An Overview of the Id Governance Framework

Version: 1.0

Abstract

The secure and appropriate exchange of identity-related information between users and applications and service providers (both internal and external) is the basis of providing deeper and richer functionality for service-oriented architecture.

Sensitive identity-related data such as addresses, social security numbers, bank account numbers, and employment details are increasingly the target of legal, regulatory, and enterprise policy. These include, but are not limited to, the European Data Protection Initiative, Sarbanes-Oxley, PCI Security standard, and Gramm-Leach-Bliley as examples.

The Id Governance initiative assists entities managing identity data with increased transparency and demonstrable compliance with respect to policies for identity-related data. It would allow corporations to answer questions such as: Under what conditions may user social security numbers be accessed by applications? Which applications had access to customer account numbers on January 27, 2007?

Filename: overview-id-governance-framework-v1.0.pdf

Notice:

This document has been prepared by Sponsors of the Liberty Alliance. Permission is hereby granted to use the document solely for the purpose of implementing the Specification. No rights are granted to prepare derivative works of this Specification. Entities seeking permission to reproduce portions of this document for other uses must contact the Liberty Alliance to determine whether an appropriate license for such use is available.

Implementation of certain elements of this document may require licenses under third party intellectual property rights, including without limitation, patent rights. The Sponsors of and any other contributors to the Specification are not and shall not be held responsible in any manner for identifying or failing to identify any or all such third party intellectual property rights. **This Specification is provided "AS IS," and no participant in the Liberty Alliance makes any warranty of any kind, express or implied, including any implied warranties of merchantability, non-infringement of third party intellectual property rights, and fitness for a particular purpose.** Implementers of this Specification are advised to review the Liberty Alliance Project's website (http://www.projectliberty.org/) for information concerning any Necessary Claims Disclosure Notices that have been received by the Liberty Alliance Management Board.

Contents

1 Introduction to the Id Governance Framework

The secure and appropriate exchange of identity-related information between users and applications and service providers (both internal and external) is the basis of providing deeper and richer functionality for service oriented architectures. Identity-related protocols have been developed that aim to give users more ability to control the consumption and flow of information between service providers on the network, yet the ability to capture and express the constraints on the use of personal data has not been addressed.

The Id Governance Framework is an initiative by the Liberty Alliance Project to provide a policy foundation for multiple identity protocols such as LDAP, SAML, WS-Trust, and Liberty ID-WSF to ensure that Enterprises are meeting governing regulatory requirements as well as terms and conditions expressed with users. While some emerging protocols have worked to improve user control and ceremony experiences, IGF further improves the ability of the user to control the use of their data by adding the ability to capture and express constraints on the user of personal data by web service providers and identity service providers.

Many governments and industry vertical segments are enacting legislation or best practice rules that address issues of identity theft, privacy, and the appropriate use of information by service providers on the Internet. As identity information is exchanged across departmental, organizational, and jurisdictional boundaries, contracts, machine policy, and audit trails between consumers and producers of identity-related data are critical to documenting the use of identity information and its secure exchange.

The intent of the Id Governance framework is to add policy enforcement to systems that produce and consume identity data in order to help all parties manage risks and provide a level of assurance to users that their privacy is being maintained by the parties to whom they entrust their information or who otherwise have access to this information.

The requirements are structured to support a layered approach to Id governance which supports the broadest possible uptake of the result of this initiative. At the foundational layer are privacy properties, consent data, and business agreement references. Privacy properties include information such as how long data is to be persisted or whether it is to used for a single session only. At the next layer are declarative statements by both consumers and custodians of identity data.

For consumers, these statements include details of the identity data sought and the various privacy promises associated with the data. For custodians, these include the conditions that need to be fulfilled when data is released; e.g., whether only specific users or groups or applications have access to the data, the obligations that consumers must fulfill, and whether consent is required.

1.1 Use Cases: Privacy-Enabled Exchange of Identity Data

The following is a simple overview of how policy is tied to the exchange of identity-related data. Items outlined in **green** in the following figures are subjects of standardization or are related to how the standard is applied to a specific protocol through a profile.

Attribute Authority – An attribute authority is typically a business entity that holds information about ***subjects*** that is considered to be authoritative. Depending on context, the attribute authority may refer to the organization that owns, controls, or is otherwise responsible for the information held, or it may refer to a service acting as an Identity Service Provider. An application that either stores ***attributes*** or propagates ***attributes*** and/or ***properties*** may also be considered an attribute authority when the application switches roles from consuming to publishing identity information. For example, an HR system collects information from a user and other sources (e.g., database) and then stores the information for later use by itself or other applications.

Examples of attribute authorities include:

> "**User Managed Attribute Service**" - the user directly controls the ***attributes*** and/or ***properties*** being handed out from the service, irrespective of whether the ***attributes*** and/or ***properties*** are self-asserted, on the desktop vs. at a server, or whether some third-party validation has been performed. Notice that this includes the case where the user, herself, directly provides the identity information.
>
> "**Third-Party Managed Attribute Service**" - the attribute service is managed by some autonomous entity distinct from the user. This entity controls who can access ***attributes*** and/or ***properties***, and the user has some legal/business relationship with the entity but does not directly control to whom the ***attributes*** and/or ***properties*** may be handed out. A common example of such a service is the various enterprise directories and databases that hold information about employees, customers, vendors, and partners.
>
> "**Autonomous Attribute Service**" - as above, but the user has no direct relationship with the service. In this case, however, legislation and corporate practice constrain to whom the user data may be provided. Credit rating and background search services are examples of this type of service.

1.1.1 Simple Attribute Exchange

To exchange information between parties, whether directly or via a user-agent (browser), a client application or consumer issues a request for attributes and properties of an identity to an *attribute authority*. The request may include promises or privileges that are requested regarding the use of data received. For example, the consumer may wish to indicate a request to propagate information to certain parties or to store or cache information for a period of time. The request may also include a reference to a legal document that defines the terms and conditions for sharing information between the parties.

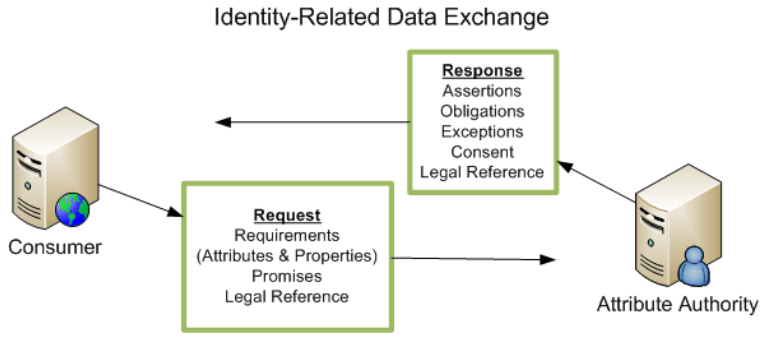

Figure 1

On receiving a request, the attribute authority responds with a set of assertions that may include meta-data such as restrictions, consent, or other legal documentation. The meta-data is used to inform the client application about any transaction-level restrictions that may apply. If the client application requested an attribute or property that was not available, allowed, or filtered due to consent, an exception may also be included documenting why the information or operation was not performed.

1.1.2 Identity-Related Data Exchange with Policy

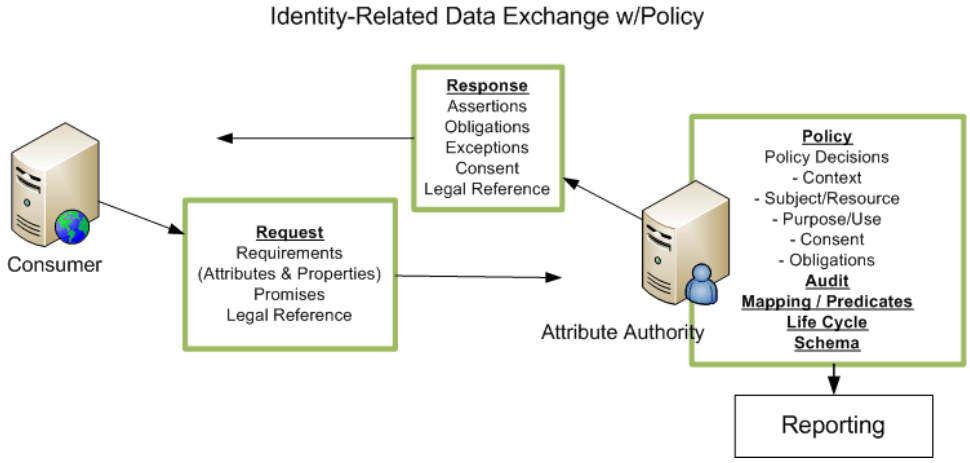

Figure 2

The decision to provide identity attributes and/or properties, filter, or modifications to a response to a client request is determined based upon policy enforced by the attribute authority. The *attribute authority policy* is able to say that, in a specific context, a particular *user* may perform a specific action (read/write) against a *subject* record. In addition to using define policy, *the attribute authority* may also need to define mappings and predicates that map client attribute requests to the schema available within the attribute authority. Finally, the attribute authority is also responsible for the life cycle of identity-related data that is collected for a specific purpose and retained only for the period of time required. Policies, audit, mapping, life cycle, and schema characteristics are functions of a governance-enabled attribute authority.

1.1.3 Declarative Applications

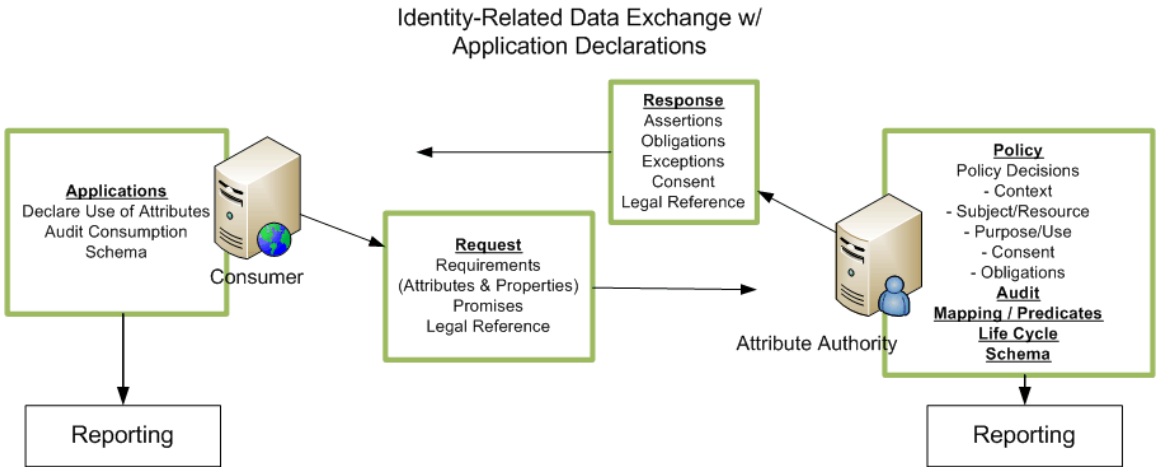

Figure 3

In order to define an *attribute authority policy* that enables a *client application* (consumer) to successfully request and receive identity-related attributes, the attribute authority should be able to understand what the client application requirements for identity-related data are (Figure 3). This information can be passed to the *attribute authority* in advance or as part of the exchange protocol. The combination of *client attributes requirements* and *attribute authority policy* gives auditors and identity managers an excellent understanding of where information is published and where it is consumed. In federated systems, it should also document the legal terms under which information was exchanged and provide attestation evidence that the rules were followed.

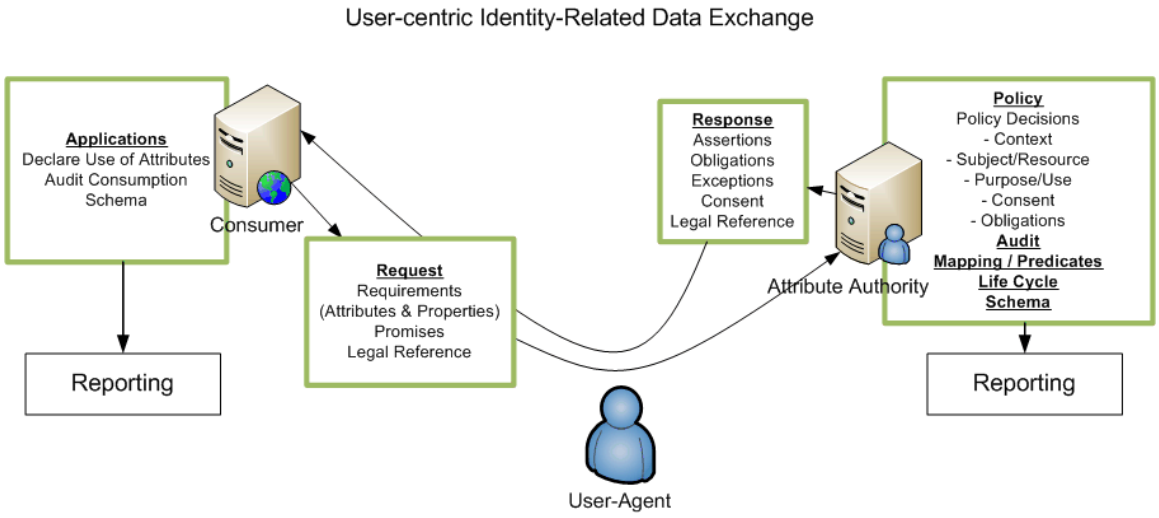

Figure 4

The governance of identity information extends not only to simple client-server relationships but also to three-way relationships (Figure 4) where a user-agent (browser) is responsible for acting as the vehicle for exchange of information and may also be responsible for influencing the flow of information and aggregation. The advantage of user-agent-based interactions is this ability to give users control over the flow of their information, adding much to the enablement of user privacy and control of disclosure of personal information.

1.1.4 Federated Exchange

Before web applications may consume information from a federated *Attribute Authority* service, it is assumed (but not always guaranteed) that the *Web Application* has some business relationship with the *attribute authority* service. Before information can be exchanged, the *Web Application* and federated attribute authority service must agree on protocol (how to access the information), the quality of information (whether it be trusted), the schema (that may be mapped), and when it can be used.

An identity governance framework suggests that a client *Web Application* will provide a declaration of *Client Attribute Requirements* which specifies the identity-related data required and the obligations and usage requirements the application has for using information (e.g., may propagate with specific business partners). On receipt of the *Client Requirements*, the *Identity Services Manager* works to identify appropriate authoritative sources of attribute data that meet the requirements of the client application. Once identified, the *Identity Services Manager* negotiates with the *Attribute Authority Administrator* to determine the appropriate *Attribute Policy*. Once defined, the *Web Application* may then request and/or update attributes from the *Identity Service*.

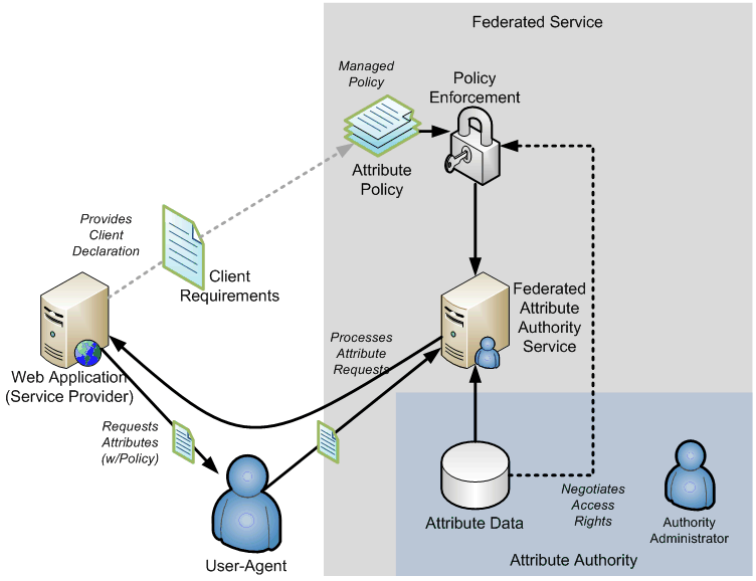

Figure 5

In the case of an interaction where a user-agent or browser is present, the user-agent (Figure 5) is responsible for propagating the specific attribute requests (e.g., Web Application Policy such as WS-SecurityPolicy). Depending on protocol, the user may be aware of this transfer, or it may be part of a referral request given to the user's browser. The *Federated Attribute Authority Service* must then use *Attribute Policy* to determine if the request is appropriate and then process the request.

1.1.5 Flexible Applications

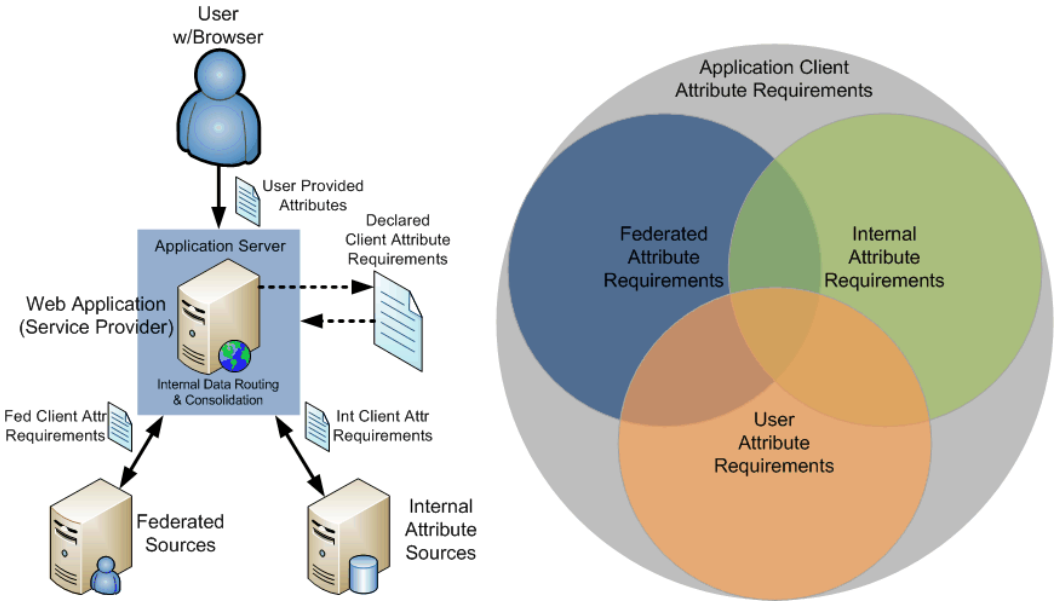

Figure 6

A web application (Figure 6) must be able to handle identity-related information that may be flowing to it from the end-user, from federated sources, and from internal attribute sources (e.g., databases and directories). When deploying an application, consider that while an application may have declared *Client Attribute Requirements*, those requirements will need to be distributed between the *User*, the *Internal*, and *Federated Sources*. This suggests the need for some kind of API, application server provider, or internal service that handles routing and consolidation between the various sources in a configurable and/or manageable way.

1.1.6 Putting It All Together

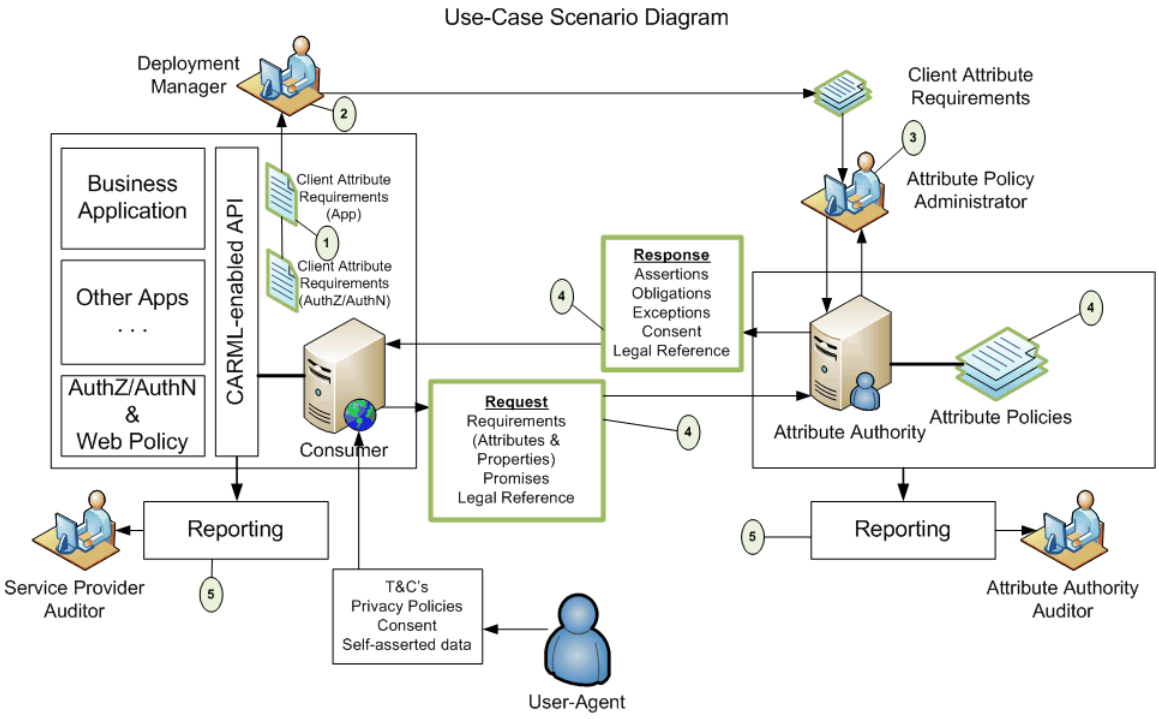

Figure 7

In the diagram, above, the relationships between the deployed application environment, the attribute authority, and the end-user are shown:

1. Developer – the developer declares the attribute requirements of the application.
2. Application Deployment Manager – determines how attributes will flow to/from the application, what information is gathered directly from the user under what Ts and Cs, and what information will come from back-end systems and federated partners.
3. Identity Services Manager/Attribute Authority Manager – Attribute authorities are contacted for permission to use information by providing an appropriate declaration. If the Attribute Policy Admin approves, then the attribute policy for the Attribute Authority can be revised to enable access by the client business application.
4. Client application – Access identity information sources using CARML declaration and AAPML policy enforced providers.
5. Audit Reporting – Auditors on both sides audit the consumption and publication of identity-related information.

1.2 Relationship to Other Projects

1.2.1 Will you develop a new identity/federation protocol from the Id Governance MRD?

Absolutely not! Instead of building a new protocol, our focus is on the management of identity data carried by all the popular identity protocols. Only the Liberty Alliance Project's Technical Expert Group can provide exact recommendations, but we may need to create profiles and recommendations for existing protocols.

1.2.2 What about existing projects like Higgins and Bandit?

We don't believe our proposals duplicate ongoing work within these projects. In fact, we plan to work closely with both projects, especially on components that involve identity data such as the Higgins IdAS (Identity Attribute Service).

1.2.3 Will you help end-users express their privacy requirements?

There are a wide variety of methodologies, dependent on geographies and vertical industry segments, for expressing business terms, privacy requirements, and collecting data and consent from users. This effort does *not* directly address these issues, but does concern itself with how this information should be modeled and expressed within applications and data custodians. The explicit collection of consent information from users is also out-of-scope. The interpretation of different consent assertion types (particularly by policy rules) is within scope.