ETSI GS INS 006 V1.1.1 (2011-11)

Identity and access management for Networks and Services; Study to Identify the need for a Global, Distributed Discovery Mechanism

Disclaimer

This document has been produced and approved by the Identity and access management for Networks and Services (INS) ETSI Industry Specification Group (ISG) and represents the views of those members who participated in this ISG.
It does not necessarily represent the views of the entire ETSI membership.

Reference
DGS/INS-006

Keywords
access, control, ID, management,
network, service

ETSI

650 Route des Lucioles
F-06921 Sophia Antipolis Cedex - FRANCE

Tel.: +33 4 92 94 42 00 Fax: +33 4 93 65 47 16

Siret N° 348 623 562 00017 - NAF 742 C
Association à but non lucratif enregistrée à la
Sous-Préfecture de Grasse (06) N° 7803/88

Important notice

Individual copies of the present document can be downloaded from:
http://www.etsi.org

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be the printing on ETSI printers of the PDF version kept on a specific network drive within ETSI Secretariat.

Users of the present document should be aware that the document may be subject to revision or change of status.
Information on the current status of this and other ETSI documents is available at
http://portal.etsi.org/tb/status/status.asp

If you find errors in the present document, please send your comment to one of the following services:
http://portal.etsi.org/chaircor/ETSI_support.asp

Contents

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to ETSI. The information pertaining to these essential IPRs, if any, is publicly available for **ETSI members and non-members**, and can be found in ETSI SR 000 314: *"Intellectual Property Rights (IPRs); Essential, or potentially Essential, IPRs notified to ETSI in respect of ETSI standards"*, which is available from the ETSI Secretariat. Latest updates are available on the ETSI Web server (http://ipr.etsi.org).

Pursuant to the ETSI IPR Policy, no investigation, including IPR searches, has been carried out by ETSI. No guarantee can be given as to the existence of other IPRs not referenced in ETSI SR 000 314 (or the updates on the ETSI Web server) which are, or may be, or may become, essential to the present document.

Foreword

This Group Specification (GS) has been produced by ETSI Industry Specification (ISG) Identity and access management for Networks and Services (INS).

Introduction

Today, discovery of identity data across domains is generally realized with two different ways.

- Discovery Service (DS): A service defined by a group of network entities (providers) which participate in a federation. Identity data (actual data or mappings) are registered in the service and can be provided to all the participants of the group. The location of the discovery service and the protocol for exchanging messages is static and known to the participants of the group (federated model).

- The "user@location" format: By using an identifier of this format, a user directly points to a network point that holds identity information about him (user-centric model). This location may hold information for only one profile of the user (id = email) or for many profiles (id = Virtual Identity [i.1]).

However both of the above ways provide limited discovery of user's identity information. For the federated model, only the identity data which exist within the federation of providers can be discovered (and-or associated). Information outside the federation cannot be discovered. Providers that participate in the federation, have previous knowledge of the location of the DS (where to ask for information), and how to exchanged data with it (how to ask for information). Efforts to locate data outside predefined federations are usually hampered by the proprietary design of the discovery services and the customized identity formats and protocols that each federation uses. For the User-centric model the use of a specific predefined format instantly excludes the discovery of identity data from providers that are not familiar with it. Even though the adoption of a globally accepted identifier would solve major identity issues, this seems to be inapplicable mainly for business reasons and severe protocol modifications in various networks and technologies.

This work item assumes that all data and attributes required to provide a service are not available within a single service provider. For example proof of residence is required to access online streaming services. An acceptable issuer of this attribute may not be known to the streaming services' provider beforehand and must be discovered.

The purpose of the present document is to investigate the current landscape on the IdM area and evaluate if there is a need for such a discovery mechanism, or whether this can be covered by existing solutions.

1 Scope

The present document will focus on gap analysis for global distributed discovery mechanism of identifiers, providers and capabilities.

2 References

References are either specific (identified by date of publication and/or edition number or version number) or non-specific. For specific references, only the cited version applies. For non-specific references, the latest version of the reference document (including any amendments) applies.

Referenced documents which are not found to be publicly available in the expected location might be found at http://docbox.etsi.org/Reference.

NOTE: While any hyperlinks included in this clause were valid at the time of publication, ETSI cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are necessary for the application of the present document.

Not applicable.

2.2 Informative references

The following referenced documents are not necessary for the application of the present document but they assist the user with regard to a particular subject area.

[i.1] Designing Advanced network Interfaces for the Delivery and Administration of Location independent, Optimised personal Services.

NOTE: See http://www.ist-daidalos.org.

[i.2] Libery Alliance Project, Projectliberty.

NOTE: See http://www.projectliberty.org.

[i.3] Kantara InitiativeTM, Shaping the Future of Digital Identity.

NOTE: See http://kantarainitiative.org.

[i.4] Libery Alliance Project, Liberty ID-WSF Discovery Service Specification.

NOTE: See http://projectliberty.org/liberty/content/download/3450/22976/file/liberty-idwsf-disco-svc-v2.0-original.pdf.

[i.5] Internet2 Middleware Initiative, Shibboleth®.

NOTE: See http://shibboleth.internet2.edu.

[i.6] DiscoveryService.

NOTE: See https://wiki.shibboleth.net/confluence/display/SHIB2/DiscoveryService.

[i.7] Eduserv, OpenAthens.

NOTE: See http://www.openathens.net.

[i.8] Microsoft Windows CardSpace™.

NOTE: See http://windows.microsoft.com/en-US/windows-vista/Windows-CardSpace.

[i.9] Higgins, Personal Data Service.

NOTE: See http://www.eclipse.org/higgins.

[i.10] OpenID®.

NOTE: See http://openid.net.

[i.11] XDI.org.

NOTE: See http://www.xdi.org.

[i.12] OASIS, Extensible Resource Identifier (XRI).

NOTE: See http://www.oasis-open.org/committees/download.php/15377.

[i.13] OASIS, Extensible Resource Identifier (XRI) Resolution Version 2.0.

NOTE: See http://docs.oasis-open.org/xri/2.0/specs/xri-resolution-V2.0.html.

[i.14] SWIFT.

NOTE: See http://www.ist-swift.org.

[i.15] STORK Project.

NOTE: See https://www.eid-stork.eu.

[i.16] M. Dabrowski, P. Pacyna, "Cross-Identifier Domain Discovery Service for Unrelated User Identities", DIM Workshop, 2008.

[i.17] Wikipedia, Domain Name System.

NOTE: See http://en.wikipedia.org/wiki/Domain_Name_System.

[i.18] Wikipedia, Dynamic DNS.

NOTE: See http://en.wikipedia.org/wiki/Dynamic_DNS.

[i.19] DNSSEC: DNS Security Extensions.

NOTE: See http://www.dnssec.net/.

[i.20] Wikipedia, DNS cache poisoning.

NOTE: See http://en.wikipedia.org/wiki/DNS_cache_poisoning.

[i.21] Handle System®.

NOTE: See http://www.handle.net.

[i.22] IF-MAP.com.

NOTE: See http://www.if-map.com.

[i.23] Jon Crowcroft, Steven Hand, Richard Mortier, Timothy Roscoe, Andrew Warfield, "Plutarch: An Argument for Network Pluralism", ACM SIGCOMM, 2003.

3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

DDNS	Dynamic DNS
DNS	Domain Name System
DNSSEC	DNS Security Extension
DS	Discovery Service
EU	European Union
GHR	Global Handle Registry
IANA	Internet Assigned Numbers Authority
IdM	Identity Management
IdP	Identity Provider
ID-WSF	Identity Web Services Framework
IF	Interstitial Function
IF-MAP	Interface for Metadata Access Points
IM	Instant Messaging
IP	Internet Protocol
ISG	Industry Specification Group
LHS	Local Handle Registry
MAC	Media Access Control
OASIS	Organization for the Advancement of Structured Information Standards
OWL	Web Ontology Language
SAML	Security Assertion Markup Language
SMS	Short Message Service
SP	Service Provider
SSO	Single Sign On
STORK	Secure IdenTity AcrOss BoRders LinKed
TCG	Trusted Computing Group
TNC	Trusted Network Connect
UDP	User Datagram Protocol
URL	Uniform Resource Locator
VID	Virtual Identity
WLAN	Wireless Local Area Network
WS	Web Service
WSC	Web Service Consumer
WSP	Web Service Producer
XDI	XRI Data Interchange
XML	EXtensible Markup Language
XRDS	EXtensible Resource Descriptor Sequence
XRI	EXtensible Resource Identifier

4 Scenarios

The vast majority of existing identity management systems can provide identity solutions only if specific requirements are met. Such requirements may be that, during a network operation, all participants trust each other, have established common protocols and formats, share or know where to find the desired information, etc. These assumptions however do not always apply and in some cases the desired identity information does not exist in places where the IdM systems presume. For these cases the interested party may need to dynamically discover and acquire the desired data.

Figure 1 illustrates a situation where a party needs to dynamically discovery identity information about a user. User "user@webstore.com" logs in on a service provider "webstore.com" with a pre-registered account and requests a specific service (e.g. an online purchase). In order to complete the transaction, provider "webstore.com" must contact other providers (e.g. payment.com, supplier.com etc) which all participate in Federation A. Among them, the "supplier.com" provider needs to validate user's age against a trusted entity before completing its part of the service. This information however does not exist in any of the providers forming the Federation A but resides in the organization "gov.com" which is member of the Federation B.

Even though "Supplier.com" and "gov.com" participate in a common federation (Federation B), existing literature fails to support the above operation. "Supplier.com" cannot autonomously discover where the desired information resides (is restricted to use only information that exist in Federation A and are associated with the username "user@webstore.com") and even if it somehow knew that the desired information existed in the "gov.com" domain, the username "user@webstore.com" means nothing to "gov.com".

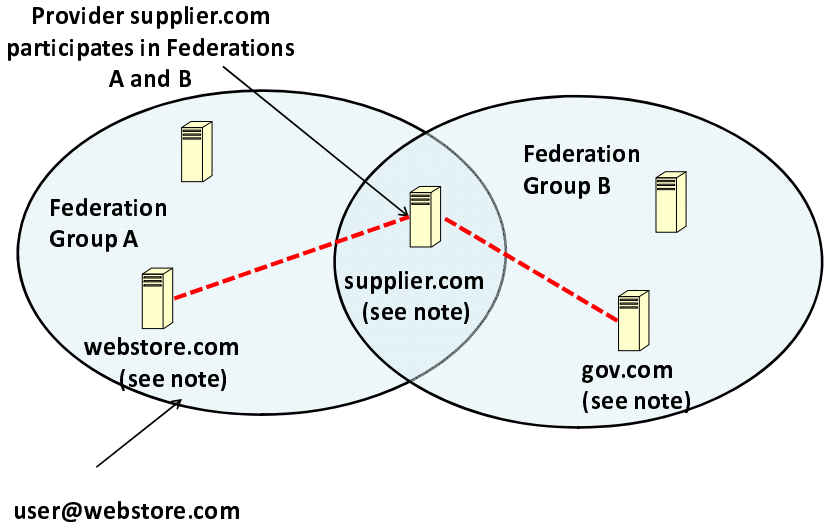

NOTE: The end user has a registered account with this provider.

Figure 1: Locating identity information about a user outside a Federation

Figure 2 presents a second scenario where a network entity (e.g. a WLAN hotspot provider) needs to dynamically discover identity data about a mobile user in a roaming setting. The user switches on his laptop and connects to a public Hot Spot. He opens his web browser and is redirected to a predefined web page (HotSpot Provider's web page) to be authenticated. If the user does not have an active account with the HotSpot Provider, existing practices provide the means to authenticate him through a collaborative party (trusted 3rd party). The selection of this collaborative party is usually held through a static procedure, mainly by asking the end user to choose among a predefined set of parties (a procedure known as WAYF - "Where Are You From"). Each case though, assumes different security requirements, policy agreements, trust levels etc which inexperienced users may not judge correctly during their selection. Furthermore the predefined set of collaborative parties may not always include the most appropriate provider for a specific given case, or include providers which are unknown to the roaming user.

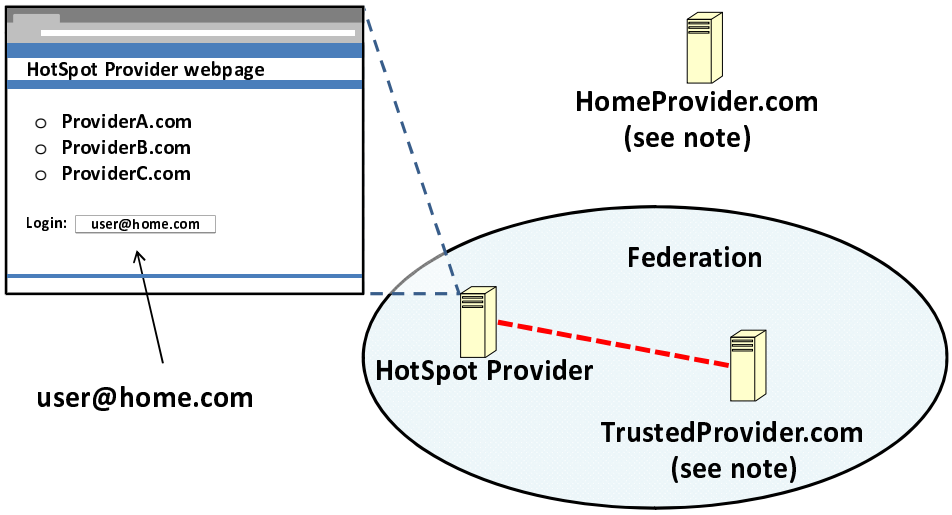

NOTE: The end user has a registered account with this provider.

Figure 2: Locating the best source to authenticate a user

The end user does not have an account at any of the providers displayed on the HotSpot Provider's webpage. Not knowing what username to select, the user submits a unique username (user@home.com) issued by his Home Provider ("HomeProvider.com"). The HotSpot Provider however, does not know the "HomeProvider.com" - or does not trust it enough to complete this specific service- and may thus need to dynamically discover information about the end user from alternative trusted locations. Such a location can be the provider "TrustedProvider.com".

It must be noted that:

- the only information available to the "HotSpot Provider", about the user, is the submitted username (user@home.com);
- the mobile user has a registered account at "TrustedProvider.com".

5 Current landscape

This clause describes the existing IdM systems and discovery mechanisms and evaluates their ability to provide identity discovery services across different administrative or technological contexts (domains, federations, protocols, formats, etc.).

5.1 Federated Identity Management Frameworks

Liberty Alliance [i.2] group has proposed Liberty Federation, a framework for federated identity management. Based on its specifications, OASIS formed the Security Assertion Markup Language (SAML) 2.0, an XML based standard for data exchange between Identity Provider and Service Providers. It must be noted that Kantara Initiative [i.3] is the evolution of Liberty Alliance and has adopted all its work and related materials.

The Discovery Service (DS) used by these initiatives is the ID-WSF Discovery Service [i.4]. Based on it, a Web Service Consumer (WSC) can find a relevant Web Service Producer (WSP) associated with a particular user identity. The DS facilitates the WSCs to discover where on the network the user's different identity attributes are located (WSP endpoint) and also enables the WSC to actually send a request for that identity to the endpoint, by issuing to the WSC an identity token to be included in the request. To be able to discover a Web service, it must initially be registered. The discovery service matches available registered services to a lookup request coming from a WSC. This request describes the desired services both in terms of functional criteria as well as ownership. Discovery has special requirements when it handles identity related data. ID-WSF V2.0 defines the following components in support:

- A data model to represent the available web services associated with an identity.

- An interface to allow WSCs to retrieve from a Discovery Service a list of web services associated with an identity based on various criteria, along with other information needed for invoking the service.
- An interface to allow WSPs to register and manage their resources (web services) at the Discovery Service.
- A model to bootstrap the service discovery process by including required information about the Discovery Service itself.

Shibboleth [i.5] is an Internet2 Middleware Initiative project that designed and implemented an open-source identity framework for authentication and authorization for federated identities. Using federated identities, domains which reside in a federation can exchange identity data and information about their users, thus achieving cross-domain single sign-on without the need for content providers to store usernames and passwords. Users' information is supplied by Identity Providers (IdPs) to the Service Providers (SPs) which want access to secure content. Shibboleth is based on Security Assertion Markup Language (SAML).

IdP discovery in the initial versions of Shibboleth [i.6] was performed by asking the user directly using the "Where Are You From (WAYF)" service. Discovery of the appropriate IdP was either held using a flat, static page that relied on a fixed, known set of possible IdP's, or through a dynamic discovery service, which was a separate app that could generate a set of options based on metadata, present those options to the user, and send the user to the selected home.

- Flat Page Discovery: Static HTML is often sufficient for discovery when there is a limited and static set of IdP's to choose from, and is simple and clean to implement.
- Discovery Service: A DS presents, highly customizable, a set of IdP's from which the user can choose. After the user makes a selection, the DS redirects the user to the SP, which then formulates the AuthnRequest based on the user's selection.

It was soon realized that despite its simplicity, the WAYF architecture was not flexible and could not always provide the best IdP for a given situation. Users were not always fully aware of what they were asked for, or were not always familiar with terms like federation, IdP etc. Furthermore independent services were not always well positioned, resulting in the formation of a confusing list of IdPs to the user. To minimize these issues a new architecture, which associates a discovery service with each SP, was adopted. With this approach the discovery service can provide assistance in selecting the IdP in a way which is targeted at its particular users base and also limits the IdPs presented to the user, only to those which the SP is willing to accept authentication from.

Athens [i.7] is an Access and Identity Management service supplied by Eduserv which provides high levels of user management capabilities and also single sign-on services to protected resources. Various versions of Athens exist, with different security options. Athens main concept is the replacement of the multiple usernames and passwords necessary to access subscription based content with one single username and password that can be entered once per session. Athens discovery service is also based on the WAYF service.

Federated IdM Systems create independent DS (often proprietary) which can only operate within the borders of a federation. Participating parties have previous knowledge of the location of this DS (where to ask for information) and the formats and protocols it uses (how to ask for information). Multiple identifiers of various formats can be registered and processed, as long as they belong to providers which participate in the federation. Identity information and formats outside the Federation cannot be handled. Any efforts for inter-federation support usually result in the formation of larger federations again able to process only the information inside the new borders.

5.2 User-Centric Identity Management Frameworks

Microsoft CardSpace [i.8], is a Microsoft-initiated solution based on WS-* specifications. It is a user-centric mechanism where users can create, delete, and modify various identity profiles, known as Information Cards, thus control the kind and amount of information revealed in the network.

Project Higgins [i.9] is another user-centric identity framework which is based on SAML. This system unifies all identity interactions across multiple heterogeneous systems through a common user interface metaphor also based on Information Cards (i-cards). It provides to the end users a single point of control over multiple heterogeneous identities preferences and relationships. The represented identities (Digital Subjects) and their Identity Attributes are exposed in a Context through a data model, described by OWL.

According to the OpenID [i.10] proposed solution, when a user contacts an OpenID-enabled web site he inserts a URL instead of his username. The OpenID site redirects the user to a site that corresponds to the submitted URL, which in turn, performs the user login operation. OpenID is fast becoming the de-facto solution for secure login in the internet as it is user-friendly, user-centric and supports features like Single Sign On (SSO).

The OASIS XRI Data Interchange (XDI) [i.11] is an effort to define a generalized, extensible service for sharing, linking, and synchronizing data over the Internet and other data networks using XML documents and XRIs (Extensible Resource Identifiers). The Extensible Resource Identifier (XRI) provides uniform syntax for abstract structured identifiers as defined in [i.12]. Since XRIs can be used across a wide variety of communities and applications (as Web addresses, database keys, filenames, object IDs, XML IDs, tags, etc.), no single resolution mechanism may prove appropriate for all XRIs. XRDS (Extensible Resource Descriptor Sequence) has been introduced as a simple generic resource description format for discovery of metadata about a resource. Specification [i.13] describes a standard protocol for requesting XRDS documents using HTTP(S) URIs, and a standard protocol for resolving XRIs using XRDS documents and HTTP(S) URIs. XRI resolution follows the DNS architecture except at a higher level of abstraction, i.e., rather than using UDP to resolve a domain name into a text-based resource descriptor, it uses HTTP(S) to resolve an XRI into an XML-based resource descriptor called an XRDS document.

Besides XRI resolution, examples of typical XRDS usage include:

- OpenID authentication for discovery and capabilities description of OpenID providers.
- OAuth discovery for locating OAuth service endpoints and capabilities.
- The Higgins Project for discovery of Higgins context providers.
- XDI.org I-name and I-number digital identity addressing services for generalized digital identity service discovery.
- The XDI data sharing protocol for discovery of XDI service endpoints and capabilities.

DAIDALOS Project [i.1] proposes a cross layer identity management system based on the management of all different profiles a user may have in the network. These profiles are linked together into multiple groups, creating the so called Virtual Identities (VIDs). Two or more VIDs cannot be associated with each other, providing strong privacy and security. With virtual identities a user can create multiple personae and gain complete control over his private information. Project SWIFT [i.14], a European Union funded project started in 2008, is also based on VIDs.

STORK [i.15] is another EU funded project which aims at designing a European eID Interoperability Platform that will allow citizens to establish new e-relations across borders, just by presenting their national eID. The eID is again a new type of identity introduced by STORK.

In [i.16] authors acknowledge the need for a widely deployed IdM discovery service that can encompass identities of different layers, formats and business areas. This work describes a valid large scale discovery mechanism which however assumes the existence of a globally trusted organization to produce and distribute the necessary public and private keys of their framework.

The identity management frameworks which follow the user-centic model tend to introduce new types of global identifiers which must be adopted by everyone, in their effort to support different contexts and technologies. These systems do not introduce any kind of discovery mechanism since the proposed global identifiers follow the format "global_id@location" and directly point to the location of all the data which are associated with them. The enforcement of one specific "global identifier" until today is rejected mainly for business and severe protocol modifications. This is a situation which is not expected to change. A more realistic approach would be the creation of a mechanism capable of handling various types of identifiers (existing and new ones) without affecting their format or interfearing with the issuers' internal network functionality.

5.3 Discovery Frameworks

5.3.1 DNS, DDNS, DNSSEC

The Domain Name System (DNS) [i.17] is a system which holds information about the address of each "zone" or site that can be accessed on the Internet. It follows a hierarchical architecture with each level of this hierarchy being "authoritative" for the level just below it. For instance, the "Root" that is maintained by the IANA (Internet Assigned Numbers Authority) contains information about all of the Top-Level domain - TLDs (.com, .org, .edu, .info, etc.) and who is responsible for them. The .edu registry contains information about all .edu Web sites.

The Dynamic DNS (DDNS) [i.18] is a network service that provides the capability for a networked device, such as a router or computer system, to notify a DNS name server to change, in real time, the active DNS configuration of its configured hostnames, addresses or other information. The most common usage of DDNS is to provide a residential user's Internet gateway that has a variable, often changing, IP address with a well known hostname resolvable through standard DNS queries.

The DNS Security Extensions (DNSSEC) [i.19] has been introduced as a security measure to protect the DNS from cache poisoning exploits (recently referred to as the Kaminsky bug [i.20]) which can allow an attacker to get in the middle of an Internet users' request to access a Web site without their knowledge. During a DNS request process only simple yes/no responses are given and there is no means to authenticate the entity claiming to own the address of the requested resource. DNSSEC introduces digital signatures to the DNS infrastructure, allowing end users to more securely navigate the Internet. The resolution process is secured by establishing a "chain of trust" that effectively asks for a secret password or "key" in order to exchange the information at each level.

By definition the DNS system was designed to provide a specific functionality: the translation of domain names to IP addresses. Modifying its existing functionality or exploiting its infrastructure to support discovery of identity related information will introduce high risks in term of security, privacy and functionality, since both systems have diverse requirements and operations.

5.3.2 HANDLE

The Handle System [i.21] is a technology specification for assigning, managing, and resolving persistent identifiers for digital objects and other resources on the Internet. It includes an open set of protocols, a namespace, and a reference implementation of the protocols. The protocols enable a distributed computer system to store identifiers, known as handles, of arbitrary resources and resolve those handles into the information necessary to locate, access, contact, authenticate, or otherwise make use of the resources. This information can be changed as needed to reflect the current state of the identified resource without changing its identifier, thus allowing the name of the item to persist over changes of location and other related state information.

The interoperable network is formed by distributed handle resolver servers (also known as the Proxy Server System) linked through a Global Resolver (which is one logical entity though physically decentralised and mirrored). Users of Handle System technology obtain a handle prefix created in the Global Handle Registry. The Global Handle Registry maintains and resolves the prefixes of locally-maintained handle services. Any local handle service can, therefore, resolve any handle through the Global Resolver. Handles are passed by a client, as a query of the naming authority/prefix, to the Handle System's Global Handle Registry (GHR). The GHR responds by sending the client the location information for the relevant Local Handle Service (LHS) (which may consist of multiple servers in multiple sites); a query is then sent to the relevant server within the Local Handle Service. The Local Handle Service returns the information needed to acquire the resource, e.g. a URL which can then be turned into an HTTP re-direct.

NOTE: If the client already has information on the appropriate LHS to query, the initial query to GHR is omitted.

Handle is a large scale system which is able to locate a specific object irrespectively of its network location. The system presumes that the requester already knows a persistent identifier about this object. Associating various identifiers to a unique object (e.g. associating multiple identities to an user) is not examined.

5.3.3 IF-MAP

IF-MAP is an initiative of the TCG's (Trusted Computing Group) Trusted Network Connect (TNC) subgroup [i.22], to standardize the way devices and applications share information with one another. It does for coordination and collaboration what IP has done for connectivity. It defines a client/server protocol and the server side is a database that supports publish, subscribe and search operations. The IF-MAP server can dynamically associate any kind of information, taken from disparate and potentially unrelated sources, about an object (user, machine, sensor etc) and create a comprehensive view about it. Such information may for example be: network location (IP), hardware ID (MAC address), session information, priviledges, etc.

There is no pre-defined global structure for an IF-MAP database; rather, the global database structure (schema) emerges as each application and system publishes information to the IF-MAP service. Systems compatible with IF-MAP can subscribe to changes in data of interest - such as a new device coming onto the network, or a user changing role, or an item moving from one location to another - and receive updates automatically.

IF-MAP system has the ability to locate various information about an object. Various data can be associated with one object. However the discovery process assumes that the requester already knows where to subscribe to the information of interest. Thus dynamic discovery of an object or information without previous knowledge of the IF-MAP server is not possible.

5.3.4 Plutarch

Plutarch [i.23] is an inter-networking framework designed to address the heterogeneity of the next generation networks. It is a new architecture for the future internet which does not rejects the existing "homogeneous" internet but embraces it as one architecture among many. It divides the world into contexts, each comprising some set of hosts, routers, switches, network links etc. Each context is defined by a certain level of homogeneity among specific concepts like addresses, packet formats, transport protocols and naming services. Distinct contexts differ in at least one of these areas. An "Interstitial Function (IF)" acts as a bridge and allows data to pass between two adjoining contexts. Within the Plutarch system, communication takes place between endpoints within contexts and may travel through more than one IF.

In the Plutarch system there are no global names, and the name resolution mechanism involves the discovery of the appropriate context in which the target host exists. This procedure is held through queries which are based on epidemic-style gossip advertisements and queries across contexts, and the result may be more than one candidate contexts.

This system is able to support communication among heterogeneous contexts and also compatible with existing networks and technologies. Its purpose however is not identity management thus important aspects like privacy (e.g. its discovery mechanism advertises queries to multiple places) or association of multiple elements (e.g. identities) are not examined.

6 Use Cases

6.1 UC1: User's identity data are scattered across unassociated administrative domains

6.1.1 Description

User "user@webstore.com" logs on his account in provider "webstore.com" and requests a specific service (e.g. an online purchase). In order to complete the transaction, provider "webstore.com" must contact other providers (e.g. paypal.com, supplier.com etc) which all participate in the Federation A. The "supplier.com" provider though needs to validate user's age against an trusted entity. The requested information resides in the organization "gov.com" which does not participate in Federation A. "Supplier.com" and "gov.com" participate in a different Federation (Federation B).

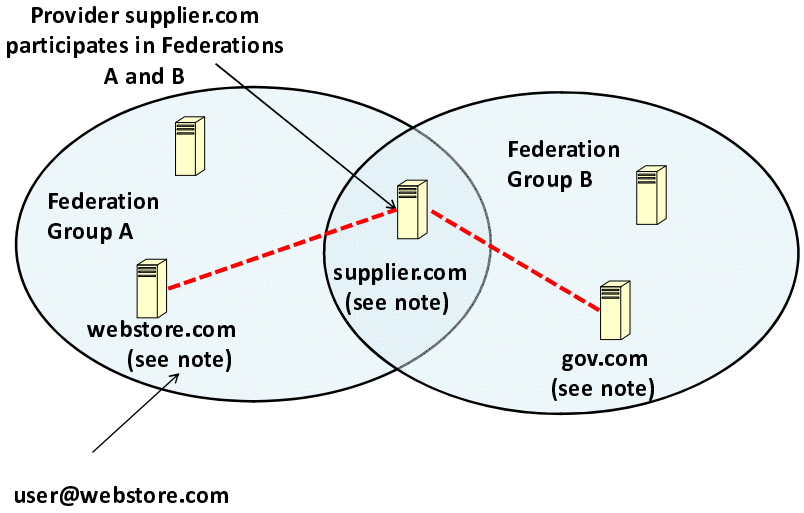

NOTE: The end user has a registered account with this provider.

Figure 3: User's scattered identity data

6.1.2 Actors

- User who wants to complete an online purchase.
- Domain "webstore.com" which provides the interface to a user to complete an online purchase service.
- Domain "supplier.com" which provides a part of the purchase service (e.g. product delivery).
- Domain "gov.com" which holds identity information about users of a specific geographical area.

6.1.2.1 Actors specific Issues

- User:
  - Wants to buy a product online.
- Webstore.com:
  - Provides the interface to complete the online service (purchase).
  - Hides all the background operations which take place in order to complete the service.
  - Collaborates with multiple other domains to complete the service.
- Supplier.com:
  - Needs to validate the consumer's age before completing his part of the purchase transaction.
  - Cannot directly contact the end-user.
  - Participates in multiple Federations and has established trust and procedures with all of them.
- Gov.com:
  - Holds specific identity data for users in a specific geographical area.
  - Is a distinguished organization and is trusted in the specific geographical area.

6.1.2.2 Identified gaps

- User:
  - Cannot complete the online purchase.
- Webstore.com:
  - Cannot complete the service.
  - Cannot provide all the required identity data to a collaborative domain.
- Supplier.com:
  - Cannot autonomously discover where the desired information resides since the identifier "user@webstore.com" is only valid within Federation A (is restricted to use only information that exists in Federation A).
  - Cannot use the identifier "user@webstore.com" outside Federation A.
- Gov.com:
  - Cannot identify the owner of the username "user@webstore.com" .

6.1.2.3 Alternative Solutions based on existing literature

- Collection and Storage of all the necessary identity information:
  - Federation A asks from all the participating providers to indicate all the identity data they need to support all the services which are provided in the federation.
    - The required data are located, collected and stored in the Federation. This solution is not feasible since the owners of these identity data (user, gov.com) may not agree that all the participants of a specific federation have access to all of their data. The fact that a specific owner (gov.com) trusts a specific provider (supplier.com) in a federation does not necessarily mean that it trusts all the providers in the federation.
  - New interfaces are constructed between the Federation and the owners of the required identity data.
    - The number of the required interfaces will be very (extremely) large, especially each time the Federation expands and incorporates additional service providers.
    - Except from the technical details, it must be noted that each new interface between a federation and the owners of an identity information, is also a new trust relation. A specific owner (gov.com) can trust a specific provider (supplier.com) but may be reluctant to trust a federation (group of many providers).
- Federation A can ask from the user to provide the required information:
  - Webstore.com may ask the missing information.
    - An intermediate entity (webstore.com) must not be able to learn irrelevant private information, about a user. In this case webstore.com learns that the user has a relation with gov.com.
  - The user is redirected to supplier.com to provide the required information.
    - This practice is not Single Sign On.

6.2 UC2: Unknown user authentication

6.2.1 Description

A user wants to access a service, provided by a Service Provider. The user does not have an active account with the Service Provider. The user is asked to provide information in order to authenticate himself with the help of a collaborative party (trusted 3rd party). A list of collaborative providers may or may not be presented to the end user in Service Provider's web page.

6.2.2 Actors

- User who wants to access a service provided by an unknown Service Provider.
- Service Provider which provides an online service which is accessible only to authenticated users.
- Collaborative parties (trusted 3rd party) which can authenticate the end user.

6.2.2.1 Actors specific Issues

- User:
  - Wants to acquire a service (e.g. network connectivity).
- Service Provider:
  - Wants to identify and authenticate the unknown user.
  - Sustains preconfigured trust and service agreements with a number of other domains (providers, organizations etc).
- Collaborative parties:
  - Sustain preconfigured agreements and have established trust with the Service Provider.

6.2.2.2 Identified gaps

- User:
  - A list of providers (collaborative 3rd parties) is presented to the user.
    - A user may not always know the very specific technical requirements of the service and select the best Identity Provider to authenticate him to the Service Provider.
    - A user cannot know the trust relations of the foreign Service Provider and his Identity Providers, to select the personae which satisfies the trust requirements for his authentication.
    - The user may not sustain registered account in any of the presented providers.
- A list of providers is not presented to the user:
  - The user cannot know which one of his Identity Providers can be trusted by the foreign Service Provider. Thus the user cannot select and submit the appropriate identifier.
  - Inexperienced users. Phishing attacks: An inexperienced user may be tricked to provide private information to a malicious entity.
  - The user (especially the inexperienced one) cannot always judge the exact amount of information which must be revealed to acquire a service.
- Service Provider:
  - Cannot autonomously locate the exact information it needs to identify the end user - is limited to use only the information provided by the foreign end user.

6.3 UC3: Contacting an offline user

6.3.1 Description

User John is registered with an Instant Messaging (IM) provider (im.com), using its customized application to send messages to his friends. He decides to send an urgent message to his friend Nick (Nick@im.com), who also has a registered profile in the same IM provider but he is offline at the moment. John has no knowledge of any other way to contact Nick. Nick, on the other hand, always carries with him his mobile phone and has the ability to receive the instant message as an SMS. But he cannot know that someone wants to send him an urgent message.

6.3.2 Actors

- User John who wants to send an urgent message.
- User Nick who is registered to the IM service but is not currently connected.
- The IM provider which can transform and forward IM to other networks and technologies.

6.3.2.1 Actors specific Issues

- User John:
  - Wants to send an urgent message to a user who is currently offline in a specific service.
- User Nick:
  - Is not currently connected to the IM service.
  - Permits the delivery of urgent messages.
  - Carries with him his mobile phone and can always be contacted.
- IM provider:
  - Can transform instant messages to other types of messages and forward them to many types of networks.

6.3.2.2 Identified gaps

- User John:
  - Cannot send to user Nick an urgent message.
- User Nick:
  - Cannot know that someone wants to send him a message (and provide information for alternative ways of communication).
- IM Provider:
  - Cannot discover alternative endpoints to send the message.

6.3.2.3 Alternative Solutions based on existing literature

- Many existing frameworks can complete a message delivery across different domains and technologies not as part of an identity solution but as an integrated service provision.
  - The identity associations are statically defined. (The end user must register - thus reveal - all the identities he wants to associate in the Service Provider).

7 Conclusion

This study examines the need for a global distributed discovery system for identities and identity related information. It is based on the assumption that the information required to provide a service is not available within a single service provider and must be dynamically discovered. Its main purpose is to evaluate existing discovery mechanisms - of identity management systems and other widely deployed systems of various purposes - and examine if these mechanisms are capable of supporting this assumption.

Clause 5 evaluates the main discovery mechanisms that exist today and concludes that in their current form they are unable to provide the functionality needed by a global discovery system for identities and identity related information. Some of the systems have the ability to provide partial solutions if specific requirements are met, but again, not on a global scale. Others, may operate on a large scale, but do not satisfy key requirements of the identity management area e.g. privacy.

Furthermore the scenarios described in clause 6, indicate that existing literature is unable to support a variety of use cases and there is a need for the development of this discovery mechanism. This system may adopt characteristics from existing systems. For example the XRI/XRDS architecture may be used to build identifiers and also as resolution mechanism to get the identifiers associated to an identity and vice-versa. However its exact architecture cannot be presumed. Only after a thorough examination of its requirements and capabilities one may decide whether this system can be based on an existing one, or if it must be built from scratch.

Annex A (informative): Authors & contributors

The following people have contributed to this specification:

Rapporteur:
Konstantinos Lampropoulos, University of Patras, Greece

Other contributors:
Spyros, Denazis, University of Patras

Simo Fhom, Hervais, Fraunhofer SIT

Wolfgang, Steigerwald, Deutsche Telekom AG

Antonio F., Gomez Skarmeta, University of Murcia

History

Document history		
V1.1.1	November 2011	Publication