Service Provider High-Level Design

Service Provider High-Level Design

This document focuses on the SP and the consumption of identity and user attribute information. This document will refer to the IAM's Common Interface, which is detailed in Archived User Provisioning High-Level Design.

IAM Enumeration

In a typical Shibboleth authentication request, an SP will use some technique to discover which IAM to use. In the user provisioning process, an SP must be capable of enumerating a set of supported IAMs. If using the snapshot or point-in-time access strategy, the SP will query each IAM in turn, retrieving identity information in accordance with the IAM's Attribute Release Policy. If using a subscription strategy, the SP will prepare to receive messages from each supported IAM.

Examples of SP / IAM relationships

1-to-1	An SP only communicates with one IAM. No other IAMs are supported.
1-to-n	An SP communicates with a subset of IAMs listed in the metadata.
Open	An SP communicates with all IAMs listed in the metadata.

SP Attribute Request Process

Snapshot Process

1. SP initiates "Snapshot" call

The SP selects the appropriate IAM and initiates a "Snapshot" call.

SP request for identity information from IAM includes the following:

IAM Endpoint	TBD
SP EntityID	Uniquely identifies the SP to the IAM
Access method	"Snapshot"

2. IAM prepares snapshot

The IAM uses an internal process for interrogating all identity repositories supporting the IAM processes and building a list of all community members associated with the given SP. This prepared snapshot is saved for later transmission by a subsequent out-of-band SP call.

To prevent unwanted information disclosure, the snapshot shall contain only those attributes allowed by the given SP's Attribute Release Policy.

Only one prepared snapshot per SP is allowed: each subsequent snapshot request shall overwrite the previous prepared snapshot.

The IAM shall allow the SP to download the prepared snapshot until the specified deletion deadline has passed.

The IAM response to the SP shall include:

Response code	Indicates whether the request was successful or if an error occurred.
Transaction IDs	See Thoughts about Interchange Formats.
Retrieval endpoint	Indicates to the SP where the prepared snapshot may be fetched
Deletion deadline	Indicates when the prepared snapshot will be deleted by the IAM

3. SP retrieves snapshot

The SP shall retrieve the snapshot prepared by the IAM during step #2 by calling the Retrieval endpoint specified in the IAM's response.

During the idenity transmission from IAM to SP, the IAM shall prevent simultaneous calls from the same SP using a locking mechanism.

If the IAM returned an error response code, the SP shall log relevant information and abort the process.

4. SP consumes snapshot

The SP parses, filters (see "Validation"), merges (see "Identity Merging"), and updates its internal representation of identity information.

Subscription Process

1. SP requests subscription

SP request for subscription to future identity changes shall include the following:

IAM Endpoint	TBD
SP EntityID	Uniquely identifies the SP to the IAM.
Access method	"Subscription"
SP listener endpoint	Where the IAM is expected to push change events

2. IAM subscribes SP

IAM responds to SP subscription request with the following:

Response code	Indicates whether the request was successful or if an error occurred.
Transaction IDs	See Thoughts about Interchange Formats.

3. Ongoing: IAM pushes changes to SP

Details of messaging are left to the messaging implementation.

4. SP unsubscribes from IAM

SP request to unsubscribe from future identity changes shall include the following:

IAM Endpoint	TBD
SP EntityID	Uniquely identifies the SP to the IAM.
Access method	"Subscription"

Changelog Process

1. SP initiates the "Changelog" call

The SP selects the appropriate IAM and initiates a "Changelog" call.

SP request for identity information from IAM includes the following:

IAM Endpoint	TBD
SP EntityID	Uniquely identifies the SP to the IAM.
Access method	"Changelog"
Transaction ID	See Thoughts about Interchange Formats.

2. IAM prepares changelog

The IAM prepares the log of changed identity records available to the given SP. Only those records changed since the last successful Snapshot, Changelog, or Subscription call shall be included in the changelog. The IAM shall provide all new change transactions that occurred since the SP-specified Transaction ID (excluding the specified Transaction ID).

If the SP specifies the Transaction ID of a change that it has already retrieved, the IAM shall return an "Expired Transaction ID" response code. Similarly, if the SP specifies a Transaction ID of a change that the IAM has pruned due to local policy, the SP will receive the same error.

The IAM response to the SP shall include the following:

Response code	Indicates whether the request was successful or if an error occurred.
Retrieval endpoint	Indicates to the SP where the prepared changelog may be fetched

3. SP retrieves changelog

The SP shall retrieve the changelog prepared by the IAM during step #2 by calling the Retrieval endpoint specified in the IAM's response.

Once fetched, the IAM shall not allow the SP to re-fetch the changelog via the same Retrieval endpoint.

Essential data that should be found in every changelog entry provided by the IAM:

Transaction ID	Uniquely identifies the change within the log. The SP may assume that this ID is sequential.
Change type	Describes the type of change. Type must be one of: insert, update, or delete.
Identity record	Contains all attributes available to the SP for the given user, if any.

During the identity transmission from IAM to SP, the IAM shall prevent simultaneous calls from the same SP using a locking mechanism.

If the IAM returned an error response code, the SP shall log relevant information and abort the process.

4. SP consumes changelog

The SP parses, filters (see "Validation"), merges (see "Identity Merging"), and updates its internal representation of identity information.

Validation

The SP shall validate that the data provided by the IAM meets the following criteria:

- Data conforms to the specified formats
- Any institution information listed in the identity information correctly matches the IAM's institution.
- Transaction IDs are consistent, as described in Thoughts about Interchange Formats.

The SP shall reject all data that does not pass validation prior to updating internal identity data.

Identity Merging

When identity information for one user appears in multiple IAM data sets, SPs may opt to unify the identity information using a identity merge process. Currently, there is no universal way to determine if two identity records strongly match, indicating the two records should be merged. This being said, the UC system does generate UCnetIDs for employees and some other classes of users. UCnetIDs, when present, will uniquely identify a user across the entire UC system. UCnetIDs will allow an employee-only SP to merge identity records, however this ID cannot be relied on for the more general SP (student, affiliate, and some non-salaried appointments).

The solution for multi-campus user provisioning does not specify a common design for merging identity records. The design of merging identity records is left to individual SPs.

Response codes

The SP shall be capable of interpreting the following response codes sent by the IAM:

Success	Call succeeded and IAM advanced the internal journal position for the SP
Error: Internal server error	IAM is misconfigured or an exceptional error occurred. The IAM is unable to complete the request
Error: Method not allowed	IAM does not support the requested access method.

Error: Resource locked	IAM has locked the resource because another SP call is in progress.
Error: Not found	IAM could not find the specified resource. This may occur if the SP does not retrieve a prepared snapshot before the deletion deadline.
Error: Expired Transaction ID	The specified Transaction ID has expired because either: the SP already retrieved the given change, or the IAM pruned the change per local policy.

References

Document	Description
Archived User Provisioning High-Level Design	Provides overall design of the User Provisioning project
IAM High-Level Design	Describes the Common Interface that the SP will use to access identity information