

Requirements on an Entity Registry and Related Components

General Principles

1. Support/advertise a strong conceptual difference between email and user ID
2. Support widely used standard authn/authz protocols for federations (OAuth2 + SAML2)
3. Support multi-site replication and synchronization
4. Support unicode (and make clear what character set is supported)
5. Avoid/disallow re-use of persistent identifiers (or define "persistent" better!)
6. Allow non-person Entities
 - a. Client /Agents
 - b. Service Accounts
 - c. Department/Organization
 - d. Internet of Things (devices, IOT)
7. Suggest various ways people can model their data; Do we want a registry that could host different models?
 - a. Relational vs LDAP
 - b. Use as delivered vs. customize
 - c. "Built in" schemas vs. common configurations (that can be customized)

Companion Doc for Data minimal requirement for items marked registry [Minimal Entity Registry Definition/Logical Design](#)

Tabulated Requirements

#	Component	Requirement	R1 - minimum registry feature	Notes
1	Identity Registry	1. Paths in and out <ol style="list-style-type: none">1. RESTful API2. Asynchronous messaging interface (PUB/SUB, etc)3. Administrative interface console	yes	Flat file removed by consensus of the workgroup .("bulk up/download"; SFTP or similar to make the file available to ingest via message or API logic)
3	Identity Registry	For new records, assign a permanent unique identifier to map between various source system identifiers. This entity identifier must be made available to the SOR as a response to the entry from the SOR.	yes	This is perhaps the single most important thing the registry must do.
4	Identity Registry	When SOR notifies Registry of an entity/person, return relevant info to that SOR . Minimally, the unique identifier assigned by the Entity Registry and the institution may extend what is returned.	yes	SOR return INFO can be extensible
5	Identity Registry	Change (add/modify/delete) notifications/events to Provision when an "attribute" changes on a Person record. Minimally registry records entity, attribute identifier, verb, old value, new value, timestamp of change	yes	This feeds the requirement 23 for Provisioning Component.
6	Identity Registry	An entity/person can have multiple simultaneous affiliations with an organization. We will use the term affiliation	yes	Need consistency on "relationship" "role" "affiliation"
7	Identity Registry	Each relationship has a "type" (affiliation) and can have its own set of data describing the individual and this relationship (start/end dates (possibly in the future), dept/center, title, who/what added the entry, affiliation type owner)	yes	Anything more than this should be handled in the Groups component. It will be deemed not a registry function.
8	Identity Registry	An entity can have multiple affiliation relationships with the same "type" value (eg faculty member who is associated with multiple academic departments)	yes	This would also be handled by the groups database. How affiliation is handled by the registry vs the grouping component is to some degree institutionally selectable. The workgroup believes that the affiliation relationships are significant enough to belong in the registry. The may be simultaneously built out in the grouping component or fed as data events from the registry to the grouping component. Choice to be made.
9	Identity Registry	Support start and end (sunrise/sunset) dates for attributes. Many attributes should support these dates. Phone, email, name(s), affiliations, etc. The date serve as triggers and allow for a live history to be built for an entity.	yes	The live history allows the data to provide all names a person has been known as or ... , data is rarely (maybe never) deleted they are simply not current info.
10	Identity Registry	The Registry does not need to hold all IAM data within it. Rather data is to be considered to be contained in one of three conceptual data containers: Entity/Person registry, Groups and Privileges, Party (person/organization) ODS/MDM data stores.	yes	These can take the forms of relational data, LDAP, etc. The need for a FAT everything registry is a waning technique. Organizations have built Perdon ODS data, person data hubs that serve all applications . These can include the IAM Entity Registry.
11	Identity Registry	Support extensible local and/or auxiliary information about entities	yes	

12	Identity Registry	Associate a "level of confidence" with various attributes (eg self-asserted, verified via gov't documents, etc)	no	This is different than the entities LOA. It is a measure on how data was collected and vetted on an attribute / sourcing level. For example: a person may self assert their name is John Doe, but at personal gov't documents indicate John Doe should really be Jonathon Doe. The level of confidence for the info is better after the vetting at HR than when self asserted.
14	Identity Registry	Email notification to user indicating change/pending change to key registry profile information.	yes	Similar features should exist for Credential management and provisioning/de-provisioning
15	Identity Registry	Support Batch purging of entries (e.g., applicants) [May require a different concept than "purge". "Permanent disable"?, should use a soft delete mechanism]. Generally this will be the ending of an affiliation like applicant it might even add an affiliation former applicant. A repetitive calling of the API/message (see #1) is the process or doing this. Institution would set up a process to take in a list and call the service. This assures that edit, triggers and all logic involved in setting individuals and communicating changes is followed.	No	Should be a soft delete concept - usage of one of the methods in item 1. Flat file batch will not be supported. The use of this is standard affiliation management
51	Identity Registry	A Person may have multiple personas that an organization may require them to "act in the role of", An easy way of switching personas should be constructed as a part of the final solution.	yes	Not R1, but data model should support extension later
35	Identity Registry	Associate multiple authentication methods with an entity in the Registry	yes	36 is same requirement
36	Identity Registry	Methods can be internal (ie managed by the organization) or external (ie rely on a different organization to perform the authentication and assert its result; eg social)	yes	Not "must support anything", but must support an external authN method
37	Identity Registry	Each Authn method should have an associated LOA - Assurance measure/value	no	not sure if minimum - but is a good requirement
61	Identity Registry	Support various management models for GUEST types (eg self-registration, require a Sponsor with specific Roles, etc)	yes	
62	Identity Registry	Support specific terms for GUEST type (eg must be renewed every N months)	yes	Must have a parameter specified duration begin end date above
57	Identity Registry	Ability to spin up "collaboration services" for campus researchers and other groups, where a campus member is designated as the collaboration administrator and can invite other participants, and can enable applications (such as file storage and email lists) for the collaboration.	maybe	Restful API / VO membership / May be related to Groups as well. ??? Is this minimum. ?
56	Identity Registry	Ability to store comments associated with any edits (including running comments)	yes	Manual over ride. stewardship/admin ... - move to Audit.
2	Identity Registry Identity Matching	As part of Registration from an SOR, invocation of Identity Matching Engine . Registry attempts to match with an existing record <ul style="list-style-type: none"> match - can positively identify an existing Registry entity/person - becomes an update no match - Can not identify a preexisting registry entity/person - becomes an add indeterminate (maybe) identifies possible collisions but match logic is not scored high enough to determine a specific match. This requires a human interaction and mapping of information. A new record will be added with a "suspect duplicate status" . A data steward (human) type responsibility to resolve these using the merging/splitting function. The institution will need to decide if provisioning is allowed to these cases prior to resolution. It is the working groups recommendation that these suspect duplicates are not candidates for service provisioning until they are resolved and no longer a suspect duplicate record. 	kh, wc, hs, eg	deleted former 17 and 18 row duplicate to this. Identity matching is required function. Might be a service call to a Identity match service , Solution to be determined,
47	Identity Registry Identity Matching	Support for finding potential duplicates ("suspect duplicates") entity/persons and adding /merging/splitting records to resolve and the resolution of these registry entries.		Moved to pair with the requirement - matching function item # 2
19	Identity Registry Identity Matching	Identity merging needs to be well managed and low impact. The assignment of provisional ids is a method for special use cases of merging		
20	Identity Registry Identity Matching	Attempts to match with an existing record in the Registry use heuristic algorithms		
21	Identity Registry Identity Matching	May rely on "attribute assurance level" when matching input values against Registry entries		
42	Identity Registry Audit	Events performed by any of these components must be recorded such that an Audit system can perform queries in various ways and see the results of those queries	kh, wc, hs, eg	
43	Identity Registry Audit	Maintain a secure permanent audit record / history of ALL changes related to an entity record.		
32	Identity Registry Authentication	Users must be able to authenticate to the Admin Console	eg	IF Admin console is not included (see #1)

33	Identity Registry Authentication	The Registry should support authentication via CAS and Shibboleth(SAML2) or other methods supported by TIER. The Identifier provided by the authentication mechanism should be used to search the Registry to find the matching record.	kh, wc, hs, eg	OAuth2, etc are valid candidates as support is considered moving forward.
34	Identity Registry Authentication	External services must be able to authenticate to the RESTful /messaging endpoints exposed by the Registry	kh, wc, hs, eg	This security function is currently being discussed(march 2017)
53	Identity Registry Authentication	Beyond WEB Only Authentication (e.g. ECP and CLI protocols) for authentication must be enabled as for Research/Collaborative computing		
40	Credential Mgmt /Storage	Provide a mechanism for (possibly) storing and propagating various secrets supporting authentication (eg passwords, personal certificates, two-factor secrets, lower quality passwords (eg synched gmail), KBA questions/answers		see 41 and 42 are these the same.
41	Credential Mgmt /Storage	Password Reset capabilities must be standardized upon and deployed in the out of the box solutions, with sufficient flexibility to meet institutional business practices. (Probably need to talk through the non-password self-service interface -- allow emailed one-time links, one-time printed tokens, 2FA and other "private token" mechanisms)		Do we need to call out the ability to manage account and passwords securely? wc
38	Credential Mgmt /Storage	Various events can raise and lower the associated LOA (eg password reset over the phone could lower a password-based LOA)		
39	Credential Mgmt /Storage	If an internal method has Identity Vetting Requirements support them in some fashion	yes	vetting/ proofing etc..
49	Identity Registry UI Console	Support for out-of-band password reset mechanism ,(SMS/email, etc)		similar to 41
13	Credential Management	Support for provisioning codes (one-time use link/code/token) for account claims		Reclassified to Credential management
45	Identity Registry UI Console	Search for users (including users who are no longer active)	eg	needed by other functions (eg password reset)
46	Identity Registry UI Console	Support for "renaming" users, and changing any of their attributes (including their various identifiers)	eg	eg: "any" is overstated for r1
48	Identity Registry UI Console	Support for creating entities in the Registry	eg	eg: Unless this is solely PoC, need some ability to create people not from SoR wc: do we need this in the POC, need to review 1.4 and 48 (are these the same)
50	Identity Registry UI Console	Support for authentication to Admin console using various authentication methods		
54	Identity Registry UI Console	Allow users to see (portions of) their records, and maintain the self-asserted attributes in their record		eg: seems an easy addition; tempted to put as R1 As : POC
16	Groups	There is a need to identify a "primary" Affiliation? (Primary affiliation calculation is a requirement to assist in handling the EduPerson Primary affiliation., calc required when individual has multiple distinct types of affiliation student and employee for example institution must decide how they handle this.)		Seems to be best handled in the Grouping tool. This could be fed to registry based on grouping result.
52	Groups	Support for authorization framework (different People/Roles authorized to see/change different attributes; LOA of authentication method affects permissions)	kh, wc, hs	Handle with Groups eg: This seems broader than "different permissions". I think this was referring to literally a general purpose privilege management service.
60	Groups	Provide support for the creation and maintenance of a type/affiliation of "GUEST" affiliation and many others on Registry records		seems like a group feature related to affiliation (s) that are loosely attached to the institution
23	Provisioning	When an "attribute" changes on an entity data was placed for provisioning to consume based on the event. Entity record an event to be provisioned with minimal field including: entity, attribute identifier, verb, old value, new value, timestamp of change.	wc, hs, eg	Likely to use a logging concept initially, think through this in more detail. An API call or a messaging channel should be the consumer. traditional connectors are valid in this use case as well. Grouping facility (grouper or something else) clearly must be a consumer of the event. The knowledge sharing of entity info seems well suited for asynchronous messaging to pub-sub style consumers. However, technology can vary by institution. (Provisioning and Connectors 23-31 wc 4/22)
24	Provisioning	Rules that specify Provisioning Operations can trigger these events (invoking specific outbound Connectors associated with specific target systems)	eg	
25	Provisioning	These events can be consumed by internal processes which then change other Attributes (eg passing an End Date causes Status to change Active to PENDING)		
26	Provisioning	These events can also be consumed by "Connectors", which then effect changes in external systems.	eg	
27	Provisioning	Semantics of a change are determined by each Connector (eg Idap vs google vs LMS, etc)	kh, eg, wc	

28	Provisioning	Receive from the Provisioning System an event describing a change in the Person record; they map that change to the appropriate sequence of events to transmit to their associated external system. (eg provisioning accounts, synchronizing passwords, changing permissions, etc)	eg	
29	Provisioning	Events contain: attribute identifier, verb, old value, new value)		
30	Provisioning	A mechanism to augment the catalog of Core Connectors must be provided to the community for inter-institutional sharing and implementation.		
31	Provisioning	A set of pre-built connectors should be supplied "out of the box" (eg ldap, AD, kerberos, Grouper, SCIM, some popular cloud based services (eg Canvas), etc), Initial for LDAP, Kerberos only	wc, hs, eg	IAM side of connector speaks messages and /or restful APIs
44	Provisioning	It MUST be possible to see the relationships between events in the different components (eg a Registry change triggers a Provisioning change triggers a Connector action)		
55	Provisioning	Support for workflows that involve administrative sign-off from specific users (eg approval for certain types of edits)		
58	Consent	The solution may enable user to be in control of their personal data stores such that when relying parties are requesting access to those data, users should have fine-grained controls over what pieces of personal data are shared with such parties.		
59	Partitioning	Partitioning is mentioned in several use cases, and is difficult to define. There are a number of underlying conditions that seem to lead to "partitioning"; these should probably be teased apart and treated individually, as none of them yet seems compelling on its own. (Most seem like a data presentation question - perhaps a locally defined attribute for an account which is then important when Connectors are invoked).		??? DO not understand this. Can anyone clarify.. Is this from investor sessions.. or ???
63	Community Documentation and Interaction	Solution extensions must be available in the form of a Marketplace or some other suitable means of presenting a catalog of available functionality, contributed by the community, for utilization by others.		
64	Community Documentation and Interaction	Solution must enable the sharing of a common documentation repository as well as a place for school practitioners and service providers to go to find useful instructions, standards, practices and guidelines for building end-to-end services based on TIER components		
65	Standards and Enforcement	The program must assert and enforce Policy Standards		
66	Policy and Performance Monitoring	Log files should be available to monitoring tools. Should be able to discern what data was seen and changed during a session, Which features were used..	kh, wc, hs	Use ELK stack eg. Agree in principle, abstain on anything but "log files exist and monitoring tools can be made to read them"