Attribute Metadata and Confidence Scoring

Workshop: Applying Measurement Science in the Identity Ecosystem
Version: 1, December 16, 2015

Information Technology Laboratory, NIST

1 INTRODUCTION

This document serves as a primer for discussions held at the "Advanced Identity Workshop" at the National Institute of Standards and Technology (NIST) in Gaithersburg, Maryland, on January 12 and 13, 2016. That workshop will convene federal agencies, commercial relying parties, and identity solution providers to collaborate on improving standards, practices, and policies around identity management.

1.1 PURPOSE

Business rules and organizational policy, specifically authorization policy, increasingly depends on the evaluation of a subject's attributes to make risk-based decisions. Enterprises that leverage automated decision support systems rely on attributes to enable a broad range of essential business functions. As enterprise domains continue to expand, architectures become further distributed, and business relationships become more complex, external entities increasingly provide these attributes. Organizations need methods for evaluating the trustworthiness of an asserted attribute in order to increase their ability to properly enforce policies.

This document introduces some of the technical, policy, and implementation considerations associated with the development of a schema by which federal agencies and other organizations could make risk-based authorization decisions that are informed by confidence in attributes.

Specifically, this document explores the following topics:

1. Defining standardized attribute metadata that organizations can use to support business decisions.
2. Establishing a scoring framework and its associated components to enable standardized attribute confidence scores.

1.2 SCOPE

This document and the forthcoming workshop will focus on addressing those concepts that apply to enabling organizations to make access and authorization decisions based on attributes provided by either internal or external entities. The document also introduces the concept of leveraging attribute metadata to enable organizations to convey the privacy requirements associated with accepting and using attribute information, as well as how the attribute information was obtained for a given subject.

Throughout this white paper, the term *attribute* is used to refer to characteristics of an entity (human or device)—often referred to as *subject attributes* or *entity attributes*. While this document does not

discuss the topics of object, contextual, and environmental attributes, the methodologies and frameworks described herein will be capable of interacting with these attributes within an access control model. NIST intends to explore the extensibility of the metadata and scoring methodologies described in this document to other attribute types at a later time.

NIST acknowledges that the addition of even a single element of metadata to any exchange or online interaction carries with it substantial scalability considerations. As a first step, this paper is focused on the development of attribute metadata to support federated identity assertions and to gain insight into the core elements required to establish confidence in attributes. Additional work will be required to determine the viability of deploying attribute metadata at scale across the heterogeneous IT systems that manage and steward identity attribute information.

2 WORKSHOP FOCUS AREAS

The content that follows introduces a set of attribute metadata and considerations for an attribute confidence framework. NIST recognizes challenges that face the development of any such construct as well as the possibility that attribute scores may not support any greater degree of confidence or interoperability. Regardless of the ultimate architecture, the goal is to establish a repeatable and adaptable process for understanding the degree of confidence that can be placed in attributes.

NIST requests that readers consider the following questions as they review this document and prepare for January's workshop:

Attribute Metadata:

- Is the proposed categorization of attribute metadata appropriate? Are the categories at the right level of granularity? Are there other categories that could more effectively segment attribute metadata?
- Are the metadata elements presented here appropriate? What additions, subtractions, or modifications would support decision-making based on attributes?
- Are concepts such as *pedigree* adequately captured by the recommended metadata values or are further descriptions required? Are the parameter values for metadata elements, such as *authoritative* or *sourced*, sufficient?
- Does the inclusion of *verification* and *pedigree* adequately address the issue of *binding* as a process for tying an attribute to an individual entity? Is there another way that this can be captured or represented in metadata? If included, how should a *binding* element be described and to what degree must it be aligned with verification and identity proofing practices?
- In what ways would the addition of attribute metadata impact an organizations infrastructure, operations, and performance? How "deep" into infrastructure would attribute metadata need to be injected? For example, would Human Resource Management Systems or relational databases need to be updated to support a metadata construct for attribute confidence?

Attribute Confidence Scoring:

- Is an attribute confidence scoring framework needed to support interoperability and trust within and across communities? Are standardized metadata elements and syntax sufficient to support these objectives?
- Do operational scoring structures exist that can be leveraged as examples for implementing an attribute confidence-scoring framework?
- How best could the framework integrate integrity protections, such as digital signatures, into confidence scoring?
- Is an attribute provider score separate and distinct from the attribute confidence score? Or, could this same information relating to trust in the provider be more efficiently captured through an expanded set of *provenance* metadata that includes the criteria that may otherwise have contributed to an attribute provider score?"
- How best can contractual, reputational, and other forms of "out-of-band" trust complement an attribute confidence scoring structure? Can a relationship over time with an attribute provider obviate the need for an attribute scoring framework?

3 ATTRIBUTE METADATA

Attribute metadata are granular elements of information about a single instance of an attribute. A relying party (RP) can use these metadata to better understand the applicability of an attribute for specific uses. These elements could include information about the pedigree of the underlying attribute itself (e.g., its authoritativeness), the processes used to create or establish the attribute (e.g., whether it is self-asserted or retrieved from a record), or the attribute's value (e.g., how often it is updated). Organizations can evaluate this metadata and apply organizational rules to make determinations of authorization to resources or benefits. Regardless of the access control methodology leveraged by an organization, integrating attribute metadata into decision support systems can enable more informed policy evaluation.

Defining the appropriate metadata and valid syntax, along with values for that metadata, may increase interoperability and a common understanding of attribute confidence across a broad ecosystem of RPs and attribute providers (APs).

In an initial effort to create a federal policy and framework for attribute exchange and management, the Federal Identity, Credential, and Access Management program (FICAM) convened federal stakeholders to explore the development of a standardized set of attribute metadata.[1] In reviewing the results of this effort and those of the Identity Ecosystem Steering Group (IDESG) and Open Identity Exchange (OIX),

[1] General Services Administration, Federal Identity Credential and Access Management. *FICAM Attribute Management Roadmap.* April 30, 2015. http://idmanagement.gov/sites/default/files/documents/FICAM%20Attribute%20Management%20Roadmap_20150430_FINAL.pdf

NIST proposes the set of attribute categorizations and metadata that appears in section 3.1.[2] Appendix A presents a mapping to other comparable efforts.

In collaboration with stakeholders across government, industry, and academia, as well as individuals, NIST would like to determine if the attribute metadata set, and its associated categories, meets organizational needs, and to identify gaps that may exist when applying attribute metadata across a broad range of access control use cases.

3.1 Proposed Metadata Categorization

While attribute metadata may have many uses for RPs, some metadata are more commonly tied to specific types of decisions. To facilitate RP decision-making and increase interoperability, we have created three categories for metadata based on common uses of metadata—accuracy, currency, and provenance—as well a fourth category for metadata with a less common or evolving role in RP decision-making:

- **Accuracy**- Metadata relevant or pertaining to the RP's ability to determine if the attribute is correct and belongs to a specific entity.
- **Currency**- Metadata relevant or pertaining to the RP's ability to determine the "freshness" of a given attribute.
- **Provenance**- Metadata relevant or pertaining to the RP's ability to evaluate the source of the attribute's value.
- **Other**- These metadata elements may provide additional information required under specific business or regulatory requirements, convey usage restrictions to organizations, or enable an overall understanding of the underlying attribute. These metadata elements are particularly important within the context where items such as *individual consented* could be used to convey and support alignment with privacy principles, objectives, or controls.

Each category of metadata elements is important for enabling the federation of attributes across a community or environment. Metadata associated with *accuracy*, *currency*, and *provenance* may facilitate cross-system trust. As these may not be the only attribute characteristics important to an RP or community, metadata in the *other* category may enable a common understanding of an attribute and how it may be leveraged within a specific context.

Table 1 presents the four categories and their definition.

Metadata Category	Description
Provenance	Metadata relevant or pertaining to the RPs ability to evaluate the source of the attribute's value
Accuracy	Metadata relevant or pertaining to the RPs ability to determine if the attribute is correct and belongs to a specific entity
Currency	Metadata relevant or pertaining to the RPs ability to determine the "freshness" of a given attribute

[2] The Open Identity Exchange, *Attribute Exchange Trust Framework Specification: Technical Specification V1.* July 02, 2013. http://openidentityexchange.org/wp-content/uploads/2014/06/OIX-AXN-Trust-Framework-Specification-1.0-7-5-2013.pdf. References to OIX are informative only and do not represent NIST endorsement of their products or services. OIX is a non-profit trade association.

Other	Those metadata elements which support interoperability of attributes by enabling standardized understanding of attribute metadata, acceptable uses, and specific business requirements

Table 1: Attribute Metadata Categories

3.2 PROPOSED METADATA

Within each category, the set of attribute metadata elements should be sufficiently complete to ease decision-making, while lightweight enough to ease consumption and avoid degraded system performance. NIST proposes an initial set of 13 metadata elements: two in the *accuracy* category, five in the *provenance* category, and three each in the *currency* and *other* categories, presented in Table 2 through Table 5 below.

Metadata	Description	Rationale
Verifier	The entity that verified the attributes value. **Origin**- Verified by the entity that issued or created the attribute value **Provider**-Verified by the attribute provider **Not verified**- Not verified	Included so the RP knows if the attribute's value has been verified, and if so, by whom.
Verification Method	The method by which the attribute value was verified as being true and belonging to a specific individual. **In-person**- The attribute was verified during an in-person session **Record verification**- The attribute value was verified against an electronic database **In-person with record verification**- The attribute value was verified both in person and against an electronic database **Not verified**- The attribute value has not been verified	Included so the RP to knows the method by which the attribute's value has been verified.

Table 2: Accuracy Metadata Elements

Metadata	Description	Rationale
Last Update	The date and time when the attribute was last updated. This metadata is used to derive the age of the attribute. **Date (time, day, month, year)**	Included to provide the RP with an understanding of how current the attribute value is. Specifically this enables the RP to determine if the date of last refresh is sufficient for a specific attribute or use.
Update Frequency	The frequency the Attribute Provider will refresh the attribute. **Real Time** **Daily** **Weekly** **Monthly** **Annually** **Never**	Included to provide the RP with an understanding of how current the attribute value is. Specifically this enables the RP to understand the rate at which an attribute is updated or refreshed.
Expiration Date	The date an attribute's value is considered to be no longer valid for its defined use. **Date (time, day, month, year) or none**	Included so the RP knows the date at which an attribute's value is no longer valid for its defined use. RPs may choose to accept attributes after they have been considered expired for their original intended use.

Table 3: Currency Metadata Elements

Metadata	Description	Rationale
Origin	The entity that issues or creates the initial attribute value. **Origin's Name or None**	Included so the RP knows the organization that originally provisioned or captured the attribute value.
Provider	The entity that is providing the attribute. **Provider's Name or None**	Included so the RP knows who is providing the attribute value or an attribute assertion.
Provider Signature	Properly formatted digital signature[3] of the organization providing the attribute. **Signed** **Unsigned**	Functional requirement to demonstrate the integrity of the attribute's value being provided from the AP to the RP. Also provides a general understanding of the attribute's chain of custody.
Origin Signature	Properly formatted digital signature[4] of the organization that issued of created the attribute value. **Signed** **Unsigned**	Functional requirement to demonstrate the integrity of the attribute value from its origin through the provider to the RP. This may be the same as the provider signature.
Pedigree	Description of the attribute's relationship to the authoritative source of the value. **Authoritative**- Created by source of authority **Sourced**- Collected by provider from one or more sources **Self-Asserted**-Asserted by an entity about themselves	Included to enable the RP to understand the authoritative nature of the attribute and the process by which it was established.

Table 4: Provenance Metadata Elements

Metadata	Description	Rationale
Individual Consented	Captures whether the user has consented to providing the attribute. **Yes** **No** **Unknown**	RPs may have specific privacy or regulatory reasons for understanding whether a user consented to the release of a specific attribute. This element enables organizations to meet that business requirement.
Description	A description of the attribute.	Provides RPs with an understanding of the attribute in order to support proper application in business decisions.
Acceptable Uses	A description of the acceptable business uses to which the attribute can be applied.	Provides RPs with an understanding of what business cases the attribute value can be used to support.

[3] Specified in Federal Information Processing Standards Publication 186-4, *Digital Signature Standard (DSS)*. http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf

[4] Ibid.

Table 5: General Metadata Elements

4 CONFIDENCE SCORING FRAMEWORK

To increase the utility of the attribute metadata and facilitate risk-based access and authorization decisions, organizations may use a scoring framework to determine confidence in the underlying attribute.

NIST has identified several possible approaches to increase trust and interoperability. While each of these requires the development of standardized attribute metadata, they vary in substantive ways from one option to the next. RPs could use these approaches as a baseline in making decisions, or can adjust accordingly based on risk tolerance.

4.1 DETERMINING METADATA ELEMENT SCORES

Establishing individual scores for each metadata element may increase the utility of these elements for determining and communicating overall confidence in the underlying attribute. The values for various metadata will vary in type and format (e.g., binary, text value, date, range). As a result, any scoring system would have to take into account the acceptable values for each identified piece of metadata as well as an appropriate weighting system to ensure that no piece of data is inadvertently elevated or depreciated in importance due to the assignment of scores. Moreover, in accepting or applying scores within an access control system, each RP could determine the attribute information that is important to its organization based on implemented risk assessment practices.

Scoring based on standardized metadata would involve the assigning of numeric values to metadata values. For example, when assigning scores to *verification method,* the acceptable values of {*not verified*, *record verification, in-person verification, in-person with record verification}*, could equate to ordinal values (i.e., 1, 2, 3, and 4), respectively, or scalar values (e.g., 0, 0.2, 0.8, 1).

4.2 DETERMINING OVERALL ATTRIBUTE CONFIDENCE SCORES

Scores for metadata elements can serve as inputs for quantifying and calculating a score that serves as a proxy for the confidence in the underlying attribute. By assigning scores to those metadata elements of concern, an RP, community of interest, federation, or trust framework has several options for developing an overall score for the underlying attribute. For instance, the overall score could be assigned as (equations are illustrative in nature):

- **Aggregated Score:** The aggregate of each metadata element. This overall score could then be compared by the RP to an established risk scale associated with their systems, applications, or other resources. In such a scheme, the RP could aggregate the individual element scores to enable organizational specific weighting, or the AP could do so based upon an established community framework for scoring. E.g.,

$$\textit{Attribute Confidence Score} = \textit{Origin Score} + \textit{Provider Score} + \textit{Pedigree Score} + \textit{Verifier Score} + \textit{Verification Method Score} + \ldots$$

- **Weighted Aggregate:** The weighted aggregate of each metadata element. Communities of interest, RPs, or even specific applications could apply weights (*a, b, c, …*) to each metadata element when producing the overall score. Attribute providers could provide the individual element scores—an RP would then apply the weighting—or the overall score in accordance with an established community framework for scoring and weighting. E.g.,

  Attribute Confidence Score = a(Origin Score) + b(Provider Score) + c(Pedigree Score) + d(Verifier Score) + e(Verification Method Score) + …

- **Category Score:** The aggregate of scores within each metadata category, without aggregating to an overall score. For example, RPs could evaluate an *Accuracy*, *Provenance*, and *Currency* score when determining confidence. RPs could also apply weights to these categories by, for instance, assigning a zero weight to currency for a birth date, as this is unlikely to change over time. In this scheme, the RP could develop the category scores after receiving the metadata elements from the AP or the AP could provide category scores. Weighting would either be done at the RP or by the AP in accordance with a defined scoring framework specified in a community framework. E.g.,

  Attribute Confidence Score = Accuracy Score, Provenance Score, Currency Score

- **Weakest Link**: The lowest score among any metadata element. The overall attribute score could follow a weakest link approach and assign overall score based on the lowest individual score of the metadata elements in the set. E.g.,

  Attribute Confidence Score = Min{Origin, Provider, Pedigree, Verifier, Verification Method, …}.

4.3 Inclusion of Attribute Provider Confidence

Determining confidence with this framework requires establishing trust in both the individual attribute itself and the organization providing the attribute. Certain elements regarding the AP are explicitly incorporated through the suggested metadata included in this paper—specifically those that relate to the attribute's *provenance.* Elements such as *pedigree, provider signature,* and *origin signature* can provide RPs with a certain degree of confidence both in the information management practices of a provider as well as insight into the chain of custody associated with a specific attribute. However, there may be certain situations in which RPs or communities may choose to implement more substantial processes for evaluating and establishing confidence in a given provider. There are several different ways to achieve this, which are discussed in greater detail below. Ultimately though, each of these approaches would likely be conveyed through the individual *provenance* metadata of attributes and considered by RPs as part of the attribute confidence scoring structure.

- **Attribute Provider Statements**: The FICAM Attribute Management Roadmap ("the Roadmap") provides a set of initial criteria for Attribute Provider management, security, and privacy practices which are intended to be captured through a questionnaire (an example of which is included in the Roadmap), and subsequently embodied in an *attribute provider statement*. RPs would then determine whether to enter into a relationship with the AP based on an evaluation

of the AP's *statement*. In these instances, the attribute's *provider* metadata could be evaluated against a white list to determine whether or not it is acceptable for use in an access control decision. Additionally, attribute provider statements could be presented in both human-readable format and machine-readable portions of the *provenance* metadata.

- **Trust Frameworks & Federations:** Today, certification of Identity Providers is often carried out through trust frameworks and identity federations to facilitate trust across a community of interest. These same processes could be established to increase AP trust through defined evaluations and multilateral legal agreements. Organizations could then append trustmarks and certification information to the attribute through its *provenance* metadata, enabling RPs to evaluate the attribute's value as well as the trustworthiness of the provider.

- **AP Scoring**: A structure could be set up to enable the scoring of attribute providers based upon evaluation of a defined set of characteristics or practices. This could be managed by a trust framework, federation, or even be done in house by RPs and factored into their access control policies and the evaluation of an attribute's appropriateness for authorization decisions. For example, a system may have a minimum threshold for attribute provider scores and attribute confidence scores based on the risks associated with a particular system (e.g., high risk systems require attributes with a particular minimum AP score and a different minimum attribute confidence score). Once scored, APs could pass their score as *provenance* metadata with the attribute. In some cases, the most appropriate source of an attribute for a specific application or instance may not be the AP with the highest overall score. When more than one attribute provider is available for use in an access control situation, RPs or communities of interest must properly weight metadata to ensure the most suitable choice to meet its needs.

Whatever approach is taken, the importance of capturing and conveying trust in the AP is essential to determining the appropriateness of a specific instance of an attribute for organizational needs.

5 CONCLUSION

The attribute metadata categorization, elements, and confidence scoring framework considerations presented in this paper will serve as one of three topics during January's Advanced Identity Workshop. NIST intends to leverage the outputs from the workshop, stakeholder feedback, and existing artifacts (such as the FICAM Attribute Management Roadmap and OIX Attribute Exchange Network Trust Framework) to produce a NIST report. Such a report would address the considerations of applying an attribute metadata and attribute confidence scoring framework and a proposed implementation approach. Feedback to this paper and discussion of this topic at the January workshop will be critical to determining the content of that report.

Appendix A: Attribute Metadata Mapping

Attribute Metadata Mapping				
Metadata	**Description + Value**	**FICAM**	**OIX**	**Comments**
Verifier	The entity that verified the attributes value. **Origin-** Verified by the entity that issued or created the attribute value **Provider**-Verified by the attribute provider **Not verified**- Not verified	None	Verification Method	Included so the RP knows if the attribute's value has been verified, and if so by whom. No directly comparable metadata element in FICAM. Originated from OIX "verification method" metadata element; retitled and modified recommended values.
Verification Method	The method by which the attribute value was verified as being true and belonging to a specific individual. **In-person**- The attribute was verified during an in-person session **Record verification**- The attribute value was verified against an electronic database **In-person with record verification**- The attribute value was verified both in person and against an electronic database **Not verified**- The attribute value has not been verified	None	None	Included so the RP to knows the method by which the attribute's value has been verified.
Last Update	The date and time when the attribute was last updated. This metadata is used to derive the age of the attribute. **Date** (time, day, month, year)	Last Update	Currency / Refresh	Included to provide the RP with an understanding of how current the attribute value is. Specifically this enables the RP to determine if the date of last refresh is sufficient for a specific attribute or use.
Update Frequency	The frequency the Attribute Provider (AP) will refresh the attribute. **Real Time**- at least every 12 hours **Daily** **Weekly** **Monthly** **Annually** **Never**	Last Update	Refresh Rate	Included to provide the RP with an understanding of how current the attribute value is. Specifically this enables the RP to understand the rate at which an attribute is updated or refreshed. Aligned with "last update" (FICAM) and "refresh rate." No values were included in FICAM documentation, suggested values align with the OIX "refresh rate" values.

Attribute Metadata Mapping				
Metadata	**Description + Value**	**FICAM**	**OIX**	**Comments**
Expiration Date	The date an attribute's value is considered to be no longer valid for its defined use. **Date** (time, day, month, year) or none	None	None	Included so the RP knows the date at which an attribute's value is no longer valid for its defined use. RPs may choose to accept attributes after they have been considered expired for their original intended use.
Origin	The entity that issues or creates the initial attribute value. **Origin's Name** or **None**	None	None	Included so the RP knows the organization that originally provisioned or captured the attribute value.
Provider	The entity that is providing the attribute. **Provider's Name** or **None**	None	None	Included so the RP knows who is providing the attribute value or an attribute assertion.
Provider Signature	Properly formatted digital signature[5] of the organization providing the attribute. **Signed**; **Unsigned**	None	None	Functional requirement to demonstrate the integrity of the attribute's value being provided from the AP to the RP. Also provides a general understanding of the attribute's chain of custody.
Origin Signature	Properly formatted digital signature[6] of the organization that issued of created the attribute value. **Signed**; **Unsigned**	None	None	Functional requirement to demonstrate the integrity of the attribute value from its origin through the provider to the RP. This may be the same as the provider signature.
Pedigree	Description of the attribute's relationship to the authoritative source of the value. **Authoritative**- Created by source of authority **Sourced**- Collected by provider from multiple one or more sources **Self-Asserted**-Asserted by an entity about themselves	Attribute Derivation, Source	Data Type	Included to enable the RP to understand the authoritative nature of the attribute and the process by which it was established. Aligns generally with FICAM's "attribute derivation" and "source" elements, as well as OIX's "data type." Suggested values differ somewhat from the source documents to convey several different options that provide insight into the process used to generate the attribute.
Individual Consent	Captures whether the user has consented to providing the attribute. **Yes**	Consent	None	RPs may have specific privacy or regulatory reasons for understanding whether a user consented to the release of a specific attribute. This element enables organizations to meet that business requirement.

[5] Specified in Federal Information Processing Standards Publication 186-4, *Digital Signature Standard (DSS)*. http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf

[6] Ibid.

Attribute Metadata Mapping				
Metadata	**Description + Value**	**FICAM**	**OIX**	**Comments**
	No **Unknown**			
Description	A description of the attribute.	Description	None	Provides RPs with an understanding of the attribute in order to support proper application in business decisions.
Acceptable Uses	A description of the acceptable business uses to which the attribute can be applied.	None	None	Provides RPs with an understanding of what business cases the attribute value can be used to support.
Not Included		Chain of Custody	None	This attribute metadata element was not included as NIST intends to address this through the inclusion of "origin" and "provider" elements, as well as future recommendations around digital signatures and other integrity protections.
		Classification Level	None	This attribute metadata element was intended by FICAM to relay the classification required to view the attribute's value. NIST chose not to include at this time since it seems better addressed through access control measures on specific systems.
		None	Availability	This attribute metadata element was not included as it would be better addressed in an attribute provider statement or service agreement. Is not specific to an individual attribute, but an attribute service instead.
		None	Geographic Coverage	This attribute metadata element was not included as the geographic coverage of the attribute provider was deemed to be inconsequential to establishing confidence in the attribute's value. May be appropriate for business considerations in selecting, but not the confidence of the attribute's value itself.
		None	Coverage Amount	This attribute metadata element was not included as the coverage provided by the AP was deemed to be inconsequential to establishing confidence in the attribute's value. If an AP is unable to provide an attribute for an entity within a certain population, then there is no attribute to have confidence in. This items seems more appropriate as business consideration for engaging with an AP.