Modeling the Federal User Identity, Credential, & Access Management (ICAM) Decision Space to Facilitate Secure Information Sharing

Thomas C. Smith, Maria E. Vachino
Anil John, Chi Y. Wu, Christopher D. Obremski
Johns Hopkins University
Applied Physics Laboratory
Laurel, MD
Tom.Smith@jhuapl.edu, Maria.Vachino@jhuapl.edu

Karyn Higa-Smith
DHS
Science & Technology Directorate
Washington, D.C.
Karyn.Higa-Smith@dhs.gov

***Abstract*—Providing the right information to the right person at the right time is critical, especially for emergency response and defense operations. Accomplishing this across sovereign organizations while keeping resources secure is a formidable task. What is needed is an access control solution that can break down information silos by securely enabling information sharing with non-provisioned users in a dynamic environment. Multiple government agencies, including the Department of Homeland Security Science & Technology Directorate (DHS S&T) are currently developing Attribute Based Access Control (ABAC) solutions to do just that. ABAC supports cross-organizational information sharing by facilitating policy-based resource access control. The critical components of an ABAC solution are the governing organizational policies, attribute syntax and semantics, and authoritative sources. The policies define the business objectives and the authoritative sources provide critical attribute attestation, but syntactic and semantic agreement between the information exchange endpoints is the linchpin of attribute sharing. The OASIS SAML standard provides federation partners with a viable attribute sharing syntax, but establishing semantic agreement is an impediment to ABAC efforts. This critical issue can be successfully addressed with conceptual modeling.**

Keywords-ABAC; DHS; data modeling; access control; ICAM

I. INTRODUCTION

DHS S&T is currently exploring the viability of using ABAC for federated access control within its IdM (Identity Management) Testbed via proof-of-concepts and pilot activities. Within DHS, authoritative sources of Identity, Credential, & Access Management (ICAM) attributes are distributed across various legacy repositories. Attribute information within these independently developed data stores has been captured in disparate formats using inconsistent naming conventions. There is therefore a compelling need to consolidate and correlate this information in the form of an Attribute Authority that supports congruent yet customizable attribute views both for internal policy decision points as well as for external partner authentication and authorization needs.

The described conceptual model identifies and characterizes the attributes of a potential DHS enterprise Attribute Authority, representing a normalized view of the identity and authorization attributes required for fine-grained access control decisions. The context of each of these attributes and the relationships between them is made explicit in order to definitively capture the semantics of each attribute. The result is a canonical concept model that facilitates the semantic mapping of identity, authorization, and security attributes between agencies, commands, and departments while clarifying and helping establish access control policies. It is based on an analysis of DHS and federal policy, and has been reconciled with the published partner attribute requirements of the Department of Defense Intelligence Community (DoD/IC), Department of Justice (DoJ), and the Homeland Security Presidential Directive 12 (HSPD-12) Backend Attribute Exchange (BAE).

II. ATTRIBUTE-BASED ACCESS CONTROL

"Access control … deals with users' access to resources and, specifically, how such access is regulated according to applicable policies." [1] Access control is composed of two discrete functions – authentication and authorization. Authentication is the confirmation of the user's claimed identity while authorization is the policy-based decision to allow an authenticated user to perform the requested action on the specified resource. [1] There are currently three predominate methods for authorizing users – Identity Based Access Control (IBAC), Role Based Access Control (RBAC), and ABAC. IBAC provides access control through a direct association of resource privileges with the unique identifiers of provisioned requestors. This model works well for some resources, particularly those where a relatively small and infrequently updated Access Control List (ACL) provides sufficient security for a resource. RBAC maps a user to a role, which is in turn mapped to permissions. This is more scalable than IBAC but presents challenges when job titles don't map exactly into assignment requirements, or when assignment requirements change. RBAC also makes it difficult to alter access control quickly in response to changing access needs such as may happen during a disaster, and it does not take into account contextual information that may be highly relevant to

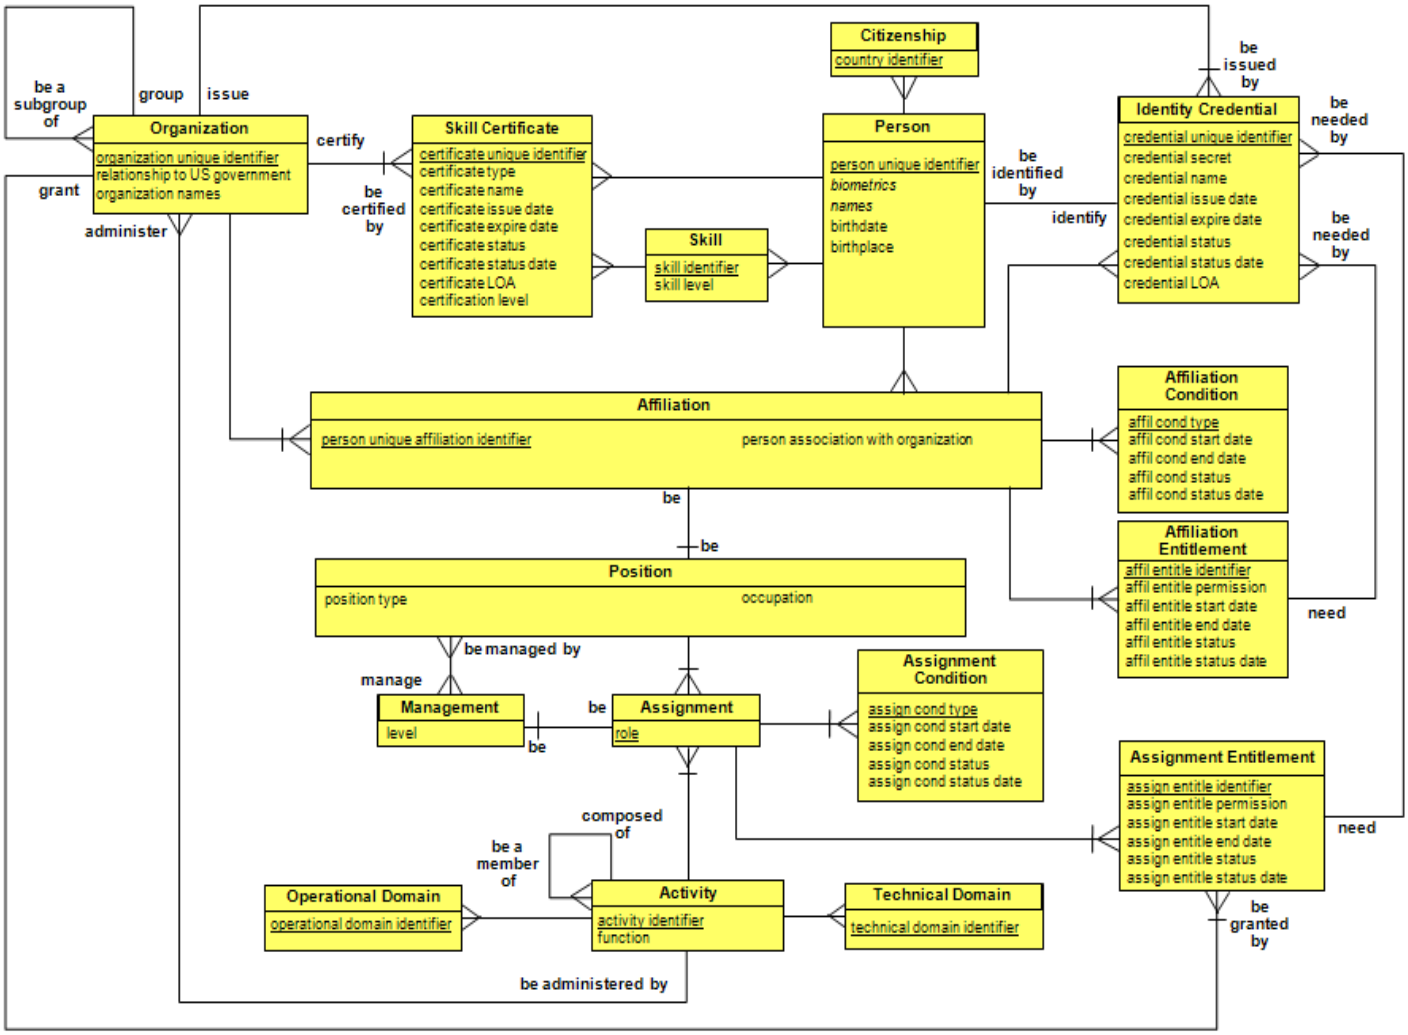

Figure 1. Federated ICAM User Attribute Concept Model

security such as the location of the user and the time of day. ABAC addresses all of these concerns and provides the additional benefits of uncoupling policy from implementation and facilitating cross-organization authentication. In ABAC, attributes are the intrinsic characteristics of an object that capture the policy requirements necessary for authorization decisions. This enables fine-grained access control. From a policy perspective the ABAC decision space identifies the requestor (subject), the protected resource, the requested action, and the environment as conceptually distinct elements. Basically, a subject requests an action on a resource in a given environment; attributes for all of these actors and elements can be used to make policy-based access control decisions, but conceptual alignment of subject attributes is critical. We will be addressing the attributes of human subjects (users) in our model.

III. APPROACH

There are fundamentally two approaches to modeling the ICAM attribute information space. The first is a bottom up approach, often driven by practical concerns, in which legacy directory services and applications are mined to discover existing attributes that could potentially be consolidated into a list that is applicable across the enterprise. However, this can easily results in a policy agnostic attribute set that does not create an information space capable of fully supporting enterprise access control policies. The second approach starts at the organizational policy level and identifies a concrete set of attributes to fully support the enterprise policies. DHS S&T has taken this approach, creating a canonical model for federal user attributes that supports policy and facilitates semantic alignment. The work of Waterman & Hammer [2] highlights DHS policy-based decision concepts and provides the foundation for the conceptual ICAM attribute model that is presented in this paper. The complete canonical model is shown in Figure 1.

Data modeling can be partitioned by the granularity of its implementation knowledge into three progressively dependent layers: conceptual, logical, and physical. Conceptual modeling is used to provide a technology neutral representation of the

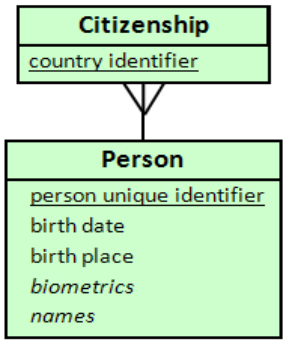

Figure 2. The Human User

data components and their relationships. A logical data model is then employed to focus the concept model by identifying assertions specific to a general technology (e.g., LDAP or RDBMS). Finally, a physical model is used to tailor a logical model with product specific implementation details. The primary advantage of this layering is that it stabilizes and maximizes the reuse of implementation knowledge across the increasing details that are captured in the specified options. This improves the consistency and adaptability across implementation choices. As the least granular approach, conceptual models define a natural balance of interrelated concepts that collectively express agreement in the form of a shared understanding about how to think about the modeled domain. They express the mutual understanding that spans organizational boundaries and establish the semantic concurrence that enables information sharing [6].

IV. FEDERATED ICAM USER ATTRIBUTE CONCEPT MODEL

A. *Person*

The ***Person*** is the individual that is requesting resource access. The Person concept includes biometrics, names, birth date, and birthplace as identifying attributes (see Figure 2). Identity uniqueness can be approached with representations of biometrics such as fingerprints and retinal scans and through combining a legal name with birth date and birthplace [2]. However, we explicitly created a *person unique identifier* attribute to guarantee uniqueness in our model. Its value can be anything, including these mentioned items, which satisfy object uniqueness. Uniqueness is highlighted in this Logical Data Structure (LDS) representation by having the attribute name underlined. Lines adjacent to a concept and perpendicular to a relationship line (as ***Skill Certificate*** in Figure 3) indicate that uniqueness for that concept requires the inclusion of the unique identifier(s) associated with the related concept.

B. *Citizenship*

The Person concept has a one-to-many relationship with the ***Citizenship*** concept, which indicates that a ***Person*** can have multiple citizenships (see Figure 2). The *country identifier* can be implemented using standards such as ISO 3166[3] and FIPS 10-4[4]. The citizenship concept can be considered as a specific form of a more generalized *user organizational affiliation* model (see Figure 4) where citizenship represents an affiliation with a government. Depending on the application such generalizations can simplify and clarify concepts, or they can

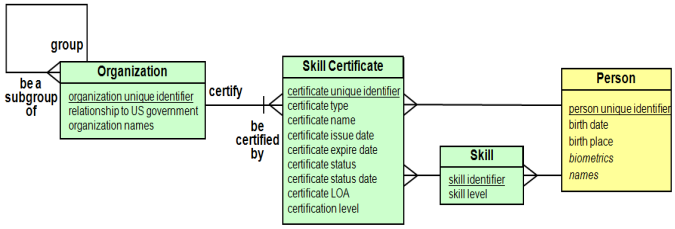

Figure 3. A Peron's Skills. New concepts are highlighted in green.

obscure meaning [5]. Conceptual data modeling should aim to expose the vernacular that is needed to clearly articulate the business rules that define policies [6]. In the federal government citizenship is an important proxy for political allegiance so we avoided the generalization in this case and exposed it as an atomic concept.

C. *Skill*

A Person can have multiple Skills, which are capabilities and proficiencies that are self-asserted unless they are associated with one or more Skill Certificates (see Figure 3).

From a policy perspective, mapping a self-asserted skill to one or more skill certificates is important because third party certifications provide greater trust than self-asserted claims. However, when a skill certificate is available an additional explicit enumeration of the implied skills is generally not required for access control decisions. We therefore created a one-to-many relationship between ***Skill*** and ***Skill Certificate*** rather than a many-to-many relationship.

D. *Skill Certificate*

The ***Skill Certificate*** concept objects are identified by a *certificate unique identifier*; however, uniqueness cannot be guaranteed without combining the issuing organization's unique identifier with the certificate's unique identifier. This is indicated by the perpendicular bar on the relationship line between ***Skill Certificate*** and ***Organization***. (If a value for *certificate unique identifier* is not explicitly available for a particular instantiation of the canonical model, a unique identifier can be derived through a combination of ***Skill Certificate*** attributes, and if necessary, additional attributes from ***Person***.) All attributes identified in the model but not underlined are useful for access control, but do not guarantee uniqueness for a particular object. The *certificate LOA* captures the Level of Assurance (LOA) ranking that quantifies trust in the certification. An LOA ranking can be arrived at by taking into account such things as the certifying organization's reputation, its certification process, and its methods for authenticating the person receiving the certificate. The *certificate name* and *certificate type* identify the skill(s) being certified, and sometimes the proficiency. The *certification level*, if present, could explicitly establish proficiency. The *certificate issue date*, *certificate expire date, certificate status* (e.g. pending, provisional, suspended, revoked), and *certificate status date* provide life-cycle details for the **Skill Certificate**.

E. *Organization*

Each ***Organization*** object is uniquely identified by its *organization unique identifier* (e.g., NIST SP 800-87[7], Data Universal Numbering System (DUNS), etc.). The *relationship*

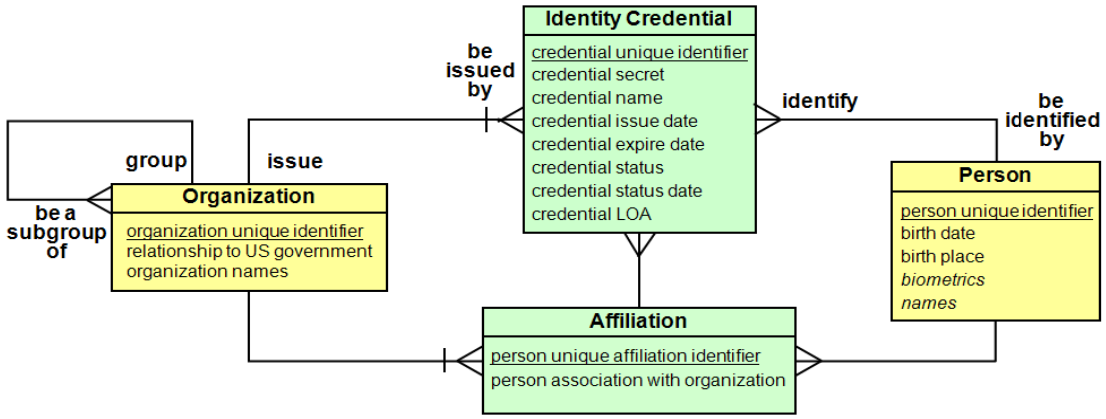

Figure 4: Person Identity Credentials and Affiliations with Organizations.

to US government attribute provides an allegiance category for policy mapping and inference. The General Services Administration (GSA) has currently established four values for this attribute when used for Personal Identity Verification (PIV) cards: Federal Government Agency, State Government Agency, Commercial Enterprise, and Foreign Government [8]. Other categories, such as Local and Tribal Government Agencies can used as required. An Organization may be composed of sub-organizations; this is represented in the LDS notation with the circular "*group"* (one-to-many) and "*be a subgroup of"* (many-to-one) labeled relationship.

F. *Affiliation*

Persons may have additional relationships with ***Organizations*** which are established through the ***Affiliation*** concept (see Figure 4). An ***Affiliation*** object is uniquely identified by a combination of its *person unique affiliation identifier* and the affiliated organization's *organization unique identifier*. The *person association with organization* attribute indicates the type of affiliation the person has with the organization (e.g. employee, contractor, board member, intern, beneficiary, or retiree). Traditionally, organizations manage different affiliate types in dispirit information spaces, but this is problematic from an access control and life-cycle management perspective [8]. We therefore maintained the ***Affiliation*** generalization and did not create separate concepts for each affiliation type. ***Affiliations*** may have conditions that either constrain or amplify their policy-based assessment. This concept is captured in the model with the ***Affiliation Condition*** as illustrated in Figure 5. Affiliation condition type (*affil cond type*) defines the condition. These conditions may include a status (*affil cond status*) (e.g. active, inactive, permanent, pending, or provisional), a status date, and lifespan information. ***Assignments*** can have similar ***Conditions***.

G. *Identity Credential*

A ***Person*** can be identified by any number of ***Identity Credentials***. Each ***Identity Credential*** is issued by an ***Organization*** and may be associated with a specific ***Affiliation***. In that case, the issuing ***Organization*** may or may not be the same as the person's affiliated ***Organization***. The ***Identity Credential*** establishes a person's identity for the purpose of user authentication. It may be a card ("something you have"), may include a *credential secret* ("something you know"), and may contain a ***Person's*** *biometrics* ("something you are").

H. *Position*

As noted above, a ***Person***'s ***Affiliation*** with an organization can *be* a ***Position*** with that ***Organization*** (see Figure 5). Figure 5 illustrates this concept specialization: a ***Position*** *is* an ***Affiliation,*** which is represented in the LDS notation with a "*be"* labeled relationship that has an identity bar adjacent to the specializing concept. The *position type* is a Human Resource designation such as military rank or General Schedule (GS) pay grade. The Office of Personnel Management (OPM) Occupational Categories or their equivalent can provide semantic agreement for *occupation*, especially when used in conjunction with *position type.*

I. *Activity*

An ***Activity*** can provide a business *function* that may encompasses multiple ***Technical*** and ***Operational Domains*** (see Figure 5). Budget and charge numbers are typical examples of top level *activity identifiers* within organizations, but cross-federation activities such as Communities of Interest (COIs) and workshops would most likely be identified by a reasonably unique name. Activities can also be composed of sub-Activities, as indicated by the reflexive relationship seen in Figure 5. The ***Organization*** which administers an ***Activity*** may or may not be the same as the ***Organization*** with which the ***Position*** is affiliated.

Technical Domain is a categorization of the technical or business area (e.g. Nuclear Regulation or Cyber Security), while the ***Operational Domain*** provides a classification of the operational activity (e.g. Al Qaeda or Organized Crime).

J. *Assignment*

An ***Assignment*** is a person's *role* (through their ***Position***) in an ***Activity*** (see Figure 5). It is possible for a Position to be assigned multiple roles for the same Activity, and to have multiple roles in multiple ***Activities***. Therefore, the *role*, the *activity identifier*, and the *person unique affiliation identifier* in conjunction with the *organization unique identifier*, are all required to guarantee uniqueness for an ***Assignment*** object.

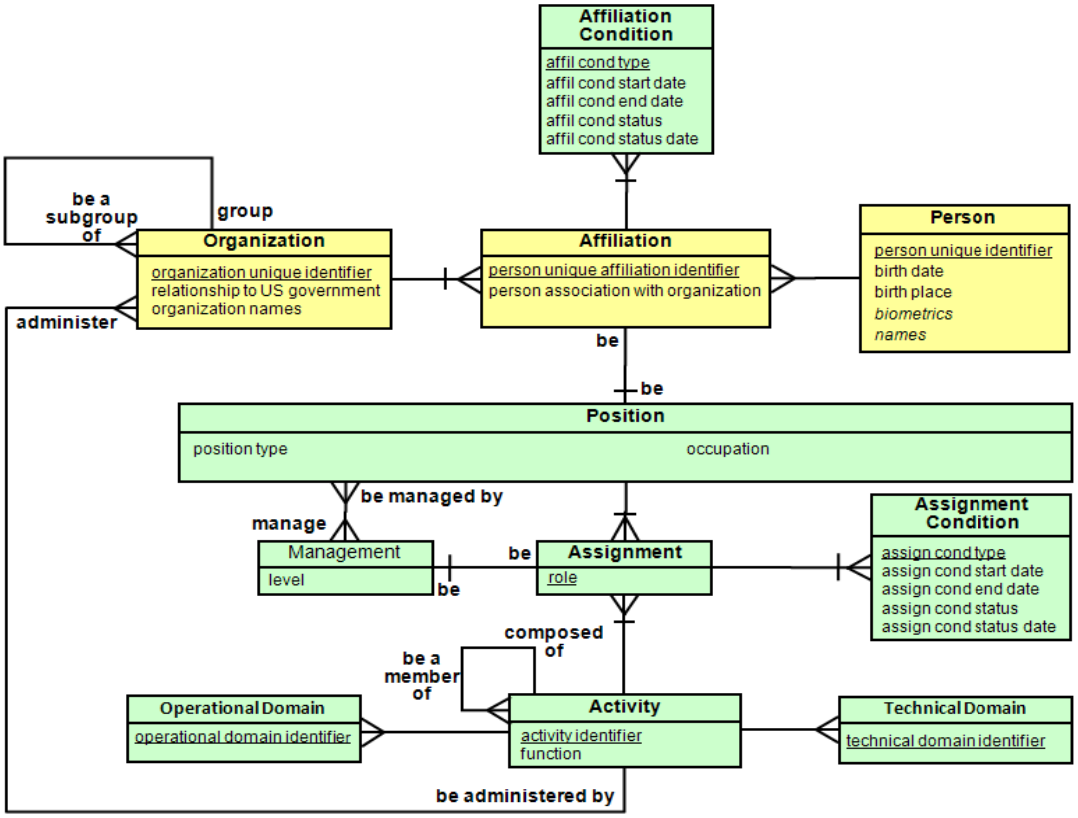

Figure 5.The Position specialization of an Affiliation and supporting Activity Assignments**.**

Assignments and ***Affiliations*** can both have associated ***Entitlements***. These entitlement concepts differ primarily by permanence (***Assignments*** change more frequently than ***Affiliations***) and by the granting ***Organization*** (***Affiliation Entitlements*** are only granted through the ***Position's*** affiliated ***Organization***). When policies are applied to entitlements they generate privileges, which are resource specific actions. Examples of entitlements are security clearances and legal jurisdiction. Privileges generated from these entitlements may include access to classified material and the ability to issue warrants. There is sometimes a fine distinction between a privilege granted by an access control policy and an ***Entitlement*** as we have defined it. Embedding privileges in an ABAC attribute model eliminates dynamic policy based access control, so the decision to include a specific *permission* as an ***Entitlement*** must be made carefully. Both ***Position Entitlements*** and ***Assignment Entitlements*** may need specific ***Identity Credentials*** to be utilized. For example, a Common Access Card (CAC) ***Identity Credential*** may be required to utilize a security clearance ***Entitlement***. It should be noted that the ***Entitlement*** *granting* ***Organization*** and the ***Identity Credential*** *issuing* ***Organization*** may or may not be the same organization.

K. Management

The ***Management*** concept is a specialization of an ***Assignment*** that adds a many-to-many relationship with ***Position*** (see Figure 5). Essentially, a ***Position*** can *be managed by* multiple ***Management*** objects (managers), and managers can *manage* multiple ***Positions*** (workers). The *management level* focuses the scope and degree of control (e.g., project manager, technical lead, or performance evaluator).

V. MAPPING ATTRIBUTE CONTRACTS TO THE FEDERATED ICAM SUBJECT ATTRIBUTE CONCEPT MODEL

Using the canonical ICAM model to examine existing and emerging federal attribute contracts provides a reasonable estimate of its ability to successfully support policy-base access control for federated users.

We mapped the HSPD-12 Federal ICAM BAE, DoD Defense Manpower Data Center (DMDC) Enterprise Attribute Service (EAS), DoJ Global Federated Identity and Privilege Management (GFIPM), DHS Federal Emergency Response Official (F/ERO), and DHS Homeland Security Information Network (HSIN) attribute contracts to the attributes and concepts in the canonical ICAM model. This illustrates the concepts important for access control decisions in each sponsoring enterprise, greatly facilitates mappings between these enterprises, and also highlights both the access control concepts that are not being utilized as well as contract attributes and concepts that are either out-of-scope (such as the inclusion of an ***Emergency Contact*** in the BAE contract) or that embed access control decisions (privileges) as access control attributes (such as Copy, Read, Update, Delete (CRUD) actions in GFIPM).

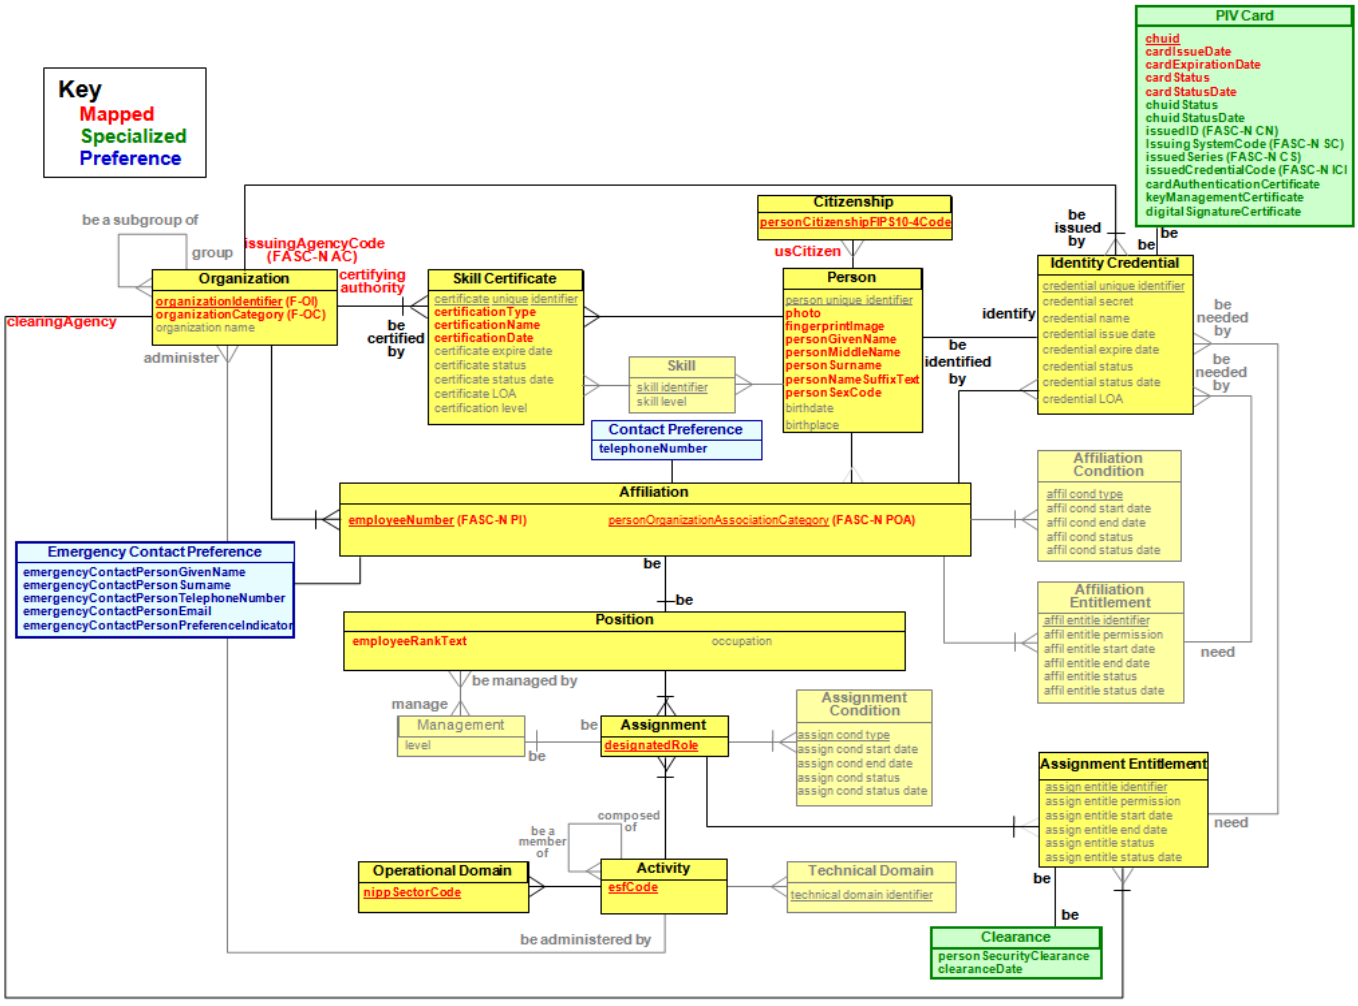

Figure 6. Federal ICAM BAE Mapping.

A. *Federal ICAM Backend Attribute Exchange Concept Mapping [8]*

The HSPD-12 Architecture Working Group (AWG) developed a BAE Architecture and Interface Specification document on behalf of the Office of Government wide Policy and the HSPD-12 Executive Steering Committee. It identifies 38 attributes that support and extend the information contained on the FIPS 201 Personal Identity Verification (PIV) Card, which is an ***Identity Credential*** (see Figure 6). They were developed for four use cases and identify 16 on-card and 22 backend attributes. These concepts map directly into the canonical model. This mapping illustrates that the BAE attribute contract identifies a single ***Organization*** and ***Activity*** and recognizes a ***Clearance*** as a specialization of an ***Assignment Entitlement***. Attributes for the ***Person***, ***Citizenship***, ***Position***, ***Affiliation***, ***Assignment***, ***Activity***, ***Operational Domain***, and ***Skill Certificate*** concepts are also included. The attribute BAE contract does not include ***Position Entitlement***, self-asserted ***Skill***, ***Management, Technical Domain, Affiliation Condition***, or ***Assignment Condition*** concepts. The BAE attribute contract also adds ***Contact Preference*** and ***Emergency Contact Preference*** concepts which are deemed not useful for policy based access control decisions in the concept model (both concepts are highlighted in blue in Figure 6).

B. *DHS Homeland Security Information Network (HSIN) Concept Mapping*[1]

The DHS has developed an LDAP schema for their next generation HSIN. These concepts map directly into the canonical model and illustrate that ***Skills*** and ***Skill Certificate, Citizenship, Operational Domain*** and ***Technical Domain, Positional Entitlement***, and ***Affiliation Condition*** concepts are not covered in the HSIN attribute contract. Communities of interest (COI) as specialized ***Activities*** are support as well as specific application ***Assignment Entitlements*** and their required ***Identity Credentials***. Attributes for the ***Person***, ***Position***, ***Assignment***, ***Assignment Condition***, and ***Organization*** concepts are also included. The HSIN attribute contract also adds ***User Contact Information*** and ***Sponsoring Organization Contact information*** concepts which are deemed not useful for policy based access control decisions.

C. *DoD Manpower Data Center (DMDC) Enterprise Attribute Services (EAS) Concept Mapping [10]*

The DoD DMDC is developing a Concept of Operations attribute set for the Global Information Grid 2.0 as an EAS extension to the existing 15 attributes supported by the current Real-time Broker System (RBS). Their approach allows an

[1] Based on LDAP schema and other information furnished by DHS OPS HISN Next Gen.

individual to have multiple user (***Affiliation/Position***) and ***Organization*** attributes sets including a primary, or administering, instance. These concepts map directly into the canonical model. One clearance per ***Affiliation*** and ***Position*** is supported as an ***Entitlement***. The ***Identity Credential*** is specialized to support CAC/PIV cards. Attributes for the ***Person***, ***Assignment***, ***Assignment Entitlement***, and ***Affiliation Condition*** concepts are also included. The mapping illustrates, that ***Skill*** and ***Skill Certificate***, ***Operational Domain*** and ***Technical Domain***, ***Assignment Condition***, and ***Management*** concepts that are not covered.

D. DHS Federal Emergency Management Agency (FEMA) Federal Emergency Response Official (F/ERO) Concept Mapping [2]

The F/ERO repository adds Emergency Support Function (ESF) (***Activity***) and National Infrastructure Protection Plan (NIPP) Sector codes (***Operational Domain***) to PIV identities (***Identity Credentials***). Limited auditing is also captured. These concepts map directly into the canonical model and illustrate that ***Skills*** and ***Skill Certificate***, ***Technical Domain***, ***Position*** and ***Assignment Conditions***, ***Citizenship***, ***Position Entitlement*** and ***Assignment Entitlement***, and ***Management*** concepts are not covered by the current F/ERO attribute contract.

E. DOJ Global Federated Identity and Privilege Management (GFIPM) Concept Mapping [11]

The Global Federated Identity and Privilege Management (GFIPM) initiative is supported through joint funding of DHS and the DOJ Office of Justice Programs, with collaboration from the Bureau of Justice Assistance and the National Institute of Justice (NIJ). GFIPM is built on the Internet2 open source Shibboleth federation capability and provides a standard set of security attributes about users' identities, privileges, and authentication details to members as a basis for trust. It defines 220 user, 10 resource, 3 action, 2 environment, and 32 entity attributes. The ICAM model addresses only user attributes and maps 132 of them directly into the canonical representation and illustrates that everything except the ***Assignment Condition*** and ***Technical Domain*** concepts are included in the GFIPM attribute contract. The GFIPM contract also includes C***ontact Informatio***n, ***Emergency Contact Information***, and ***Point of Contact*** attributes, but these are not recognized access control concepts.

Of the 220 user attributes, 88 define pre-determined boolean privileges. Having policy decision persisted in the PIP (Policy Information Point) rather than determined by the PDP (Policy Decision Point) is inherently problematic from an ABAC policy, management, and scaling perspective. As noted above, attributes in an ABAC solution should include the intrinsic attributes that establish policy requirements and not policy decisions.

[2]This information is based on slides and discussions provided by the FEMA Office of National Capital Region Coordination (NCRC) Craig Wilson and Toni Cieri.

VI. CONCLUSION

The complete canonical Federated ICAM Subject Attribute Concept Model establishes a technology-neutral representation of the policy-based user access control data components and their relationships. The model focuses on the underlying attributes of a potential Federal Enterprise Attribute Authority. It identifies and characterizes the relevant informational concepts and features that support policy-based access control decisions for users, representing a normalized view of the identity and authorization attributes that can be utilized for fine-grained ABAC decisions. It is an informational model that is uncoupled from specific policies and processes. As a conceptual model it provides a framework for logical and then physical access control models. It also provides a framework for creating access control policies. Lastly, in addition to providing semantic alignment, mapping existing access control policy attributes into the conceptual model makes explicit the concepts being utilized for access control decisions, and which are not being utilized.

ACKNOWLEDGMENT

Thanks to <insert tech writer here> for her technical writing help.

REFERENCES

[1] Adaptive Access Control Emerges, Allan & Perkins, Publication Date: 11 August 2009/ID Number: G00169295.

[2] Defining User Attributes For Authority-Based Access Control, Waterman & Hammer, COI/OAT/DHS, HSHQDC-06-C-00110, May 15, 2007.

[3] http://www.iso.org/iso/country_codes/background_on_iso_3166/what_is_iso_3166.htm.

[4] http://www.itl.nist.gov/fipspubs/fip10-4.htm

[5] Generalized and detailed data models: seeking the best of both worlds, J. Maguire, 2-06-2009, Gartner/Burton Group.

[6] Data Modeling: A Necessary and Rewarding Aspect of Data Management, J.Maguire, 5-18-2008, Gartner/Burton Group.

[7] http://csrc.nist.gov/publications/nistpubs/800-87/sp800-87-Final.pdf

[8] HSPD12 Backend Attribute Exchange Architecture and Interface Specification, Version 1.0.0, May 15, 2008.

[9] Managing Identities of Contract Workers, Version 1, G. Gebel, L. Rowland. Gartner/Burton Group, March 3, 2009.

[10] DOD Enterprise Attribute Services Concept of Operations, DRAFT version 1.1, DMDC, 3-29-10.

[11] DRAFT GFIPM Metadata 2.0 Specification, Global Security Working Group, April 1, 2010.