Glossary of the InCommon Trusted Access Platform Reference Architecture

Introduction

This glossary describes the functional components of the TAP Reference Architecture.

Identity Sources

While not a part of the TAP Architecture, identity sources provide person-related data into the TAP architecture. This data likely includes the following types of information:

- Name information
- Sufficient demographic information to disambiguate a person from others
- Address or location information
- Contact information (email address, phone, etc)
- Data related to institutional affiliations (student data, employee data, etc)

Identity source information flows into the TAP architecture and allows TAP to create the appropriate accounts, groupings and other data structures to enable rule-based provisioning and access control.

Campus System

Student Information Systems, Human Resources Systems and other institutional data systems are frequent repositories of person-related information. These systems provide the TAP architecture with information about the people contained within. In addition to basic demographic information, these systems contribute affiliation information that help to form a complete picture about an individual's relationship with the institution.

Guest / Self Registration

In addition to Campus Information Systems, many institutions provide a mechanism for people to self-register in order to access certain campus services. Visitors may self-register for wireless access, or prospects may self-register for a campus tour. This user-initiated process provides another registration path into the TAP architecture.

Invitation Service

An invitation service provides a mechanism for an existing member of the campus community to invite someone to establish a relationship with the institution. A researcher can invite a collaborator from another institution to establish a relationship and thus be provisioned access to an appropriate set of institutional resources.

Entity Registry Components

The TAP Entity Registry is principally responsible for supplying person-related information to the other TAP architectural components. The Entity Registry has several key responsibilities:

- Aggregate person information from multiple data sources,
- De-duplicate information from multiple sources in order to present a single representation of a person and their various relationships to the institution, and
- Standardize person information for consumption by downstream components in the TAP architecture.

Within the Entity Registry, there are several components that work together to establish a structured repository of identity data. These components include:

Person Registration and Update Service

This component provides the interface between identity sources and the rest of the TAP architecture. The Person Registration and Update Service provides API and messaging interfaces for identity sources to register and update identity data within the Entity Registry.

Person Match / De-duplication Service

This component is used by the Person Registration and Update Service to determine the likelihood that an incoming identity record matches an already existing record. This component will return to the Person Registration and Update Service an evaluation of ‘Exact Match’, ‘Possible Match’, or ‘No Match’. The Person Registration and Update Service can then use this information to register a new person, link new affiliation information to an existing person, or trigger an institutional workflow to resolve an inconclusive match.

Unique Identifier Creation Service

This component handles any necessary assignment or update of identifiers required to register identity information to the Entity Registry. Examples might include internally unique identifiers used by the Entity Registry as well as institutionally significant identifiers such as ID card numbers, badge numbers, login identifiers, etc.

Master Person Store

The Master Person Store is the persistent storage repository for Entity Registry data. The Master Person Store contains demographic, affiliation, contact and account data relating to Entity Registry subjects. This data is accessible to other TAP components through both messaging and API interfaces. Data within the Master Person Store is used to drive grouping, provisioning and attribute release to service providers.

The Master Person Store may be implemented as a standalone data repository that serves the TAP architecture, but may also be implemented as an interface to an existing institutional Person Store (e.g. LDAP, AD or other institutional repository) that serves a similar purpose.

Groups Service

The Groups Service uses affiliation information from the Entity Registry to drive creation and maintenance of automatically maintained (data-driven) groups. These groups can be used to drive mailing lists, authorize users to applications, or trigger provisioning or deprovisioning workflows. In addition to automatically maintained groups, the Groups Service provides a mechanism for managing manually-maintained (ad-hoc) groups.

Automatically Maintained Groups

This component uses affiliations (collections of attributes representing a person's relationship with the institution) in the Entity Registry to form dynamic, data-driven groups based on a person’s affiliation data. For example, a person receiving an ‘employee’ affiliation in the Entity Registry can be dynamically added to an ‘all employees’ group, a group based on their employing department, a group based on their title or job type, and other specific groups based on affiliation data. These groups memberships can be used to populate mailing lists, authorize users to applications or trigger dynamic provisioning workflows.

Similarly, changes to affiliation information can trigger removal from data-driven groups. A person losing an employee affiliation can be removed from employee-related groups, automatically removing authorization and triggering deprovisioning workflows.

Automated grouping can use any attribute data that's associated with a person to form dynamic groups. Class rosters, lists of students by major and employees by title or department are all examples of groups that can be created and maintained dynamically based on attribute data.

Manually Maintained Groups

Manually maintained groups provide a mechanism for the community to express ad-hoc grouping relationships that are not represented in any institutional data repository. Members of a working group or committee may not have a specific affiliation that designates them as such, but a manually maintained group can be created and populated in order to provision resources or grant access based on group membership.

A manually maintained group membership can then be used as an attribute about a person, so can be combined with automated grouping to provide a rich authorization and grouping structure. For example, an application can use an authorization group that is driven primarily by an automatically maintained group containing affiliations that are authorized (e.g. students, faculty, etc), but can also contain a manually maintained group of users that may not meet the dynamic criteria but have been granted access manually. Additionally, grouping math can subtract a manually maintained list of 'suspended' users even if they're included in an automatically maintained group that would grant them access.

Groups Data Store

The Groups Data Store is the persistent store of information provided by the Groups Service. This component provides both API and messaging-based interfaces for other TAP components to receive information about an identity subject’s group memberships.

Provisioning Service

The Provisioning Service component provides the TAP architecture with a mechanism to take action to provision accounts and access either dynamically based on affiliation data changes, or manually based on a request and approval workflow. Likewise, the Provisioning Service works to dynamically deprovision services and remove access when data events occur that impact institutionally defined provisioning rules (e.g. employee termination).

Group-Based Provisioning

The Group-Based Provisioning component performs provisioning and deprovisioning actions based on group data defined within the Groups Service. For example, a person that has been dynamically added to the ‘all employees’ group can be dynamically provisioned any resources that should be provided to all employees. When that person is dynamically removed from the ‘all employees’ group, resources can be dynamically deprovisioned.

Resource Catalog

The Resource Catalog defines the set of services, applications and other resources that a user might request using Request-Based Provisioning. The Resources Catalog presents a menu of possible options based on a user’s group memberships, and routes approval based on workflow defined for each catalog item. For example, an institution may limit access to a particular software application to only those users that express a need, in order to limit licensing cost. Once a user has requested, an approval workflow can be triggered (if required) and access can be granted using Request-Based Provisioning.

Request-Based Provisioning

The Request-Based Provisioning component provides a mechanism for users to request access to a service or resource that is not provisioned dynamically. Request-Based provisioning could invoke an approval workflow that triggers an approval request to a supervisor or resource owner. Request-based provisioning can also provide an audit trail detailing the path by which a person was approved access to a resource. For example, a Business Intelligence application might have a set of data views that are not granted to all users, but can be granted to specific users on request and after an appropriate set of approvals. A user could request a particular Business Intelligence Dashboard from the Resource Catalog which should trigger an approval workflow that routes first to the employee's supervisor and then to a data custodian for approval.

Request-based provisioning can be combined with group-based provisioning to ensure that the requesting user has any additional group memberships required for access. In the above example, the Business Intelligence Dashboard can require not only an approval flow, but also membership in a group of users that that have a sensitive data compliance form on file and another group that have completed sensitive data training within the last year.

Approval Workflow

The Approval Workflow component manages the process by which a requested catalog item is approved or denied. For example, an employee seeking to access a the Business Intelligence Dashboard described above may request it in the Resource Catalog, which triggers an approval workflow that routes first to the employee’s supervisor and then to the data custodian. The Approval Workflow component provides an audit trail that records the path by which a user was granted access to a resource.

Approval workflows may also be triggered outside of the context of a user's request. For example, an institution's Internal Audit department may have an auditing requirement that states that access to a particular financial application must be reviewed on an annual basis. To meet this need, an automated attestation workflow can be triggered annually to record a supervisor's approval that this access is still needed. The attestation process and resulting approvals can be logged and audited, serving as a business record to meet the Internal Audit department's requirement.

Provisioning Connectors

Provisioning Connectors provide the interface between the provisioning system and institutional resources such as applications or infrastructure. For example, a provisioning rule may establish that all employees are to be provisioned accounts in the campus Active Directory Service. The Active Directory Provisioning Connector will interface with Active Directory to dynamically create the account.

Authentication and Federation Services

Authentication and Federation Services provide a means to interface service providers with the TAP architecture to perform single-signon, federated authentication and authorization and to deliver attribute information for consumption by applications.

Single Signon Authentication (SSO AuthN)

This component provides a mechanism for users to authenticate once into a single institutional authentication process, and to have that authentication respected by a number of applications. Single Signon Authentication may also incorporate additional forms of authentication such as multi-factor authentication (MFA) to achieve a stronger authentication session for applications that require it.

SAML and Oauth Identity Providers (IdPs)

These components work with SSO AuthN and Attribute Resolver services to provide standards-based mechanisms for interfacing applications with institutional SSO. Vended applications and cloud providers that support either the SAML or Oauth protocols can be interfaced with institutional SSO and can receive user attributes through the SAML and Oauth IDP components.

Relying Party Data (aka Metadata)

The Relying Party Data component keeps track of the relationship between authentication services and service providers. This component tracks metadata about each service provider including the security components to validate the service provider and the attribute release policies that determine which attributes should be provided to each service provider. Relying Party Data represents an agreement between the institution and a service provider (or federation of service providers) about the strength at which users should be authenticated and the data about users that should be released from the institution to the service provider.

Consent Service

The Consent Service component engages the user in the attribute release process by providing a mechanism to prompt the user for a decision about whether or not it is appropriate to release a particular data element to a requesting application. Where the Relying Party Data above represents an agreement between the institution and a service provider, the Consent Service engages the user in the transaction.

For example, a user authenticating to a cloud service requiring release of 'email address' from the institution can be prompted whether or not they choose to release their email address to the application, or can be made aware that this data is being released. The Consent Service can provide an audit trail to meet record keeping requirements around management of 'opt-in' and 'opt-out' data release policies.

As privacy regulations mature and application integration becomes increasingly distributed, user consent is expected to become an increasingly important component of the identity management ecosystem.

Attribute Resolver

The Attribute Resolver component translates between the internal data structures in the TAP architecture and the attributes that are delivered to service providers. The attribute resolver maps specific internal data constructs to normalized attributes so that service providers do not need to be aware of the inner workings of the TAP architecture to consume attribute information about users that are accessing their services.