusted network, or in different networks. Datagrid captures the problem of identity translation by maintaining identifier translation tables.

6.1.2 GUPster

GUPster is a reference implementation of the 3GPP GUP framework that provides a privacy-conscious single point of access to subscriber data.

Subscriber data is integrated[2] from multiple sources (distributed across networks) and exported as XML, according to an agreed upon (yet extensible) schema.

GUPster offers a rich access control mechanism that lets subscribers define how data can be shared (by who, for what purpose, in what context, etc.).

Internally, GUPster uses a unified language to describe queries, mappings and access control, which guarantees processing efficiency and correctness.

6.1.3 USDS

The Unified Subscriber Data Server consists of the HLR, HSS, AAA and Datagrid. The HLR, HSS, and AAA functions are implemented within a common super distributed architecture based on the Common Operations (COPS) framework that provides a separation between distributed “control functions” and distributed “data functions.”

This simplifies the use of the HSS, HLR, and AAA data across different access paradigms and protocols to allow simple and cost-effective universal roaming. The

[2] Both Datagrid and GUPster offer integration of subscriber data. Yet they differ because Datagrid assumes that subscribers will belong to few classes for which data mapping are very similar while GUPster has been designed to support per user mappings.

COPS framework is also used in the management of application profile data (e.g., personal contact lists, buddy groupings, etc.)

6.1.4 Lucent IMS client

The Lucent IMS client is a software toolkit to build IMS clients. It offers state of the art SIP stack with some authentication methods. The Lucent IMS client has already been used for applications such as Lucent iLocator, Lucent Active PhoneBook, and an interactive gaming application.

6.1.5 LWS hosting solutions

I should talk about that too. The Lucent Worldwide Services (LWS) team can provide data hosting for SPs who don't want to be bothered about maintaining User data. LWS can host the data in Liberty compliant manner and hence function as a true IdP. This is a key value add for non-telecom, non-technology companies who want to focus on their core competencies.

6.2 Key technologies from Sun

6.2.1 Sun Java™ System Access Manager (for IDP/SP)

Access Manager (AM) provides open, standards-based authentication, policy-based authorization and auditing within a single framework. It offers single sign-on (SSO) as well as enabling data federation. Built using J2EE™ architecture, AM is highly distributed and scaleable. Support of rich Application Programming Interfaces (APIs) and Service Provider Interfaces (SPIs) make AM extensible and customizable.

Some key components are listed below.

Authentication: AM can be easily configured to integrate with existing authentication mechanisms, such as: LDAP, Unix, Windows, SecureID, Radius, MSISDN, JDBC, X.509 certificates, and NTLM. AM's authentication subsystem is extensible, allowing arbitrary authentication mechanisms to be added. Authentication chaining allows arbitrary authentication mechanisms to be grouped together for stronger authentication, and distributed authentication allows separation of credential gathering user interface from the actual authentication server behind a firewall.

Authorization: Both web and arbitrary non-web based policies can be modeled. Policy definition and evaluation APIs are available for applications to use. SPIs allow addition of new subjects and conditions. Out of the box support is provided for:

- Subjects: LDAP groups, roles, Organizations, AM managed groups and roles
- Conditions: time of the day, date/time range, authentication mechanism, authentication levels, user profile attributes and active session attributes.

Federation: Full support for Oasis SAML 1.0, 1.1 and 2.0 and Liberty Alliance ID-FF 1.1 and 1.2 and ID-WSF 1.0 specifications. Circles of Trust (COTs) can be created, and a single instance of AM can support multiple IdPs and SPs. Provides protocol support for SSO, account-linking, unlinking, single logout, discovery service, as well as personal profile and interaction services.

Web Services: AM supports the Liberty ID-WSF stack. Public APIs and plugins allow an existing service to be encapsulated as a Liberty enabled web service.

Agents: A variety of agents supplied by Sun allow existing applications to be quickly SSO and federation enabled without requiring any changes to them.

6.2.2 Sun Java System Federation Manager (Liberty compliant stacks for SP)

Federation Manager extends federation to partners in a hub-and-spoke architecture. It allows spoke partners to more efficiently leverage the core security and identity infrastructures of the hub provider.

FM supports OASIS SAML 1,0, 1.1 and 2.0 and Liberty Alliance ID-FF 1.1 and 1.2 and ID-WSF 1.0 specifications. Its lightweight design and toolkit focus makes it an ideal candidate for enabling participation in a standards based federation, while preserving investments in existing systems.

6.2.3 J2ME™ technology (to enable Liberty on phones, set top boxes, etc.)

Sun is leading various Java specifications that are relevant to end-to-end federated identity management. This includes JSR 279 and JSR 177.

JSR 279 (Service Connection API for Java ME) defines a new high-level API for connection services via frameworks supporting identity based services, discovery, and authentication, such as the Liberty Identity Web Services Framework (IDWSF).

JSR 177 (Security and Trust Services API for J2ME) provides some extensions for USIM / ISIM (through JavaCard).

6.3 Lucent/Sun solution for identity management

We revisit the abstract architecture and show how it can be mapped to Lucent / Sun products.

6.3.1 User

On the user side, Liberty offers numerous profiles that require very little from the user's terminal. Usually a regular web client will do.

More sophisticated mechanisms can be brought to the picture using various Sun's Java technologies living either on the terminal or on the SIM card.

Lucent IMS client solution can also be used to enhance IMS terminals with identity management and data federation capabilities.

6.3.2 Service Provider

On the service provider side, Sun Federation Manager is a nimble toolkit that permits development of Liberty compliant applications using traditional Java Server Pages technologies, running on any web container (e.g. Apache Tomcat, Sun Web server, etc.).

6.3.3 Data Federation

For data federation, a Liberty front end needs to be provided on top of an existing data store. Here again, Sun Federation Manager makes it easy to enhance existing data stores with Liberty compliant interfaces. For HSS data, Lucent Datagrid can be extended into a Data Service Template exporting user related data.

6.3.4 Identity Federation and Authentication

For identity federation, one needs to support the various ID-FF protocols such as single-sign-on, single-logout, etc. Sun Access Manager is a good candidate for this role.

Note that authentication can be handled by Sun Access Manager directly or outsourced to other components, e.g., corporate Active Directory, LDAP directory, AAA server (e.g. Lucent Navis Radius), IMS Generic Authentication Architecture (GAA), etc.

6.3.5 Unified Subscriber Data

For the storage of user profile information (i.e. authentication data and user data), Lucent USDS permits transparent integration of HSS user data for data federation.

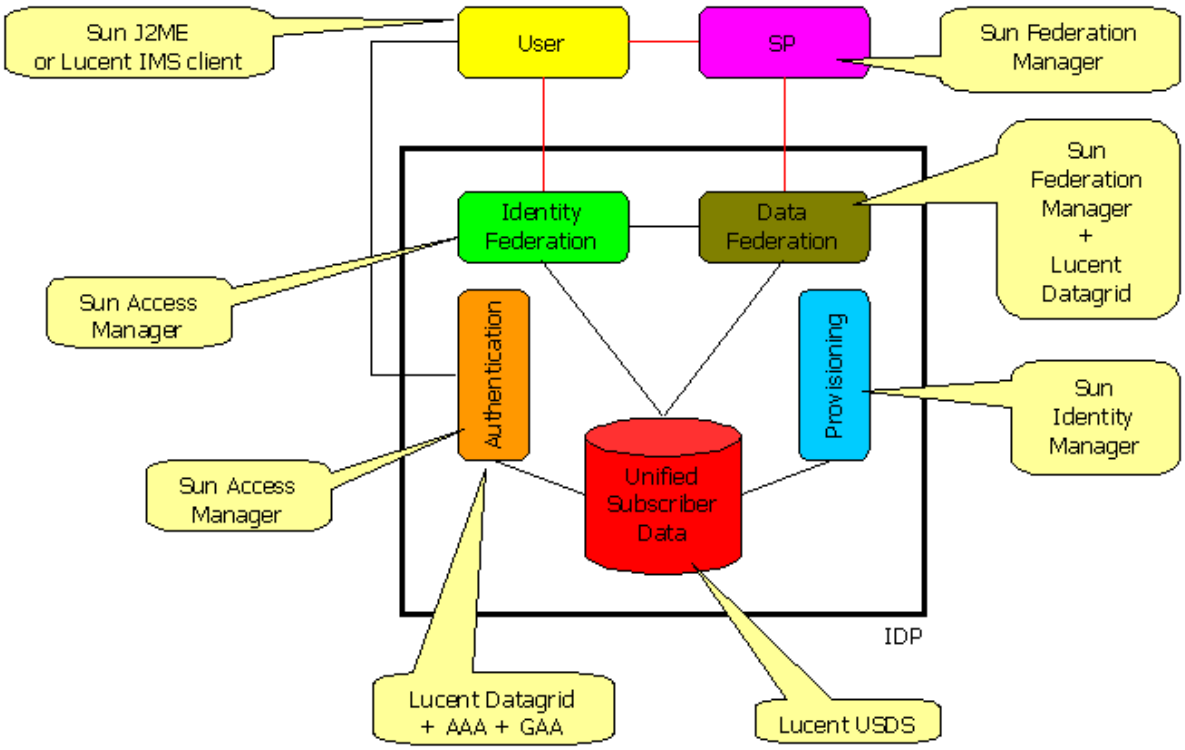

Figure 6.3-a: Lucent/Sun identity architecture.

Lucent and Sun assets complete the fundamental building blocks of the Identity ecosystem. The subscriber data will be hosted in the Lucent USDS platform. Lucent Datagrid along with Sun Federation Manager will provide data consolidation and federation capabilities. Sun Access Manager will perform Identify federation and also enable Datagrid and various Lucent AAA, GAA modules in Authentication and assertion.

7 USE CASES (VALIDATION)

In this section we describe various use cases that we are using to validate the design decisions presented above. For each of them, we present how they can be implemented using components from Lucent and Sun.

7.1 Use case #1: basic authentication and identity

For this use case, we construct an IdP using Sun Access Manager or Sun Federation as the front end and rely on either Datagrid or Lucent AAA server for the actual authentication. Depending on which method we choose, the interfaces will be different.

The two architectures are presented in the figure below.

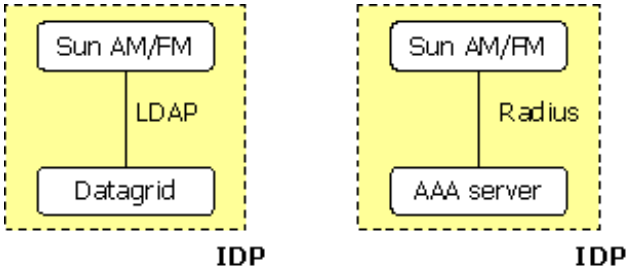

Figure 7.1-a: Basic authentication.

7.2 Use case #2: SMS weather service

A typical example of such a service is a "snow forecast" where a user can ask to receive "snow forecast" via SMS.

In this scenario, the user asks the snow forecast service to send her one (or more) SMS. The user contacts the service via WAP and authenticates with the operator's IdP. The IdP generates a pseudonym to identify the user to the SP. The SP will use this opaque identifier later to contact the SMS gateway to deliver the SMS to the user. During the whole interaction, the snow forecast service does not know anything about the user.

For meaningful snow forecast, the current location of the user is needed. This can be achieved in many ways.

7.2.1 Embedded location information

The location information is embedded in the SAML assertion generated by the IdP (using an attribute assertion). The IdP will use DataGrid to retrieve location information from the mobile network.

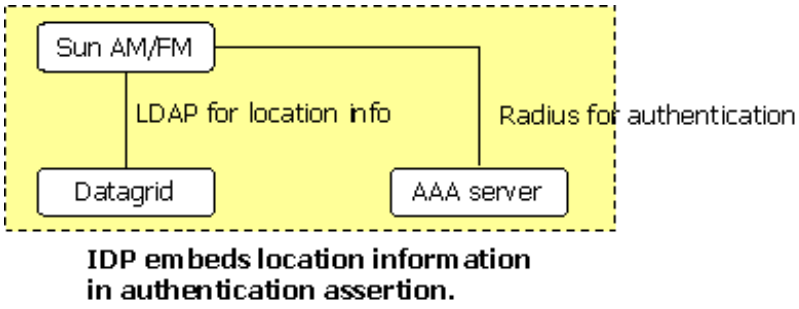

Figure 7.2-a

7.2.2 Location information retrieved separately

The location is not directly embedded in the original SAML assertion. The SMS service either has already stored some preferences about the user (anonymously identified by the pseudonym) because this is not the first the user accesses the service or the SMS service will ask the IdP (or a Data Service) for the location of the user using the pseudonym.

The architecture will look like:

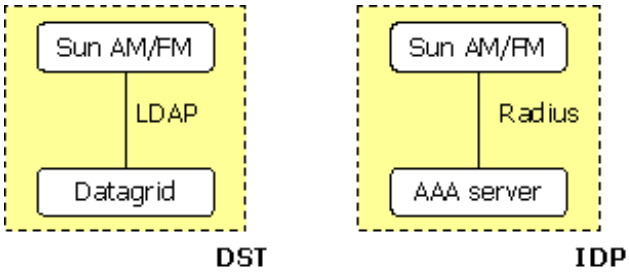

Figure 7.2-b: DST architecture

7.3 Use case #3: interaction service

The role of the interaction service (part of Liberty architecture) is to make it possible to reach the user when validation or authorization is required: e.g. "your current location is about to be disclosed to service provider X. Are you OK with that?"

By taking advantage of the information available from the network (e.g. list of devices owned by the user with their respective capabilities, presence, location, call status, etc.), the interaction service should be able to reach the user in real time.

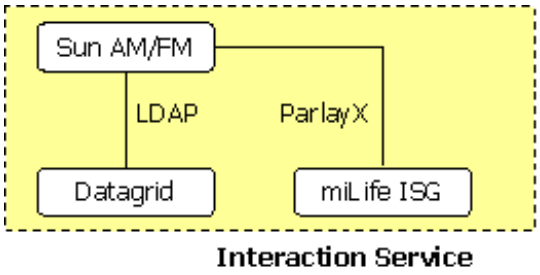

Figure 7.3-a: Interaction Service.

The various SPs (and even the IdP) may want to use the Interaction Service to contact the user. The Interaction Service will need to use a telecom application server to actually contact the user: e.g. ISG or IMS app server.

8 CONCLUSION

Identity management and Data Federation pose a big challenge and, at the same time, offer a tremendous opportunity for network operators.

A challenge first because we are trying to extend the benefits of federated identity and data to the anytime, anywhere experience of mobile devices, and provide to wireless networks the same advantages that may be available from a family set top box or a shared desktop in a cyber café. But also an opportunity because network operators can build on the special relationship they have with end users to build a new set of rich and personalized services.

In this whitepaper, we have tried to explain the rules of this new “game” and to show how the various players (end users, network operators and service providers) can all benefit from it.

We have also shown how the various aspects and roles of this game can be implemented using a mix of technologies from Lucent and Sun.

By combining the recognized expertise of both companies, we think we can bring to our customers the right building block products to help them create their own identity ecosystem, add value to their own customers and partners, and generate some new revenue.

9 GLOSSARY

End-user (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”, e.g., individual users, groups of individuals, organizational entities such as corporations, or a component of the Liberty architecture.

Service provider (SP)

A SP is typically a Liberty-enabled website providing services and/or goods, e.g., e-commerce site, SMS gateway service, weather forecast, horoscope, calendar, etc.

Identity provider (IdP)

An IdP is a Liberty-enabled system entity that manages identity information on behalf of Principals and provides assertions of Principal authentication to other providers.

SSO (single sign-on or simplified sign-on)

SSO enables a user to authenticate once and gain access to resources or services with no need for re-authentication. The motto could be ”authenticate once, be authenticated everywhere”.

Identity federation

During identity federation, two (or more) separate identities (from separate domains) are correlated (or associated) as pointing to the same principal.

Identity federation is a requirement for SSO.

The act of federating (and/or de-federating) identities is the responsibility of the user, i.e. the principal owning these identities.

Data federation

In data federation, data coming from multiple organizations is combined according to standard templates so that the data can be conveniently shared and passed around amongst the organizations’ websites/servers

Identity-based service

An identity service is a web service whose operations are indexed by identity, i.e. invokable and discoverable based on the user identity, e.g., calendar, presence, location, personal profile.

Circle of trust (CoT)

A CoT is a grouping of service providers and identity providers who have business relationships and operational agreements based on the Liberty architecture.

End users can experience the benefits of SSO inside a given CoT.

The existence of a CoT is a requirement for SSO.

Authentication assertion

An authentication assertion is issued by an IdP as a proof of authentication for a given user. This assertion is usually targeted to the SP that requested the user authentication in the first place.

An authentication assertion usually contains information such as: the identity of the user, the context of authentication, the validity, etc. It can be extended with profile information about the user: e.g. location information, gender, etc.

An authentication assertion is represented to a SAML assertion, i.e. a signed XML document compliant with the <AuthenticationAssertion> SAML element.

Authentication assertions are typically transported over HTTP messages.

Authentication artifact

This is a special kind by an authentication assertion for cases when bandwidth is limited and the verbosity of assertions is not acceptable.

An artifact is a pointer to an assertion that needs to be de-referenced by the SP by sending an artifact resolution request to the IdP.

Pseudonym

In Liberty, user identities are almost never represented using real user identities (e.g. #SS, login, phone number), for obvious privacy reasons.

As defined by Liberty, a pseudonym is ”an arbitrary identifier assigned by the identity or service provider to identify a Principal to a given relying party so that the name has meaning only in the context of the relationship between the parties.”

Note that the user does not necessarily need to be aware of the pseudonym(s) used to represent her.

Authentication context

As its name implies an authentication context represents the circumstances and the details of an authentication: e.g. type of authentication, time, expiration time, etc. Such a context is critical for SSO because services do not have the same security requirements. A low security authentication cannot be reused by SSO for a service that requires a high level of security.

Session

A session is a time interval (with a beginning and an end) during which the entities that are part of the session share a trust relationship. When a user authenticates to an IdP, a session may be established. If the user comes back to the IdP within the limits of the session, she does not need to authenticate again.

Sessions can be terminated by one (or more) party(ies) from the session or due to timeouts.

10 ADDITIONAL RESOURCES

Specifications

[1] Liberty Alliance specifications

[2] SAML 2.0

[3] WS-Federation

[4] 3GPP Generic User Profile (GUP) (stage 3)

[5] 3GPP Generic Authentication Architecture (GAA)

Proposed and/or deployed identity systems

[6] Microsoft Passport

[7] Shibloeth

[8] LID

[9] openID

[10] Passel

[11] SXIP

[12] YADIS

Press articles and white papers

[13] Identity in Action. Oct 2005. InfoWorld.

[14] Identity Management. March 2004. The Open Group.

[15] Identity Management: Time for Action. July 2005. Ovum.

[16] Overview And Netegrity Support For Liberty Alliance Functionality. May 2004. Netegrity Technical White Paper.

[17] Deploying Mobile Web Services with Liberty's ID-WSF. June 2004. . Sun, Nokia.

Research papers

[18] The Laws of identity. Kim Cameron. May 2005.

[19] Analysis of Liberty Single-Sign-on with Enabled Clients . Birgit Pfitzmann and Michael Waidner, IBM Zurich Research. IEEE Internet Computing, Nov-Dec 2003.

[20] On Adaptive Identity Management: The Next Generation of Identity Management. Marco Casassa Mont, Pete Bramhall, Joe Pato. July 2003. Technical report HPL-2003-149.

[21] Identity Management: Setting Context. Joseph Pato. April 2003. Technical report HPL-2003-72.

Success stories

[22] A GSM Operators Service Network and the Subscribers Identity

Other valuable resources

[23] The Identity Corner blog

[24] Discovering Identity blog

[25] Burton Group page on identity management