

# Attribute Assurance Considerations

Whitepaper Draft

28 March 2014



## Contributors:

Arthur Friedman, NSA

Vincent Hu, NIST

Richard Kuhn, NIST

Paul Jacob, Booz Allen Hamilton

Bryan Marzoli, Booz Allen Hamilton

Rebecca Nielsen, Booz Allen Hamilton

This page intentionally left blank

# 1 Overview

## 1.1 Purpose

Attribute Based Access Control (ABAC)<sup>1</sup> has evolved as the preferred logical access control methodology in the Department of Defense and Intelligence Community in recent years, as well as many other agencies across the federal government<sup>2</sup>. Within ABAC, attributes are used to make critical access control decisions, yet standards for attribute assurance have just started to be researched and documented. This whitepaper outlines factors influencing attributes that an authoritative body must address when standardizing attribute assurance and proposes some notional implementation suggestions for consideration.

## 1.2 What is meant by Attribute Assurance

Attribute Assurance brings a level of confidence to attributes that is similar to levels of assurance for authentication (e.g., guidelines specified in NIST SP 800-63 and OMB M-04-04). There are three principal areas of interest when considering factors related to Attribute Assurance.

**Accuracy** establishes the policy and technical underpinnings for semantically and syntactically correct descriptions of Subjects, Objects, or Environmental conditions.

**Interoperability** considers different standards and protocols used for secure sharing of attributes between systems in order to avoid compromising the integrity and confidentiality of the attributes or exposing vulnerabilities in provider or relying systems or entities.

**Availability** ensures that the update and retrieval of attributes satisfy the application to which the ABAC system is applied. In addition, the security and backup capability of attribute repositories need to be considered.

Similar to a Level of Assurance (LOA), a Level of Attribute Assurance (LOAA)<sup>3</sup> assures a relying party that the attribute value received from an Attribute Provider (AP) is accurately associated with the subject, resource, or environmental condition to which it applies.

An Attribute Provider (AP) is any person or system that provides subject, object (or resource), or environmental attributes to relying parties regardless of transmission method. The AP may be the original, authoritative source (e.g., an Applicant). The AP may also receive information from an authoritative source for repackaging or store-and-forward (e.g., an employee database) to relying parties or they may derive the attributes from formulas (e.g., a credit score). Regardless of the source of the AP's attributes, the same standards should apply to determining the LOAA.

## 1.3 Level of Assurance Vectors

While there are parallels between LOAAs and credential Levels of Assurance (LOAs), there are also some special considerations that apply to attributes. Similarities to credential LOAs include attributes that remain unchanged during the life of a digital credential. However, some attributes

---

<sup>1</sup> Gartner recently predicted that “by 2020, 70% of enterprises will use attribute-based access control (ABAC) as the dominant mechanism to protect critical assets, up from less than 5% today.”

<sup>2</sup> Examples of guidance include NIST Special Publication 800-162, *Guide to Attribute Based Access Control (ABAC) Definition and Considerations* and Intelligence Community Policy Guidance (ICPG) 500.2, *Attribute-Based Authorization and Access Management*.

<sup>3</sup> LOAA could also be referred to as a Measure of Confidence (e.g., MOC1 through MOC4).

may change regularly or over time. The frequency of refresh, for example, is a consideration in LOAA that does not have a corresponding model in conventional credential LOAs.

Noted in public literature are some discussions distinguishing between Validity and Ownership Assurance. While there is certainly merit in the distinction that could be incorporated into additional analysis on assurance levels, keeping the LOAA defined at a single level will sufficiently handle the majority of use cases and additional complexity will slow adoption.

Note that while examples in this whitepaper use Identity (Subject) attributes for examples, the same concepts extend to all attribute types including Resources (Objects) and Environmental.

## 2 Accuracy

The confidence we have in an attribute's true value is affected by the care the AP takes in both obtaining the value and maintaining the value while in its possession. This whitepaper identifies two characteristics that influence accuracy.

- Attribute Source Due Diligence
- Attribute Integrity

### 2.1 Attribute Source Due Diligence

Similar to Identity Proofing, Due Diligence measures how well the AP identifies and validates the source of attributes. This applies regardless of whether the AP is the original source of the attribute or is acquiring the attribute from another source.

In Due Diligence, we make a distinction between truthfulness or consensus on the attribute's value and authoritativeness of information. The focus needs to be on how confident we are that the attributes represent the underlying entity, resource or condition. For example, we may strongly disagree with a specific credit score, but may be confident that it did come from a specific credit reporting agency. The lowest level of confidence is self-reported information from a person or non-person entity that has not been independently vetted.

**Table 1 Due Diligence Examples**

Due Diligence Examples		
Low	Medium	High
<ul style="list-style-type: none"> <li>• Self-Reported</li> <li>• Third-party Public Source</li> </ul>	<ul style="list-style-type: none"> <li>• Identity Proofing (medium)</li> <li>• Authenticated Source</li> </ul>	<ul style="list-style-type: none"> <li>• Derived Attributes (independent of underlying factors – original source)</li> <li>• Identity Proofing (high)</li> <li>• Authenticated Source with SLAs</li> </ul>

### 2.2 Attribute Integrity

Attribute Integrity determines how securely the AP provides attributes to Relying Parties similar to Credential Strength and Validation Path considerations for PKI credentials. In other words,

how does the AP assure that the attribute that it intends to send is the attribute the relying party actually receives? This involves evaluating both Data-at-Rest and Data-in-Transit security.

For Data-at-Rest, we need to evaluate the security framework for the actual attribute store and how well the AP protects the information or attribute-generation processes in the attribute store. Factors or capabilities that need to be evaluated include:

- Is file or whole-disk encryption employed?
- Is object-level encryption employed?
- What measures are taken to detect unintended alteration of attribute values?

For Data-in-Transit, we need to evaluate how the AP authenticates the relying party and how securely the attributes and their values are transmitted. Factors or capabilities that need to be evaluated include:

- Level of assurance of authentication credentials. For example, password tokens vs. high-assurance digital credentials.
- Security protocols are used for transmitting both attribute requests to the AP and attribute values to the relying party? For example, transmitting in the clear without encryption versus PK-enabled TLS sessions.

In addition to the AP's own data storage and transmission security, we must also consider the security arrangements in a second-tier receipt of attributes – i.e., when one relying party forwards attributes as an AP to another relying party. The level of assurance with respect to attribute integrity should be the lowest of the first-tier and second-tier providers.

For higher levels of assurance, digitally-signed attributes (Crypto-binding) provides a hash of the attribute so that relying parties can be confident that an attribute was not altered or tampered with before it is received and has not been accidentally or maliciously changed while in the relying party's possession. This would also raise the assurance level in second-tier transmission of attributes.

### 3 Interoperability

Sharing of information is critical to mission success as well as collaboration between government Departments and Agencies with industry partners. Interoperability standards and protocols that all entities agree to are essential to enabling this cooperation. Agreed-upon standards in both attribute syntax and semantics must be developed to ensure successful interoperation of systems.

#### 3.1 Syntax

How successfully we share attributes depends on how universally standards are developed and adopted. Many standards exist now but not all are widely used. For example, SAML is supported by the Organization for the Advancement of Structured Information Standards (OASIS) and has a high level of adoption. Including SAML 1.1 and 2.0 in the list of standards would be a priority. Additional standards need to be surveyed against their current and planned adoption to develop a proposed set of supported standards.

In addition to specific representation standards, standards and protocols governing authentication and transmission must be tied to assurance levels. Ensuring the integrity of an attribute's value

and its metadata and keeping it free from tampering or corruption is required for a higher level of assurance. For example, an AP may adequately vet its stored attribute information and provide a high level of protection within its enclave, but if the attribute information is sent via an unencrypted and unsigned mechanism (email, HTTP) as opposed to a more secure mechanism (signed SAML assertion, TLS, etc.), then the assurance level drops. Examples are found in Table 4.

### 3.2 Semantics

Given the broad spectrum of entities that will interoperate with each other, synonyms and homonyms are inevitable. Standard dictionaries tied to industry-specific namespaces need to be published. In addition to standardizing attribute definitions, metadata about attributes needs to be standardized to support assurance levels and to drive risk-based authentication decisions. The “metadata about metadata” could include some individual assurance level components that are incorporated into the overall LOAA.

**Table 2 Semantics Example**

Entity Applicability	Person
Name	Clearance
Value	Secret
Level of Attribute Assurance	1
Assurance detail - Due Diligence	Self-Reported
Assurance detail - Refresh	Pulled
Assurance detail - Last updated	12/31/2013
Attribute source	USAJOBS.gov

Provenance information could be captured and shared. An example of this would be the “attribute source” metadata sample. In the example above, the specific “Person” attribute may be sufficient for pre-loading data into a request form, but insufficient for access to a sensitive system since the clearance level is self-reported and not drawn from an authoritative source.

## 4 Availability

A relying party needs information on how often an attribute’s value is pulled or obtained as well as how securely it is processed to have confidence in the AP’s ability to provide the attribute’s value when it is needed. This whitepaper identifies two characteristics that influence availability.

- Refresh
- Operations

### 4.1 Refresh

Unlike conventional credentials, we want to measure how often attribute values are updated or validated. There are two vectors that need to be considered in measuring the impact of a refresh rate on a specific attribute. The first is **proactive** acquisition. We want to know, for example, whether the information is being passively pushed from another source to the AP or whether the

attribute values are being pulled on a schedule proactively. Independently of the frequency – which is also important – pulling attribute values on a schedule or on demand gives us assurance of how current and, therefore, how applicable the attribute value may be to an authentication or authorization decision. The second is **sensitivity** to change. For example, eye color or height is a relatively stable personal attribute but someone’s weight and, therefore, their appearance could change more often or more quickly.

The following table is a notional mapping of LOAA to Due Diligence considerations.

**Table 3 Refresh**

Proactivity →	High	Medium	High	High
	Medium	Low	Medium	High
	Low	Low	Low	Medium
		High	Medium	Low
		Sensitivity →		

## 4.2 Operations

Operations measures how secure the AP’s internal processes and procedures are with respect to both intentional attacks and unintentional errors or failures. The key document that governs the effect of Operations on LOAA should be an Attribute Practice Statement (APS). The highest level of assurance would be an APS that is audited for compliance with policy. Lower levels of assurance would apply to APs that self-report adherence to policy or who do not publish their operation’s practices.

A notional APS was developed for the Identity Ecosystem Steering Group ([www.idecosystem.org](http://www.idecosystem.org)) and is included in Appendix A as an example of the factors that could be used for establishing the LOAA of an AP. It is based on Internet Engineering Task Force (IETF) RFC 3647 (“Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework”) and includes additional points that would apply to AP operations.

## 5 Assurance Level Mappings

In order to maintain consistency with previous published policy, directives and guidance regarding levels of assurance for identity credentials and authentication, below is a notional mapping between credential LOAs and attribute LOAAs.

Table 4 Assurance Level Mappings

	LOA	LOAA	Due Diligence	Refresh	Operations	Attribute Integrity
1	Little or no confidence in the asserted identity's validity.	Little or no confidence in the attribute value's validity.	<ul style="list-style-type: none"> <li>Self-reported</li> </ul>	<ul style="list-style-type: none"> <li>None</li> </ul>	<ul style="list-style-type: none"> <li>No published practices</li> </ul>	<ul style="list-style-type: none"> <li>No Data-at-Rest encryption</li> <li>No Data-in-Transit encryption</li> <li>Receipt by unsigned email</li> </ul>
2	Some confidence in the asserted identity's validity.	Some confidence in the attribute value's validity.	<ul style="list-style-type: none"> <li>Self-reported from validated source</li> </ul>	<ul style="list-style-type: none"> <li>Values pushed to provider</li> </ul>	<ul style="list-style-type: none"> <li>APS</li> <li>SLA</li> </ul>	<ul style="list-style-type: none"> <li>Data-at-Rest AES 128 protection</li> <li>Data-in-Transit Suite B protection</li> <li>Receipt by signed email</li> </ul>
3	High confidence in the asserted identity's validity.	High confidence in the attribute value's validity.	<ul style="list-style-type: none"> <li>Identity Proofing – medium</li> </ul>	<ul style="list-style-type: none"> <li>Values pulled on a schedule</li> </ul>	<ul style="list-style-type: none"> <li>APS</li> <li>SLA</li> </ul>	<ul style="list-style-type: none"> <li>Data-at-Rest AES 256 protection</li> <li>Data-in-Transit Suite B protection</li> <li>Signed SAML assertion</li> </ul>
4	Very high confidence in the asserted identity's validity.	Very high confidence in the attribute value's validity.	<ul style="list-style-type: none"> <li>Identity-Proofing - high</li> </ul>	<ul style="list-style-type: none"> <li>Values pulled on-demand</li> </ul>	<ul style="list-style-type: none"> <li>APS with audited policy compliance</li> </ul>	<ul style="list-style-type: none"> <li>Data-at-Rest AES 256 protection</li> <li>Data-in-Transit Suite B protection</li> <li>Signed SAML assertion</li> </ul>

## 6 Closing

Key to establishing assurance and interoperability standards for attribute assurance is the development of APSs. The level of confidence we have in attribute providers is often established on unverified assertions of validity that are not based on commonly agreed-upon standards. The act of developing an auditable APS will provide an impartial assessment of the AP's standards of operation and the confidence level we have in the provided attribute.

In addition, privacy is a growing area of concern and oversight that needs to be considered in policy and technical discussions to avoid unintentional exposure of personal information that would violate law or other codes of conduct. While not directly affecting the functional requirements of LOAA or attribute syntax, privacy needs to be included in policy guidance and influence practical implementations. While automated enforcement of privacy limitations is problematic, at a minimum, privacy standards should be incorporated into APSs. APSs should include standard verbiage for:

- Ownership – Who can decide how and when to share the attribute's value
- Personally Identifiable Information (PII) – Is the attribute or attribute collection considered to be PII?
- Selective relying parties – who may receive the information

APSs should be established at a minimum LOAA of 3 (High Confidence) for providing PII information based on the data-at-rest and data-in-transit encryption protection standards noted in Table 4.



Clearly, additional research and stakeholder outreach is necessary to begin mapping out a roadmap for developing an attribute assurance framework. Inclusion of government and commercial participants could identify specific use cases that would benefit from this framework and be used to guide further development.

## 7 Appendix A

### ATTRIBUTE PRACTICE STATEMENT

#### Draft Outline Loosely Based on RFC 3647

#### 1: INTRODUCTION

The focus of this section is to provide a high level overall introduction to the Attribute Provider's (AP) system. This section follows RFC 3647 with some modifications.

- **Overview:** provide a narrative description of the AP's overall operating parameters
- **Document Name and Identification:** Identify one or more policy object identifiers (OID) that the AP will use to identify attribute/value pairs it supports
- **AP Participants:** Identify participants who have trusted roles on the AP system, including the management authority under which the AP operations
- **Attribute Community:** Describe the intended community for which attribute/value pairs will be stored and shared by the AP
- **Relying Party Community:** Describe the intended community to which attribute/value pairs will be provided
- **APS Administration:** Provide name and contact information for the entity that will be administering the Attribute Practice Statement. If the AP will consume attribute/value pairs from any other APs, this section should include information for how those external AP's processes will be evaluated to ensure they meet the standards defined in this APS
- **Definitions and Acronyms:** Although this section shows up in section 1, having a long list of acronyms and definitions up front often breaks up the flow of the document, so recommend just having a reference here to appropriate appendices where this information can be found

#### 2: REPOSITORY

The focus of this section is to identify the attributes and the semantic meaning of the attribute/value pairs that the AP will host, as well as high level requirements related to the AP repository. If the AP wants to have the ability to change the set of attribute/value pairs without having to formally update the APS, then this section should reference the document where this information can be found, and how that document will be kept in compliance with other sections of this document (specifically Section 3). This section follows RFC 3647 with some modifications.

- **Attributes and Values:** Identify the attributes that the AP will publish in its repository. For each attribute, describe the semantic meaning of the attribute, the possible values for the attribute (which may be a specific list or a description of what the value will look like), and whether the attribute will always have a value, can only have one value, or can have many values

- **Repositories:** Identify general access restrictions on repositories. Also identify any commitments for repository availability. Reference Section 4 for more specifics regarding authentication and authorization for relying parties to access attribute/value pairs

### 3: VERIFICATION AND MAINTENANCE

The focus of this section is on the processes the AP uses to determine the values for attributes that it publishes. If different attributes are verified using different processes, all of the processes should be described so that a potential relying party can determine which attribute/value pairs meet its level of confidence requirements. If attribute/value pairs are sourced from an external AP and hosted by the AP described in this APS, describe those processes in this section. This section does not follow RFC 3647 but has been adapted to address requirements for attributes that are distinct than requirements for certificates.

- **Identification:** Describe how attribute/value pairs are linked to a specific digital identity, for example by listing the identifier that will be used by the AP and how this identifier is determined
- **Enrollment:** Describe how entity identifiers are enrolled by the AP
- **Attribute Value Determination:** Describe the process for determining attribute values
- **Attribute Value Verification:** Describe the process used to verify the correctness of attribute values
- **Attribute Value Refresh:** Describe the frequency that attribute values are re-verified to maintain correctness, and the process used to refresh the values
- **Attribute Deletion:** Describe the circumstances and process for deleting attribute/value pairs from the AP repository

### 4: ACCESS MANAGEMENT

The focus of this section is on defining who is authorized to access attribute/value pairs and the processes the AP uses to authenticate authorized relying parties. If different attributes have different access policies, describe each policy and which attributes it applies to. If entities have any ability to manage who may access their attributes, also describe how this is done. This section does not follow RFC 3647 but has been adapted to address requirements for attributes that are distinct than requirements for certificates.

- **Who May Request Attribute Values:** Describe who is authorized to request attribute/value pair information.
- **Validation of Authorized Relying Parties:** Describe how relying parties are authorized to request attribute/value pair information.
- **Authentication of Relying Parties:** Describe how authorized relying parties are authenticated prior to providing attribute/value pair information.
- **Relying Party Access Removal:** Describe the circumstances and process for removing relying party authorizations.

## 5: FACILITY, MANAGEMENT, AND OPERATIONAL CONTROLS

The focus of this section is to define the facility, management, and operational controls instituted by APs. This section follows RFC 3647 but omits subsections that are not applicable to APs.

- **Physical Controls:** Describe physical controls including site location and construction, physical access, power and air conditioning, water exposures, fire prevention and protection, media storage, and waste disposal. Also describe backup requirements.
- **Procedural Controls:** Describe procedural controls including trusted roles, number of persons required per task, identification and authentication for each role, and roles requiring separation of duties
- **Personnel Controls:** Describe qualifications, experience, and clearance requirements, background check procedures, training requirements, retraining frequency and requirements, job rotation frequency and sequence, sanctions for unauthorized access, independent contractor requirements, and documentation supplied to personnel
- **Audit Logging Requirements:** Describe security audit requirements including types of events recorded, frequency of processing log, retention period of audit log, protection of audit log, audit log backup procedures, audit collection system (internal or external), notification to event-causing subject, and audit log assessments
- **Records Archival:** Describe the types of records archived, the retention period for the archive, protection of the archive, archive backup procedures (if any), requirements for time-stamping of records, archive collection system (internal or external), and procedures to obtain and verify archive information
- **Compromise and Disaster Recovery:** Describe incident and compromise handling procedures, recovery when computing resources, software, and/or data are corrupted, and business continuity capabilities after a disaster
- **Termination:** Describe notification, archive, and other processes that will be implemented in the event that the AP terminates operations, including what happens to information contained in the AP repository, backup, and archive.

## 6: TECHNICAL SECURITY CONTROLS

The focus of this section is to identify technical security controls implemented by the AP. This section follows RFC 3647 but omits subsections that are not applicable to APs.

- **Computer Security Controls:** Describe specific computer security technical requirements and computer security rating (if applicable)
- **Life Cycle Technical Controls:** Describe system development controls, security management controls, and life cycle security controls
- **Network Security Controls:** Describe network security controls
- **Time Stamping:** Describe any time stamping requirements and mechanisms used

## 7: ATTRIBUTE SYNTAX AND DISTRIBUTION STANDARDS

The focus of this section is to identify the technical standards and formatting used for attribute/value pairs, and the technical mechanisms supported by the AP for providing attribute/value information to relying parties. This section does not follow RFC 3647 but has been adapted to address requirements for attributes that are distinct than requirements for certificates.

- **Attribute Syntax:** Describe the syntax used by the AP when providing attribute/value pairs to relying parties. Where possible, leverage existing standards and reference these standards.
- **Attribute Distribution:** Describe the mechanism(s) used by the AP to provide attribute/value information to relying parties. Include information for how data integrity of information is preserved in transit (e.g., digitally signed SAML assertion, provided during authenticated SSL/TLS session)

## 8: COMPLIANCE AUDIT AND OTHER ASSESSMENTS

The focus of this section is on requirements for a periodic independent audit of the AP's activities to ensure they are complying with this APS. This section follows RFC 3647.

- **Frequency and Circumstances of Assessment:** Describe how often compliance audits are performed. Also describe the reasons that may lead to an aperiodic compliance audit
- **Identity/Qualifications of Assessor:** Describe the qualifications a compliance auditor must possess
- **Assessor's Relationship to Assessed Entity:** Describe the independence of the compliance auditor
- **Topics Covered by Assessment:** Describe the topics covered by the audit. All aspects of the APS should be covered
- **Actions Taken as a Result of Deficiency:** Describe the actions that will be taken if an audit deficiency is identified, including activities to address the deficiency
- **Communication of Results:** Describe how and to whom the results of the audit will be communicated, including any notifications that will be provided to relying parties or other entities

## 9: OTHER BUSINESS AND LEGAL MATTERS

The focus of this section is to address business and legal matters related to the operations of an AP service. This section follows RFC 3647 with some modifications.

- **Fees:** Identify fees that will be charged to entities for having their attributes verified and hosted by the AP, and fees that will be charged to relying parties for obtaining attribute/value pairs from the AP
- **Financial Responsibility:** Identify insurance coverage, other assets, insurance or warranty coverage for entities or relying parties, and fiduciary relationships

- **Confidentiality of Business Information:** Identify the scope of business information collected or managed by the AP, which information will be considered confidential or not confidential, and the AP's responsibility to protect business confidential information
- **Privacy of Personal Information:** Describe how the AP will maintain privacy of personal information contained in attribute/value pairs. Include information treated as private, information not deemed private, responsibility to protect private information, notice and consent to use private information, disclosure pursuant to judicial or administrative process, and other information disclosure circumstances. Since the purpose of the AP may be to disclose personal information to authorized relying parties, special attention should be placed on this section so that entities are aware of what rights they do and do not have with regards to information contained in attributes managed by the AP
- **Intellectual Property Rights:** Identify who owns intellectual property rights to information stored by the AP
- **Representations and Warranties:** Identify any representations and warranties provided by the AP to entities whose attributes are hosted by the AP, relying parties, or other parties
- **Disclaimers of Warranties:** Identify any disclaimers or other restrictions on warranties
- **Limitations of Liability:** Identify any liability limitations for the AP
- **Indemnities:** Identify any AP indemnities
- **Term and Termination:** Describe term, termination, and effect of termination and survival
- **Individual Notices and Communications with Participants:** Identify any individual notices and communications with participants
- **Amendments:** Describe the procedure for amendment, notification mechanism and period, and circumstances under which any identified attribute policy OIDs must be changed
- **Dispute Resolution Provisions:** Describe any dispute resolution provisions, especially any disputes that may arise with entities over attribute values
- **Governing Law:** Identify the governing law for the AP
- **Compliance with Applicable Law:** (Editor's note, this is in RFC 3647 but it's kind of silly, applicable law applies regardless of what this document says) State that the AP will comply with applicable laws
- **Miscellaneous Provisions:** Address assignment, severability, enforcement, and force majeure provisions
- **Other Provisions:** Identify any other provisions.

## 8 Appendix B

### REFERENCES

1. NIST Special Publication 800-63-1, *Electronic Authentication Guide*, December 2011
2. OMB M-04-04, *E-Authentication Guidance for Federal Agencies*, December 16, 2003
3. NIST Special Publication 800-162, *Guide to Attribute Based Access Control (ABAC) Definition and Considerations*, January 2014
4. ICPG 500.2, *Attribute-Based Authorization and Access Management*, November 2010
5. Internet X.509 Public Key Infrastructure Certificate Policy and Certification Practices Framework, 2003