

NISTIR 7657

**A Report on the
Privilege (Access)
Management Workshop**

NIST/NSA

Privilege Management Conference Collaboration Team

NIST IR 7657

A Report on the Privilege (Access) Management Workshop

NIST/NSA
Privilege (Access) Management Workshop Collaboration Team

March 2010



U. S. Department of Commerce
Gary Locke, Secretary

National Institute of Standards and Technology
Patrick D. Gallagher, Director

Certain commercial entities, equipment, or materials may be identified in this document in order to describe an experimental procedure or concept adequately. Such identification is not intended to imply recommendation or endorsement by the National Institute of Standards and Technology, nor is it intended to imply that the entities, materials, or equipment are necessarily the best available for the purpose.

Preface

This document is based on the discussions and conclusions of the Privilege (Access) Management Workshop held on 1-3 September 2009 at the Gaithersburg, Maryland, facilities of the National Institute of Standards and Technology (NIST), sponsored by NIST and the National Security Agency (NSA). This document includes additional material resulting from in-scope comments made by workshop participants and the public during the review periods for this document. An overview of the workshop is available in the published proceedings of the workshop. [NISTIR 7665 - Proceedings of the Privilege Management Workshop, September 1-3, 2009]

Participants at the workshop generally agreed that access management is the umbrella under which to consider privilege management. At the same time, many workshop participants felt that the term “privilege management” was not needed at all, since all aspects of the discussions held in the various tracks could be described without use of the term. Yet, the term “privilege management” was being used in several contexts, with differing meanings, and there was a strong desire to clarify its meaning. Contributing to the reason to use the term was the definition of “privilege management” that appeared in the draft document¹ produced by the Identity, Credential, and Access Management (ICAM) Subcommittee just months earlier [FICAM-09 - Federal Identity, Credential, and Access Management.] That proposed definition seemed to be closely related to the area being examined at the workshop. Also, the view of privilege management expressed in this document generally aligns with the architectural and service framework for privilege management presented in the FICAM document. Both the FICAM document and this report treat privilege management as a subset of access management.

The results of the workshop, as described in this report, show that the central topic of the workshop turned out to be attribute and policy management. Whether attribute and policy management should be called “privilege management” is an open question at this point. Looking at the definitions of “privilege management” in the FICAM document and in this report, it appears that they address different levels of concern in the area of identity, credential, and access management. The FICAM definition appears to view privilege management as a governance and business process, while this report’s definition focuses on computer-based management of attributes and policies. As the reader can easily discover, it is possible to substitute “attribute and policy management” for “privilege management” throughout this report without damage to the content. The question arises, then, as to whether a definition of “privilege management” as found in the FICAM extends to the area of access management covered in this report or should be limited to the governance and business process level. It remains for future deliberations, such as a follow-on workshop, to examine the issues involved and resolve such questions.

The discussion in this document is not comprehensive, dealing principally with those ideas, points, gaps, and concerns derived from presentations and discussions at the workshop. In particular, it does not address assurance issues associated with the topics covered because the workshop’s scope specifically excluded assurance considerations, in order to achieve a useful first step in exploring privilege management. We believe, however, that this report provides a good basis for further exploration of the topics and issues and, in particular, for a follow-on workshop.

¹ *Federal Identity, Credential, and Access Management (FICAM) Roadmap and Implementation Guidance*

Table of Contents

Introduction	1
A Context for Thinking About Privilege Management	2
Definitions and Standards	12
Access Control Methods	14
Basic Methods	14
Enhancements	15
State of the Practice	15
Considerations for Implementing Access Control	16
Policies and Requirements	19
Research Agenda	21
Conclusion	25
Summary	25
Recommendations	26
A Way Forward	26
Bibliography	28
Annex A: Authorization and Attributes Glossary	30
Annex B: A Survey of Access Control Methods	32
Annex C: Authoritative Attribute Source and Attribute Service Guidelines	33
Annex D: Advanced Capabilities for Privilege Management	34
Annex E: The Policy Machine	35
Annex F: Security Framework for Privilege Management	36

List of Figures and Tables

Figure 1. High-Level View of Access Control	3
Figure 2. High-Level View of Real-Time Access Control	3
Figure 3. Authentication Management and Privilege Management	5
Figure 4. High-Level View of Relationships as a Venn Diagram	6
Figure 5. Information Managed by Privilege Management.....	7
Figure 6. Interfaces of Privilege Management – View One	9
Figure 7. Interfaces of Privilege Management – View Two	10
Table 1. Factors to Consider for the Selection of an Access Control System.....	16

National Institute of Standards and Technology Interagency Report on the Privilege (Access) Management Workshop

Introduction

This National Institute of Standards and Technology (NIST) Interagency Report (NISTIR) on the Privilege (Access) Management Workshop is organized as follows:

- **A Context for Thinking About Privilege Management:** This section describes the full scope of enterprise-level access control and management, showing how privilege management fits under the umbrella of access management.
- **Definitions and Standards:** This section examines the need for definitions and standards, with a focus on eXtensible Access Control Markup Language (XACML).
- **Access Control Methods:** This section identifies current, distinguishable access control methods and focuses on the attribute-based access control method.
- **Policies and Requirements:** This section presents considerations about digital policy management.
- **Research Agenda:** This section identifies issues in several topical areas of privilege management, including policy and attribute management, standards, and several others.
- **Conclusion:** This section gives a brief summary of the document and provides a list of recommendations.
- **Bibliography:** This section provides references and other recommended reading.
- **Annex A: Authorization and Attributes Glossary**
- **Annex B: A Survey of Access Control**
- **Annex C: Authoritative Attribute Source and Attribute Service Guidelines**
- **Annex D: Advanced Capabilities for Privilege Management**
- **Annex E: The Policy Machine**
- **Annex F: An Alternate View**

A Context for Thinking About Privilege Management

This section describes enterprise-level access control and privilege management, both of which come under the umbrella of access management. At the enterprise level, access management encompasses all the practices, policies, procedures, data, metadata, and technical and administrative mechanisms used to manage access to the resources of an organization. Access management includes access control and privilege management as well as other related capabilities such as identity management. Considering things at the enterprise level ensures that all elements of privilege management are included so that the needs of all organizations, large and small, can be met.

Privilege management at the enterprise level is usefully viewed in relation to enterprise-level access control. *Access control* ensures that resources are made available only to authorized users, programs, processes, or systems by reference to rules of access that are defined by attributes and policies. *Privilege management* is the definition and management of attributes and policies that are used to decide whether a user's request for access to some resource should be granted. In this context, resources can be both computer-based entities (files, Web pages, and so on) and physical entities (buildings, safes, and so on), and users requesting access to resources can be people, processes running on a computer, or devices. Please note that this description of privilege management is a working definition for the purposes of this report. Any definition, to be an approved, agreed-upon term, must go through a formal review process by bodies such as the Authorization and Attribute Services Committee (AASC) and NIST. As noted in the Preface of this report, work needs to be done to formalize any terminology beyond that being proposed by organized committees such as the Identity, Credential, and Access Management (ICAM) Subcommittee and the AASC.

To have a clear notion of the meaning and scope of privilege management, we start by considering how access control works at a high level, as shown in Figure 1.

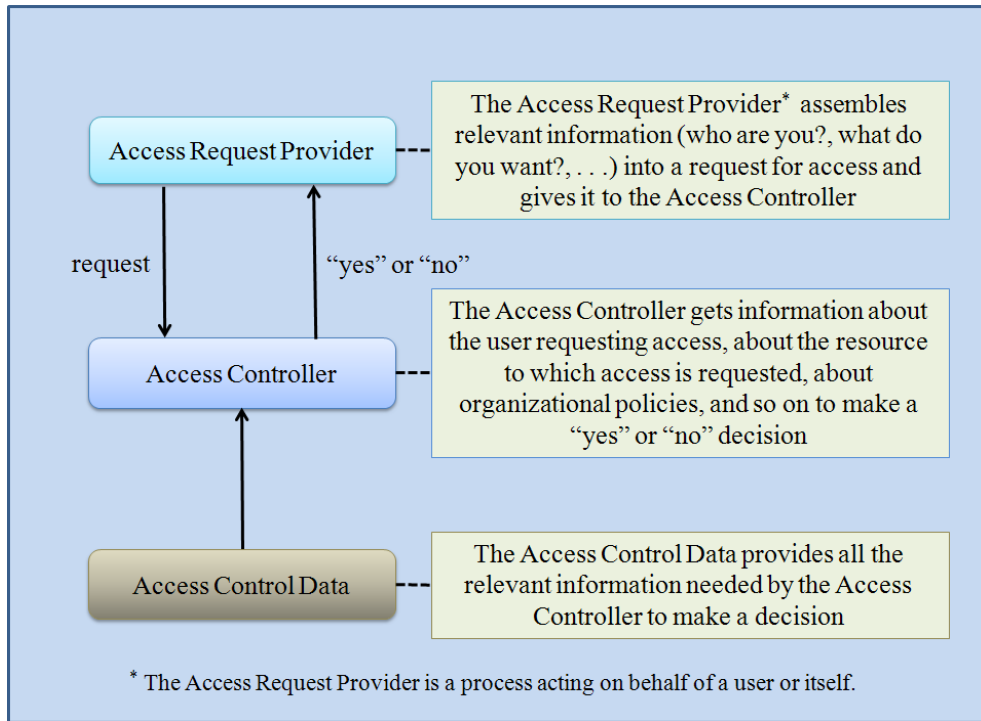


Figure 1. High-Level View of Access Control

Figure 2 depicts the real-time framework for access control in more detail, introducing terminology that is used in this report.

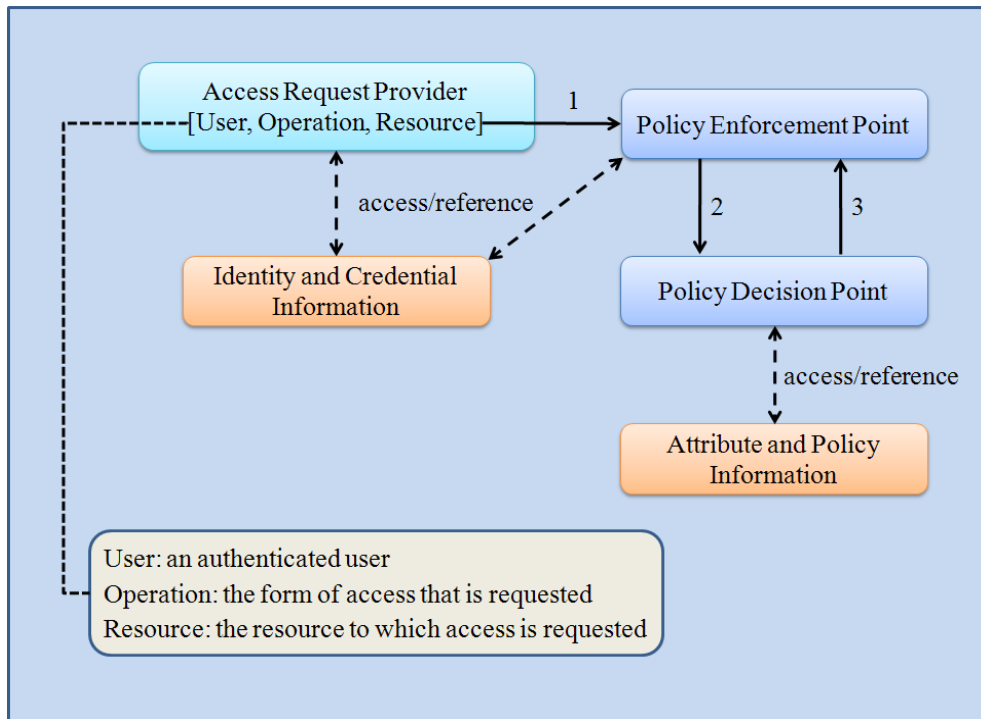


Figure 2. High-Level View of Real-Time Access Control

In Figure 2, the Access Controller of Figure 1 is split into two parts—Policy Enforcement Point and Policy Decision Point—and Attribute and Policy Information replaces Access Control Data. The meanings of the Policy Enforcement Point and Policy Decision Point (as described in Annex A: Authorization and Attributes Glossary) are as follows:

- **Policy Decision Point (PDP):** A *system entity* that makes *authorization decisions* for itself or for other system entities that request such decisions.
- **Policy Enforcement Point (PEP):** A *system entity* that requests and subsequently enforces *authorization decisions*.

“Attribute and policy information” is being used in a very general sense and is intended to have the same scope as access control data. Thus, it includes any form of information that can be used for access control. For example, it includes traditional access control lists (ACLs.) For an ACL, the attribute might be a group or user name while the policy² is implicit. In this context, “policy” denotes digital policy—policy that can be processed by computer. The rationale for this scope of the terminology is to enable discussion without having to deal with the details of the many forms of access control data, while at the same time distinguishing the main categories of access control data—attributes and policies.

A *user* is a person, process, or device. The term “user” is defined for the context of this report, as suggested by RFC 4949, and shares connotations of meaning with the terms “subject,” “system entity,” and “system user” as defined in RFC 4949 and with the terms “subject” and “user” as defined in CNSSI-4009. *Attributes* are distinguishable characteristics of users or resources, conditions defined by an authority, or aspects of the environment. Attributes might provide, describe, or be contact information, membership in communities of interest, roles within a community of interest, sensitivity of data, permission bits, location of the user or the resource, properties of the user session, conditions in the enterprise network or in the environment, priorities associated with individuals, status of resources, current bank account balance, and so on. *Policies* are rules that specify how to use attributes to render an access decision. A policy might specify that a user’s signature authority must equal or exceed signature-level-two in order for the user to authorize a monetary account withdrawal.

The view of real-time access shown in Figure 2 does not reflect assurance mechanisms and other entities that might enter into the activity, except for the high-level reference to credentials. So, for example, in a real system, the policy enforcement point might use a credential validation service to convert authorization credentials³ into attributes that it then provides to the policy decision point. As noted in the Preface, however, assurance is not being addressed in this document.

As suggested in Figure 2, the Access Request Provider uses identity and credential information that is relevant to the context in which the request is being made. The level of trust associated with the identity’s credentials can vary widely as can the form of the credentials, and the same holds true for attribute and policy information. In addition, the policy enforcement point may

² Policy: The process requesting access to the resource is allowed read access if the “read” permission bit is enabled.

³ Authorization credential is an attribute assertion digitally signed by the issuer so that it can be cryptographically validated. An attribute assertion is a statement made by an attribute authority that an entity possesses a particular set of attributes.

access (pull) credential information as well as having it provided to it (pushed) (see, for example, RFC 3281).

The question naturally arises, “Where does privilege management fit in this view?” The attribute and policy information must be created and maintained, and this is the business of privilege management. Similarly, the identity and credential information is the business of authentication management. This is depicted in Figure 3.

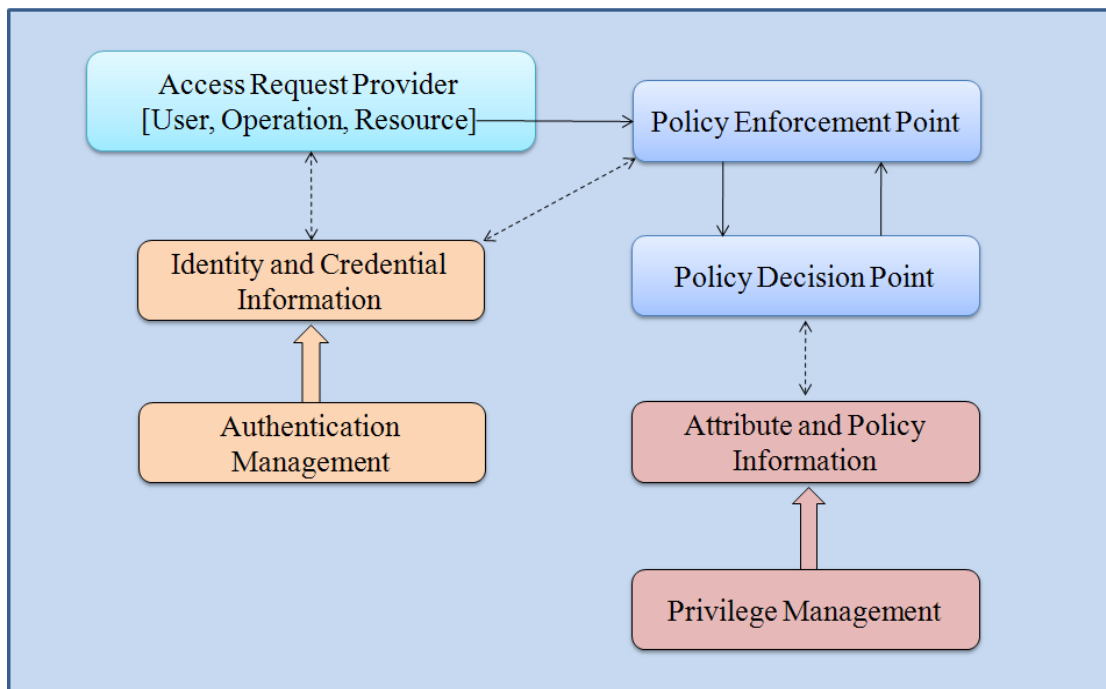


Figure 3. Authentication Management and Privilege Management

Authentication management deals with identities, credentials, and any other authentication data needed to establish an identity.⁴ Privilege management creates, stores, and manages the attributes and policies needed to establish criteria that can be used to decide whether a user’s request for access to some resource should be granted. Access control uses the data made available by authentication and privilege management, plus other information provided by the access request provider, such as the form of access requested, to make an access control decision.

Access management, which includes privilege management and access control, encompasses the science and technology of creating, assigning, storing, and accessing attributes and policies, of using those attributes and policies to decide whether a user’s request for access to a resource should be allowed or denied, and of enforcing the access control decision.

⁴ Assurance of integrity and authenticity of the information generally will be required as well; assurance is not addressed in this report.

In the foregoing description:

- “science” is included because many aspects of access management are amenable to scientific treatment, theory, and structure.
- “technology” is included because hardware and software for access management define actual usage and the state of the practice.

Figure 4 shows a high-level view⁵ of the relationships among access management, access control, privilege management, and attributes and policies.

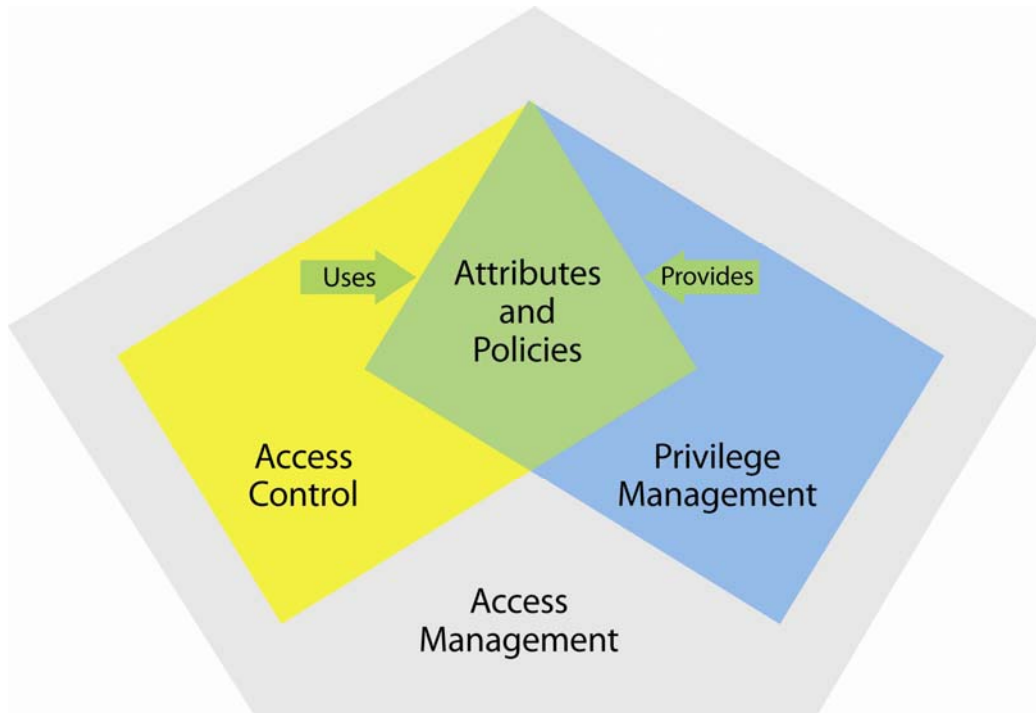


Figure 4. High-Level View of Relationships as a Venn Diagram

Privilege management, which is the focus of this report, can usefully be viewed in more detail, as shown in Figure 5.

⁵ Another, somewhat broader view of relationships, which goes beyond the scope of this report, is shown in the suggested figure on relationships in Annex F: An Alternate View.

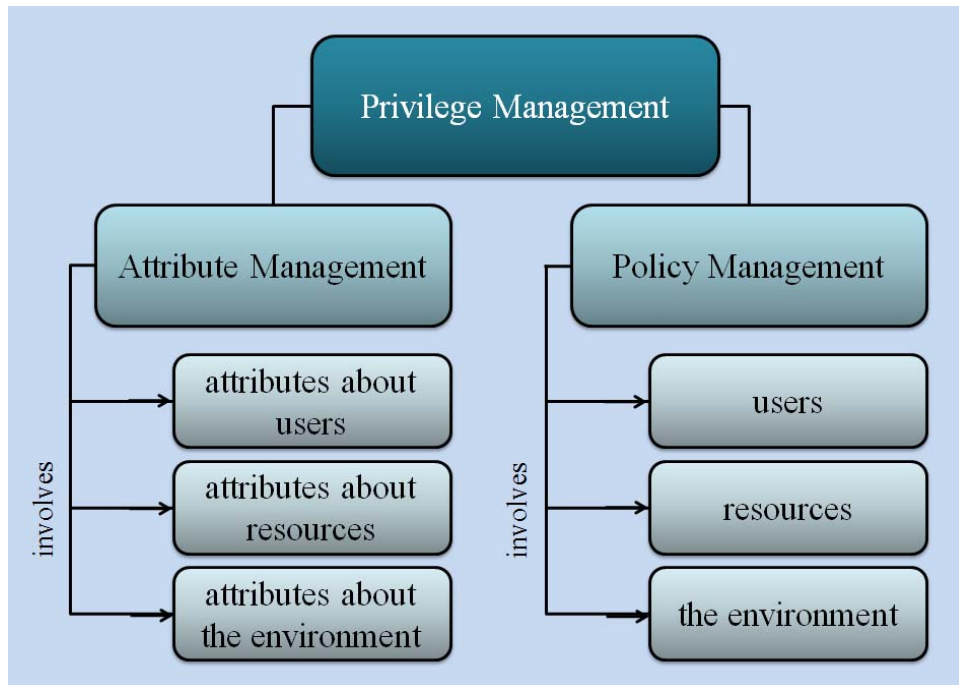


Figure 5. Information Managed by Privilege Management

As depicted in Figure 5, privilege management is conceptually split into two parts—attribute management and policy management—since they might have differing governance. Policy, for example, might be managed at a higher level of an enterprise than attribute management. They are also different kinds of entities or structures, to which different standards apply. Attributes are characteristics of entities, while policies are rules that specify how to use attributes in making an access control decision. Attributes describe aspects of people, resources, and environments. Some attributes are basic (who or what an entity is) and some might be enterprise-specific (the current balance in a savings account or the country that a client is in at the moment). Some attributes that we can think of may be beyond the state of the practice. Such attributes can be input to identification of needed research. Policies describe the rules that a policy decision point uses to determine whether a request should be granted. Policies can range from simple to complex. An example of a simple policy is the requirement that a requester’s clearance equal or exceed the classification of the requested resource. An example of a complex policy in a military context is the requirement that a requester’s clearance equal or exceed the classification of the requested resource, that the requester has a documented need-to-know with regard to that resource, that the Information Operations Condition [INFOCON] level⁶ is at 3, 4, or 5, and that the Defense Readiness Condition [DEFCON] level is at 3, 2, or 1. The same kind of policy, but in a banking context, might be stated as follows: A bank employee’s request for access to a customer’s account is granted if the bank employee is at least a branch manager, there is a certified record of need-to-know, and the account has not been locked due to a regulatory issue. Partitioning attribute management into three parts—management of user attributes, management of resource attributes, and management of environment attributes—is also a conceptual

⁶ See www.answers.com for descriptions of the INFOCON and DEFCON levels.

partitioning for two reasons: First, to make clear that attribute management must address all the elements that go into making a decision; and second, because the three elements—user, resource, and environment—are recognizably different entities. Similarly, policy management may involve rules about users, resources, the environment, or any combination to define access control policies.

Another way to decompose attribute/policy management is by functionality. Both attribute and policy management must provide governance, provisioning, and access.

- Governance: deciding which attributes and policies are needed to control access to what resources by which users. Since provisioning can be expensive, selection of attributes and policies should be done with careful consideration of the entities to be covered. Governance decisions involve judgment based on enterprise needs, weighing costs and benefits.
- Provisioning: defining the attributes and policies, including semantics and representations, and instantiating them.
- Access: providing for storage and retrieval of attributes and policies. This can be complex but is a generic data management task.

An important aspect of privilege management is the way it interfaces to the rest of the world and, in particular, to the access control components of an enterprise. Standards for interfaces are particularly powerful, providing flexibility in configuring functionality, wider choice of vendors, and upgradability without disruption when properly implemented.

In the next two figures, activity is partitioned into real-time activity and administrative-time activity. Another way to think of this partitioning is as shown earlier in Figure 4: access control uses and privilege management provides attributes and policies. Also, as shown in Figure 7, privilege management may participate in the real-time activity.

There are two basic ways to view the interfaces of privilege management within the context of access control. One way is shown in Figure 6, in which the interfaces are depicted as double-headed wide arrows inside the area marked as “administrative time activity.”

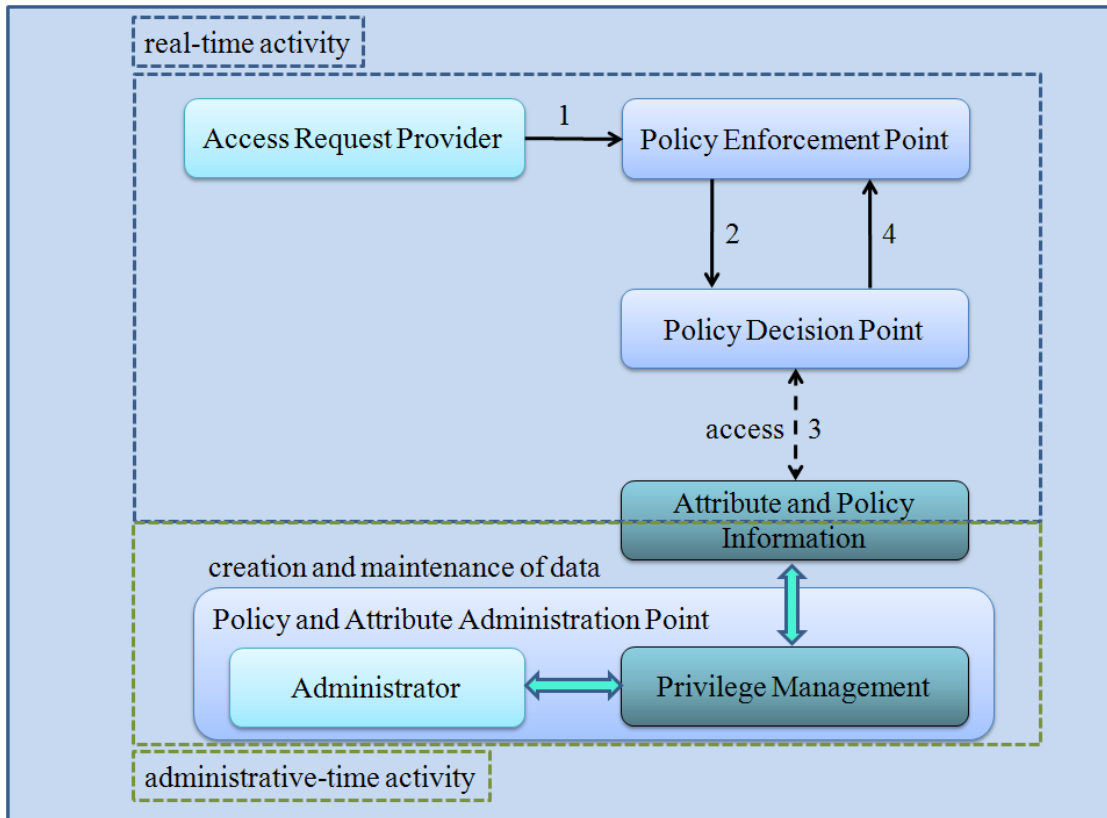


Figure 6. Interfaces of Privilege Management – View One

In this architecture, the policy decision point accesses the information it needs directly in the relevant information stores. That information is created, maintained, and placed there by the privilege management system. Thus, privilege management interfaces to the attribute and policy information repository, or repositories, but does not interface to the policy decision point. The other interface, between administrator and privilege management,⁷ supports the human-computer interface. The administrator is shown as a separate entity to make explicit the human-computer interface; as implemented, the administrator’s functions might be an integral part of the privilege management system, might be on a dedicated workstation, might be realized through a Web interface, and so on. Also, “administrator” should be understood to represent also a system user who can change access privileges on a system having a discretionary access control mechanism.

⁷ The XACML standard, Version 1.0, defines a Policy Administration Point as follows: “The system entity that creates a policy or policy set.” Since privilege management in this document deals explicitly with attributes as well, we use the term “Policy and Attribute Administration Point” to refer to the privilege management and administrator functions.

This architecture has features worth noting:

- The information needed by the policy decision point can be efficiently obtained; for example, all stores⁸ might be of the same type and might be directly, locally accessible.
- The management system has a single role to play; in addition, it can manage information in multiple stores of different types.
- Compared to the next architecture, fewer interfaces are involved; thus, fewer standards are needed.

A second way to view the interfaces of privilege management within the context of access control is shown in Figure 7. The interfaces are again shown as double-headed wide arrows.

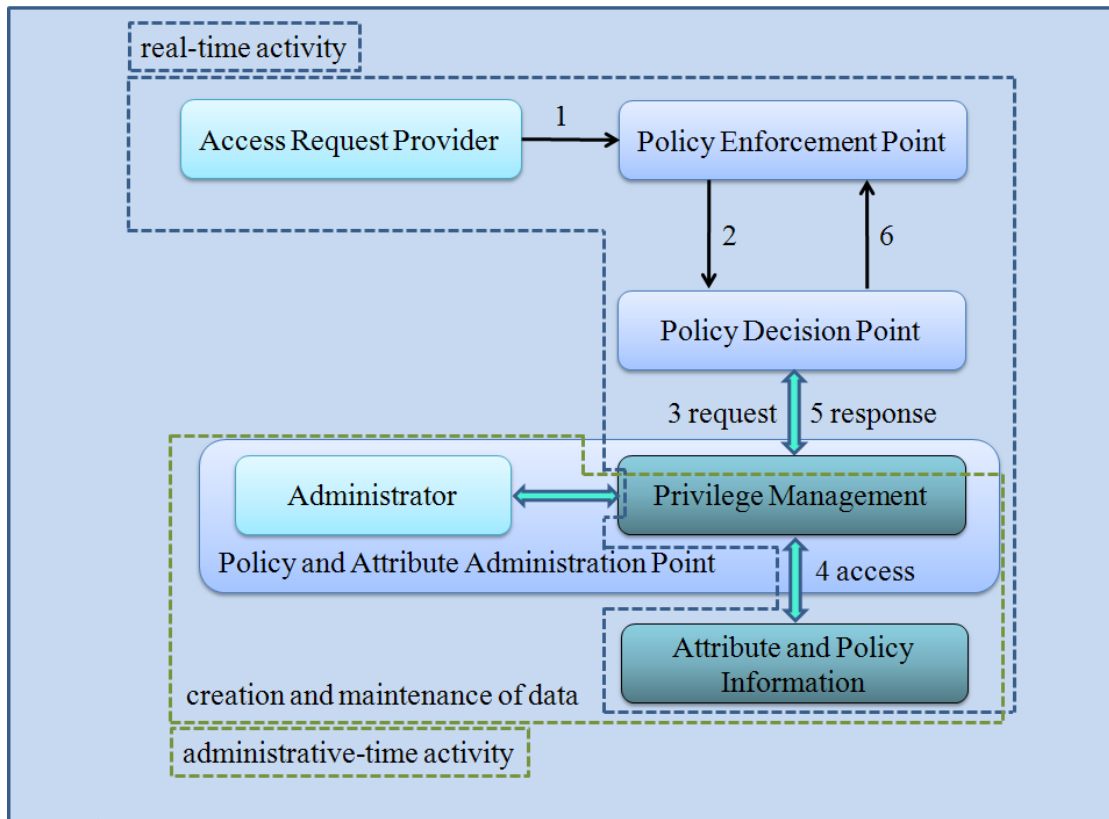


Figure 7. Interfaces of Privilege Management – View Two

In this architecture, the policy decision point sends a request for information to the privilege management system, which acts as a retrieval agent for the policy decision point.

⁸ Attribute and policy information can conceivably be stored in a number of repositories, each of which might differ from the others.

This architecture has features worth noting:

- The policy decision point need not have any knowledge of the stores⁹ in which the information it needs is kept; the stores can be of any type and information for a given class of information can even be kept in multiple, different stores.
- A single privilege management system can provide its services to multiple policy decision points.
- Real-time inputs to the policy decision point can also be provided by the privilege management system; in the first architecture, the policy decision point would have to obtain such inputs, complicating its functionality.

Advantages and disadvantages of these architectures depend on factors in the context in which the architectures are applied: for example, how they are implemented, what supporting infrastructure is available, and what organizational policies come into play. Implementation of either of these architectures, and any variants of them, would clearly benefit from appropriate interface standards. Standardized interfaces would facilitate integration of products from various vendors.

Managing and accessing attributes, which is central to any architecture, can be considered a data management capability. At the enterprise level, system management budgets, flexibility, robustness, performance, and correctness are major risk factors for organizations. The risk can be ameliorated by use of proven data management systems. Managing and accessing attributes at the enterprise level is functionality that generally should not be included in each policy decision point or privilege management system. Instead, like a directory, it can be provided by a data management capability that is shared among many policy decision points, as in Figure 5, and many privilege management systems, as in Figure 6. Such an approach insulates policy decision points and privilege management systems from the complexities of the potentially many stores supporting attribute management and access. Consumers—that is, policy decision points and privilege management systems—have a logical view, and the underlying data management system maps to the physical stores, caches, and services that return the attribute data. Management of attributes must also deal with the issues arising from multiple values of attributes having the same name. Again, translating and mapping multiple instantiations of the same attribute is better done as part of a data management capability rather than by individual policy decision points.

The foregoing discussion sets the stage for considering the following topics—definitions and standards, access methods, policies, and research agenda.

⁹ Although shown as a single entity in the picture, Attribute and Policy Information might be kept in multiple stores.

Definitions and Standards

The ability to control access to sensitive data in accordance with policy is a fundamental security requirement. Definitions enable discourse and standards facilitate implementation in support of enterprise-level access management and privilege management in particular.

Definitions generally in use, as in NISTIR-7298 Glossary of Key Information Security Terms and those explicitly referenced in Annex A: Authorization and Attributes Glossary, serve the purpose of effective discourse. However, no generally accepted definition of privilege management appears to be in use. The Identity, Credential, and Access Management (ICAM) Subcommittee of the Federal Information Security & Identity Management Committee is developing implementation guidance for access management, which includes a draft definition of privilege management, as follows:

“Privilege Management is the definition and management of policies and processes that define the ways in which the user is provided access rights to enterprise systems. It governs the management of the data that constitutes the user’s privileges and other attributes, including the storage, organization and access to information in directories.” [FICAM-09]

This definition is in accord with the view taken in this report; some architectural details in the referenced document relative to privilege management differ. Both this report and the referenced document view privilege management as one part of a larger access management capability.

The eXtensible Access Control Markup Language (XACML¹⁰), an OASIS¹¹ standard for managing access control policy, provides some of what is needed to support enterprise-level privilege management. It includes a policy language and a query language that results in a Permit, Deny, Intermediate (error in query), or Not Applicable response. XACML queries, which are typically in the Security Assertion Markup Language (SAML) format, are sent to a Policy Enforcement Point (PEP), located at the file server or Web server, which forms a request to the Policy Decision Point (PDP). The PDP determines the answer based on policy and sends back its determination to the PEP. Both the PEP and PDP might be the same application in the same server or distributed across a network.

XACML does not define the following:

- Creation and maintenance of policy;
- Policy enforcement;
- Attribute collection, maintenance, and retrieval;
- Resource attributes; and
- Authority delegation and trust management.

Both technical standards and processes for privilege management need further research and development. One area of concern is life cycle management and governance for attributes.

¹⁰ Released in 2003 and based on XML, the Sun-developed XACML was designed to become a universal standard for describing who has access to which resources.

¹¹ OASIS (Organization for the Advancement of Structured Information Standards) is a not-for-profit consortium that drives the development, convergence, and adoption of open standards for the global information society.

Consider the attribute class called “organization.” If free form text is allowed, one Air Force instantiation might be “USAF,” another might be “Air Force,” and another might be “US Air Force.” Policy logic cannot handle every possible value that might be used. What is needed is a manageable set of attributes and their values as well as a way to add new ones—in other words, governance. Without governance over attributes, policies that use them can become intractable.

Access Control Methods

This section discusses access control methods and systems and identifies a number of factors to consider when planning an enterprise's access control system. This report pays quite a bit of attention to access control even though the focus is privilege management. This is because the characteristics and features of an access control system determine the requirements for the associated privilege management system. At the same time, an organization's capabilities to perform privilege management functions, if they cannot meet those requirements, limit the capabilities of the access control system to what is feasible.

Practitioners and researchers in computer security generally distinguish among access control policies, mechanisms, and models. Policies are high-level requirements that specify how access is managed, expressing who or what should be granted or denied access to what resources. Mechanisms implement policies in a system. For example, policies might be captured as rules that a policy decision point uses, in addition to attributes about the requester and the resources, to determine whether an access request should be granted or denied. Models are used to describe and, in some cases to prove, security properties of an access control system. In short, a policy specifies access control requirements, a mechanism implements the requirements, and a model proves or describes things about the system that uses that policy and mechanism. The interested reader will find extensive discussion of these points in [FERR-07].

However, for our purposes in this report, we do not need to distinguish models and mechanisms from each other; it suffices to think about access control methods—that is, ways of doing access control. So, while an access control list (ACL) is a mechanism and role-based access control (RBAC) is a model, both are methods of access control. For the remainder of this report, then, we will refer to both models and mechanisms as methods of access control. Also, it is important to note that we use the term “policy” somewhat differently than general practice as described above. In other words, while in general usage “policy” is a high-level requirement, not necessarily expressed in computer-readable form, in this report we have defined “policy” as a rule that specifies how to use attributes to render an access decision, implying that the rule can be used by a computer.

Basic Methods

There are three access control methods that are clearly distinguishable from one another: identity-based access control (IBAC), role-based access control (RBAC), and attribute-based access control (ABAC). In the IBAC method, identity is the key determinant of whether an access request should be granted or denied. IBAC is often supported by access control lists (ACLs). In the RBAC method, the role of a requester is the key determinant for access. In the ABAC method, attributes associated with the requester, with the resource to be accessed, and with the environment or current situation are used in combination, as determined by relevant rules, to decide whether access should be granted. Although there has been much socialization of the term ABAC, there are still many variations in its description. Toward the end of standardizing its meaning, the following definition is offered:

Draft Definition of ABAC: ABAC denotes access control based on attributes and policies. *Attributes* are distinguishable characteristics of users or resources, conditions defined by an authority, or aspects of the environment, and *policies* specify how to use attributes to determine whether to grant or deny an access request.

IBAC is relatively simple to understand and implement on monolithic systems but does not scale well, especially with regard to large enterprises and cross-domain or federated situations. RBAC better supports implementation of least privilege and separation of duties but, like IBAC, does not scale well as the number of resources increases. Whatever access control can be defined with IBAC or RBAC can also be defined with ABAC. In addition, the ABAC method can provide more complex access control than can be accomplished with IBAC or RBAC. However, the ability to define more sophisticated access control comes with the additional administrative and managerial burdens imposed by complexity. With the ABAC method, the policies to be supported must be known in order to assess the trade-off between capability and complexity. All three methods have high-utility places within the access control space; their strengths and weaknesses should be considered when determining which to apply to a given set of requirements.

Additional background on some access control methods and systems can be found in [NISTIR-7316] and [FERR-07], and a survey of methods is presented in Annex B: A Survey of Access Control Methods.

Enhancements

Policy-based access control (PBAC) is frequently referred to as a distinct access control method. However, policy-based access control is inherent in ABAC since attributes without rules cannot serve to guide an access control decision. For example, simply knowing that a user has a signature level of Fiscal-2 does not allow a policy decision point to determine whether that user is allowed to access a financial record of type Fiscal-B. The policy decision point also needs to know how Fiscal-2 is related to Fiscal-B; that is, it needs a policy, expressed as a rule, to make that determination. For example, the rule might be **users with Fiscal-2 can read only Fiscal-A and Fiscal-B records**. If the user asks for read permission, the request is granted under this rule. If the user asks for modify permission, the request is denied under this rule. Although inherent in ABAC, PBAC as a method stresses the importance of policies and the potential power for expression they can provide for access control.

Recently, risk-adaptable access control (RAdAC) has been proposed as a means of adapting access control to changing conditions. It extends the ideas of PBAC by introducing environmental conditions and risk levels into the access control decision. For enterprises whose operations need to adjust to environmental or situational changes, RAdAC appears to be a method of great promise. However, RAdAC-based solutions cannot be strictly automatic but must involve user judgment and actions. Thus, policy and management capabilities will be of great importance to any RAdAC-based solution.

State of the Practice

In the commercial world, RBAC is the de facto access control implementation at the enterprise level because RBAC is what most solutions support. See [FERR-07] for a discussion of RBAC and enterprise-level security management. As discussed in [FERR-07], one obstacle to RBAC is the initial complexity involved in setting it up, a process known as role engineering—defining roles, user-role assignments, permission-role assignments, and role hierarchies. Some researchers have suggested rule-based approaches on user and resource attributes as a way of avoiding this obstacle; in other words, an ABAC approach. However, ABAC is considered the cutting edge of access control, and there is not much support for ABAC in the commercial space. It has seen slow adoption because its many-to-many relationships are difficult to represent. ABAC and

anything more complex or richer in policy is on the horizon for most organizations. Most products cannot represent or handle complex digital policies. Most implementations are identity-focused, localized, and not scalable to the enterprise level.

Considerations for Implementing Access Control

The ideas associated with PBAC and RAdAC make it clear that full realization of ABAC’s potential requires better attribute and policy management capabilities at the enterprise level. Policy needs to be reconciled across domains. For example, without reconciliation of policies, a person might be disallowed from accessing the contents of a file in one domain, but might still be able to gain access to the file’s contents indirectly through another domain that gives the person access. In a federated environment, whether an organization can enforce protections on its data depends on reconciliation of policy. In addition, policy needs to support hierarchies and it needs to be consistent across the enterprise. These are important concepts in which development should be invested. Although the ABAC method suffices for expressing access control, it depends on quality information and policies for effectively realizing an enterprise’s possibly complex requirements. Consistency in the meaning and use of attributes across the enterprise is important as well as ensuring that attributes come from an authoritative source. Significant progress in understanding the requirements for authoritative sources has been made; the reader is referred to Annex C: Authoritative Attribute Source and Attribute Service Guidelines.

Table 1 provides points to consider regarding the requirements for an organization’s access control system.

Table 1. Factors to Consider for the Selection of an Access Control System

Administrative review	Audit log review can be useful for discovering the source of errors, usage patterns, attempted policy violations, and so on; the access control system can play an important role in log generation, for example by logging granted and denied access requests.
Bypass	If some or all of the access control decisioning and enforcement will be done at the application level, is the risk of bypassing the mechanism ¹² commensurate with the risk tolerance of the enterprise?
Complexity versus simplicity	A good balance between complexity and simplicity of the access control system’s architecture provides what is needed in functionality at the lowest cost in mechanism. The simpler the architecture of the access control system, the less that can go wrong, the easier to identify and fix errors, and the lower the cost of making an access request.

¹² For example, bypass through authorized or unauthorized privileged operations in the operating system.

Delegation of administrative capabilities	It may be necessary or convenient for the access control system administrator to delegate privileges to other administrators. Will the access control system being considered make it easy and secure to do that?
Ease of administration	Consider whether the access control system will require additional technical support, for example, in order to use special languages such as XACML, or will a simple Graphical User Interface (GUI) suffice to compose and administer access control policy.
Existing standard	Standards can provide useful guidance in terms of usage and implementation. Is there a published standard that supports the model or mechanism of the candidate access control system?
Least privilege	Every user and process should have the least set of privileges needed to perform the task at hand. The implementation of this principle has the effect of limiting damage that can result from system error or malicious events.
Management	Consider management of the access control system during its life cycle: will there be a need to change access rules during operation, what will be the extent of policy changes/evolution, are equipment/software upgrades likely, and so on?
Resource/user discovery	It may be important to an enterprise's access control system administrator to be able to easily discover (by query, for example) which resources a given user has access privileges for or which users have access privileges for a given resource.
Operational/situational awareness	Will there be a need for the policy decision point to take into account operational/situational factors in making access control decisions? This sort of capability needs to be supported not only by the policy decision point but by the privilege management system as well.
Performance	Will the access control system be able to process user requests for access within a time that is consistent with the operational needs of the enterprise? This can be evaluated in part by the complexity of the decision-making algorithm, by modeling, and by prototyping.
Policy combinations	It may be convenient or necessary to be able to combine access control policies, for example combining separation of duty and workflow policies. Will the access control implementation enable this capability?

Policy support	The chosen access control method should support the enterprise's access control policy or policies that must be followed, such as mandatory access control, discretionary access, separation of duty, workflow, and so on.
Privilege management	Will privilege management capabilities be available to adequately support the envisioned access control system?
Safety	The access control implementation should include a mechanism for preventing leakage of privileges through either constraints or confinement. ¹³
Scope of control	The access control system must support the needed scope of control—for example, operating system, applications, single logical machine/agent, multi-machines/agents, single organization/enterprise, distributed multi-organization/enterprise, multi-domain environment, or federated environment.
Transparency	The protection system interface to the users should be easy to use so that the protection mechanism is accessed correctly. Any complexity in the protection system should be invisible to the user.

¹³ Definition and discussion of “confinement” and “constraints” is beyond the scope of this report.

Policies and Requirements

As suggested in the discussion of access control methods, policy capabilities for applications supporting both federal government and commercial enterprises are a key element for effective privilege management at the enterprise level. Creating and enforcing policies digitally can involve many attributes. Managing many attributes is problematic and how to represent the attributes in policy is complicated. Thus, organizations have a need to determine which attributes are essential and how to represent them. Attributes play very important roles in the more sophisticated access control methods, and it is important that they be understood and harmonized across the board to the extent possible. In federated or other sharing environments, organizations need to coordinate what attributes are used in intersecting or shared policies and how to represent and provision them. A person can have multiple identities and be a member of multiple groups, and groups can intersect on the attributes that they use. Hence, some attributes can be common and acceptable for use across multiple groups.

Developers of policies for an organization should carefully consider what attributes are needed to implement the policies and should document for future reference the process by which they make this determination. Developing good policies requires an understanding of users and a clear determination of whether and how laws and policy rules apply to them. It also requires understanding the characteristics of the data involved and the organization's need to share that data and with whom. Organizations that must share data need to express their sharing requirements in high-level agreements pursuant to their legal obligations, and then implement the policies and rules that match those agreements. In addition, protocol and interface specifications are critical to the harmonization of access control among organizations, especially in a federated environment.

Many organizations will also have a need to accommodate unanticipated users. Access management and privilege management based on the ABAC method can provide this capability. One of the advantages of the ABAC method is that it does not depend on knowledge of requesters in advance of their requests for access. If the attributes associated with a requester meet relevant criteria for access, access can be granted. Thus, ABAC is particularly useful for situations in which organizations or resource owners want unanticipated users to be able to gain access when they meet relevant criteria. This ability to determine access without the need for a predefined list of individuals is important in large enterprises where people might join or leave the organization arbitrarily.

When considering implementations of ABAC-based solutions, attention should be paid to classical operational metrics such as availability, throughput, reliability, and so on. An enterprise-level implementation of ABAC will be sensitive to network topography, latency, and throughput considerations. Consideration should also be given to the extent of administration required: for the relevant policies, this includes technical difficulty and amount of context to be understood in order to write the policies; for harmonization, the number of attributes and organizations involved; and for provisioning, the number of attributes and the number of resources involved.

Organizations that intend to enhance enterprise-level access management, including privilege management in particular, should carefully consider implementing all aspects of this space to the same level of capability to avoid the risk of having poor results because of weak links. For example, distributed access control based on the ABAC method is not likely to work well without adequate supporting infrastructure.

Research Agenda

This section presents issues or concerns, in several topical areas, that can guide the research agendas of government, industry, and academia. They arise from the preceding sections of this report. In most cases, gaps are identified but no specific research agenda is proposed, specific agendas being the proper province of researchers in industry, academia, and government.

Policy and Attribute Management

- A recommended general focus for research is policy management. Risk-adaptable policies are a desired ultimate goal but much work still needs to be done to provide foundations and infrastructure to realize their potential.
- Digital policies need to be able to specify rule precedence to deal with cases where one set of policies might conflict with other policies, or in which certain laws supersede other laws. In the legal world, there are “levels of authority” by which some laws trump other laws. Policy management needs to be able to represent this to enable correct access control. The laws and requirements of one jurisdiction often are in direct conflict with laws and requirements of other jurisdictions, so there needs to be a way to determine which laws (and therefore which policies and rules) should apply to access control decisions.
- Any language or construct that is used to express policies and rules associated with laws needs to be able to fully implement them, rather than the current situation in which digital policies and rules only partially implement laws.
- The expression of policy in digital form must be capable of enabling compliance with the Federal Information Security Management Act (FISMA), the Health Insurance Portability and Accountability Act (HIPAA), and Sarbanes Oxley (SOX) standards.
- The expression of policy in digital form must support reconciliation of policies in a federated environment. Federated organizations need to have the protections on their data enforced by member organizations that they share their data with. For example, a person might be disallowed from accessing the contents of a file by the owning organization, but might still be able to gain access to the file’s contents indirectly through another organization unless policies have been reconciled. As noted earlier, XACML may fall short in this regard. Other languages, such as Attempto, Semantic Web, and RuleML should be assessed for their suitability to address the reconciliation issue.
- Harmonization of the attributes used in shared or intersecting policies among organizations may not be easy to achieve, or it may not be feasible to harmonize all of those attributes. Automated support for harmonization and a method for dealing with attributes that cannot be harmonized are needed.
- Federated data base management systems are within the state of the practice. In theory, they could provide capability for managing and accessing attributes at the enterprise level. However, their greater capabilities can involve high cost and complexity of administration and management. A product study would benefit organizations attempting to implement enterprise-level access control systems; the study should determine whether vendor offerings, such as federated directories, could directly meet the need.

- Research is needed to address the issues associated with multiple values of attributes having the same name. Although a complete solution may not be possible, improvements to the current situation should be possible. Determining whether same-named attributes have the same semantics, translating and mapping multiple instantiations of the same attribute, and establishing standards for well-known attributes are some possibilities.

Standards

- There are two general concerns regarding standards. Standards for some of the desired operational capability for enterprise privilege management, especially concerning cross-domain or federated systems, simply do not exist. And adoption and use of standards needs to be encouraged in both federal programs and among vendors.
- Well-defined, relatively simple, and easily implementable standards would go a long way in helping vendors create products that could create and manage digital policies, attributes, and access control methods. Industry, academia, and government need to work cooperatively to develop standards and foster their adoption. From a vendor's point of view, a standard tends to limit opportunity for competitive advantage. Vendors would more likely accept and support standards defined in such a way that they could achieve true compliance and still be able to implement proprietary features for competitive advantage without breaking interoperability.
- Standards are desirable that define the following interfaces:
 - Interface between Privilege Management system and Privilege Management Administrator's Workstation
 - Interface between Privilege Management system and attribute and policy stores
 - Interface between Privilege Management system and Policy Decision Point
- A standardized access control framework that could be used to express and enforce any policy (or at least most policies), besides providing the usual benefits of standards, such as interoperability, would also inform privilege management in at least two ways. First, it would provide the standards needed to express policy. Second, it presumably could define the interface between the policy decision point and the attribute and policy store (see Figure 6) as well as an interface between the policy decision point and the privilege management system (see Figure 7).
- A standardized metric that could be used to evaluate the advantages and disadvantages of various access control methods for a given context would be helpful to enterprise administrators. Use of the metric could inform what would be needed in terms of privilege management to support a chosen access control method. Table 1 in the section on access control methods might be a start in this direction. Once a metric is established, a test bed could be used to demonstrate compliance.
- A standard, possibly an extension of XACML, is needed for the creation, maintenance, and dissemination of policy, for attribute collection, maintenance, and retrieval, and for policy enforcement. This standard must be able to express complex digital policies to capture the intent of complex real-world laws and regulations with accuracy and completeness.

- How can digital policies be implemented among partner organizations when the policies/rules under which they are legally required to operate are different? For example, how can policies among coalition countries with different data sharing requirements be represented? This appears to require the creation of better standards to define what types of data can be shared and with whom, as well as the tagging of data elements so that policy machines can make determinations about whether the data elements meet a particular policy and therefore whether they can be shared.
- A standard is needed for federation agreements that specify at least an attribute practice statement and a trustworthiness metric for attributes.
- XACML has an emerging standard for the interface between a Policy Decision Point and a Policy Enforcement Point. However, a tangible codification of architecture in this area would be helpful. Standards specifying protocols and formats for the interfaces discussed earlier (see Figure 6 and Figure 7 above) would be beneficial. Also needed is implementation guidance for secure communication (for example, TLS/SSL) at those interfaces, as well as between a Policy Decision Point and attribute/policy stores.
- Looking ahead to a method like risk-adaptable access control¹⁴ (RAAdAC), the community would benefit from a standardized way to identify and describe environment attributes.
- The community also needs a standardized way to identify and describe resource attributes to cover properties such as resource type and access methods.
- LDAP, XPATH, and SQL are available and suitable for defining/accessing attributes at a low level but no standard for life cycle management and governance of attributes exists for use by the U.S. government.

Methods

- The community could profitably develop an attribute management method by understanding and articulating the required functions. In this regard, SNMPv3 might be a useful starting point for a model architecture.
- There has been much socialization of the term ABAC, but there are still many variations in its description. The community would benefit from a clear definition of ABAC that can be formalized into a standard. The proposed definition presented earlier in this report might be an appropriate starting point for discussion and refinement.

State of the Practice

- Technology offerings cover some of what is needed for successful enterprise-level access management but are unable to create, consume, and enforce many policies that are important to government and industry. In addition, they are unable to dynamically adapt digital policy to a changing risk environment.

¹⁴ See Annex B: A Survey of Access Control Methods for a discussion of this and other methods.

Automation

- The state of automation appears to be lagging business needs in many areas. For example, while some organizations may have a clear idea of the policies that would serve their needs, they generally do not have automated tools to create and manage those policies. This is partly a human-computer interface problem but also requires appropriate supporting infrastructure.

Miscellaneous Topics

- Resource discovery and search
- Content protection
- Controlling programs in execution

Conclusion

Summary

This report presents an enterprise-level context for thinking about privilege management, starting with a real-time framework for access control, which drives privilege management functionality. Considering things at the enterprise level ensures that all elements of privilege management are included so that all organizations' needs can be met, from the smallest to the largest organizations.

Privilege management is conceptually split into two parts—attribute management and policy management—since attributes and policies might have differing governance and they deal with different kinds of entities or structures, to which different standards apply. Attributes are characteristics of entities, while policies specify how to use attributes in making an access control decision.

One basic view of privilege management's interfaces to other components of access management places attributes and policies between privilege management and access control. In this view, access control functionality includes direct retrieval of attributes and policies needed to make an access control decision. A different basic view is that access control makes requests of privilege management for the attributes and policies it needs. In this view, privilege management functionality retrieves attributes and policies and delivers them to access control, per its requests. Implementation of enterprise-level access management, and privilege management specifically, would clearly benefit from appropriate standards for the interfaces between major subsystems. Standardized interfaces would facilitate integration of products from various vendors.

Existing definitions serve the purpose of effective discourse, but there is currently no generally accepted definition of privilege management. This report uses a definition of privilege management¹⁵ that aligns with the draft definition of privilege management made by the Identity, Credential, and Access Management (ICAM) Subcommittee of the Federal Information Security & Identity Management Committee.

The eXtensible Access Control Markup Language (XACML), an OASIS standard for managing access control policy, provides some of what is needed to support enterprise-level privilege management. Also needed are capabilities for creation and maintenance of policy, policy enforcement, attribute collection, maintenance, and retrieval, and definition of resource attributes.

In the area of access control methods, this report concludes that a formal standardized definition of attribute-based access control (ABAC) is needed. Although there has been much socialization of the term ABAC, there are still many variations in its description. Toward the end of standardizing its meaning, a definition of ABAC was offered as a starting point.

There are three access control methods that are clearly distinguishable from one another: identity-based access control (IBAC), role-based access control (RBAC), and attribute-based access control (ABAC). Although policy-based access control (PBAC) does not need to be considered a separate access control method, the ideas associated with views of it make it clear

¹⁵ *Privilege management* is the definition and management of attributes and policies that are used to decide whether a user's request for access to some resource should be granted.

that full realization of ABAC's potential requires better attribute and policy management capabilities at the enterprise level. An organization should consider the many factors regarding requirements for its access control system that are briefly described in Table 1 above.

Policy capabilities for applications supporting both federal government and commercial enterprises are a key element for effective privilege management at the enterprise level. Policy developers should carefully consider what attributes are needed to implement the policies, gain an understanding of users and of the characteristics of the data involved, and consider whether there is a need to accommodate unanticipated users.

Recommendations

Based on the view of privilege management captured in this report, the following recommendations are made:

- **Privilege Management Workshop** Convene the second workshop on Privilege Management, a follow-up to the NIST-NSA Privilege (Access) Management Workshop. Defining and socializing resource and environment attributes, including metadata, is a recommended track for the workshop. Increased access to protected resources while denying access to adversaries makes the need for confidence in the correct operation of access management and related capabilities a high priority; thus, an important theme for the next workshop is assurance.
- **Policy Management** Convene a workshop on digital policy management to bring together the practitioners in all aspects to capture the current state of the practice and emerging research and to develop a research agenda. An important consideration for such a workshop is the complete digital policy life cycle. The approach used for the NIST-NSA Privilege (Access) Management Workshop could be used as a model for a digital policy management workshop.
- **Attributes** With regard to developing an Attribute Management Method, expand on the Authoritative Attribute Source work (Annex C: Authoritative Attribute Source and Attribute Service Guidelines).

A Way Forward

From the findings reported, it is also possible to outline a way forward for program managers of federal, Department of Defense (DoD), and other enterprise systems, expressed in terms of desired end states for their enterprise systems.

- The system is able to accommodate a growing population of users with increased diversity and can accommodate unanticipated users.
- The system supports sharing of information among a diverse set of users with a diverse set of access restrictions; to support this, the system implements fine-grained access control.
- The system has provisions to accommodate federated mechanisms to ensure seamless operation across traditional boundaries.
- The system includes risk-adaptive capabilities to respond to situational awareness of conditions—increased threat, political drivers, time-sensitive needs, and so on.

- Increased access to protected resources while denying access to adversaries makes the need for confidence in their correct operation (assurance) a high priority.

Given a desired end state, the way forward is to work toward achieving that end state. Additional topics about moving toward a desired end state are identified in Annex D: Advanced Capabilities for Privilege Management.

Bibliography

- [AASC] Authorization and Attribute Services Committee, December 8, 2009, *Authorization & Attributes Glossary*, Version 17, Multi-corporation committee.
- [AFIT-08] Air Force Institute of Technology, March 2008, *Composable Distributed Access Control and Integrity Policies for Query-Based Wireless Sensor Networks*, ADA478636, Air Force Institute of Technology, Wright-Patterson Air Force Base.
- [CNSSI-4009] Committee on National Security Systems, National Manager, June 2006, *National Information Assurance (IA) Glossary*, CNSS Instruction No. 4009, CNSS Secretariat (I01C) - National Security Agency - 9800 Savage Road - STE 6716 - Ft Meade MD 20755-6716.
- [DEYO-08] DeYoung, H., 2008, *Logic for Reasoning About Time-Dependent Access Control Policies*, DoD-XA.
- [FERN-05] Fernandez, R., 2005, *Enterprise Dynamic Access Control (EDAC)*, ADA461545, DoD-XA.
- [FERR-05] Ferraiolo, D. F., S. Gavrila, V. Hu, and D. R. Kuhn, 2005, *Composing and combining policies under the policy machine*, pages 11-22, Proceedings of the tenth ACM symposium on access control models and technologies, ISBN: 1-59593-045-0, Association for Computing Machinery, New York, NY, USA.
- [FERR-07] Ferraiolo, D.F., D.R. Kuhn, and R. Chandramouli, 2007, *Role-Based Access Control*, Information Security and Privacy Series, Second Edition, ISBN 13: 9789-1-59693-113-8, ARTECH HOUSE, INC., 68 Canton Street, Norwood, Massachusetts 02062.
- [FERR-07b] Ferraiolo, D. F., V. Atluri, and S. Gavrila, 2007, *The Policy Machine: A Standards-Driven Enterprise-Wide Access Control Enforcement Mechanism*, pages 87-92, Defense Standardization Program Journal, Defense Standardization Program Office, 8725 John J. Klingman Road, Stop 6233, Fort Belvoir, VA 22060-6221.
- [FICAM-09] Identity, Credential, and Access Management (ICAM) Subcommittee, June 15, 2009, *Federal Identity, Credential, and Access Management (FICAM) Roadmap and Implementation Guidance*, DRAFT Version 0.2, Federal Information Security & Identity Management Committee (ISIMC).
- [GREG-07] Gregory, M., March 14, 2007, *A Mechanism for Risk Adaptive Access Control (RAdAC)*, Presentation in PDF format, National Information Assurance Research Laboratory (NIARL).
- [HENG-05] Hengartner, U., 2005, *Access Control to Information in Pervasive Computing Environments*, ADA457117, DoD-XA.
- [HU-01] Hu, V., D. A. Frincke, and D. F. Ferraiolo, 2001, *The Policy Machine for Security Policy Management*, pages 494-503, Computational Science -- ICCS 2001, ISBN 3-540-42233-1, Springer-Verlag Berlin, Heidelberg, New York.
- [ILLI-07] Illinois University, June 1, 2007, *PolicyMorph: Interactive Policy Transformations for a Logical Attribute-Based Access Control Framework*, ADA482461, Illinois University at Urbana-Champaign, Department of Computer Science.

- [McGR-UN] McGraw, R. W., undated, *Risk-Adaptable Access Control (RAdAC)*, Presentation in PDF format, publisher not indicated; available 12-9-2009 at http://csrc.nist.gov/news_events/privilege-management-workshop/radac-Paper0001.pdf.
- [McGR-09] McGraw, R. W., September 2009, *Risk Adaptable Access Control (RAdAC)*, Presentation in PDF format, Information Assurance Architecture and Systems Security Engineering Group, National Security Agency, Ft. Meade, Maryland.
- [NIST-09] National Institute of Standards and Technology, November 2, 2009, *Proceedings of the First Workshop on Privilege (Access) Management*, National Institute of Standards and Technology, Gaithersburg, Maryland.
- [NISTIR-7298] Kissel, R. (ed.), April 25, 2006, *Glossary of Key Information Security Terms*, NIST IR 7298, National Institute of Standards and Technology, Gaithersburg, Maryland.
- [NISTIR-7316] Hu, V.C., D.F. Ferraiolo, and D.R. Kuhn, September 2006, *Assessment of Access Control Systems*, Interagency Report 7316, National Institute of Standards and Technology, Gaithersburg, Maryland.
- [OASD-10] The Office of the Assistant Secretary of Defense for Networks and Information Integration / DoD Chief Information Officer, January 6, 2010, *Department of Defense Privilege Management Roadmap*, approved for public dissemination, Office of the Assistant Secretary of Defense, United States. Although approved for public dissemination, at the time of this writing it was not available on a public Web site. It is available as a companion publication to this document. This document in PDF format provides the roadmap as an attachment.
- [RFC-4949] Shirey, R., August 2007, *Internet Security Glossary, Version 2*, RFC-4949, Informational, The IETF Trust.

Annex A: Authorization and Attributes Glossary

This annex contains the draft Authorization and Attributes Glossary prepared by the Authorization and Attribute Services Committee [AASC]. Although not officially adopted by NIST as a reference glossary, it is recognized as relevant background material, providing a departure point for moving forward. Recommended additional glossaries, all of which are accessible to the public, are the glossary by the Committee on National Security Systems [CNSSI-4009], the Internet Engineering Task Force glossary [RFC 4949] (was RFC-2828) and the NIST glossary [NISTIR-7298].

The glossary document is an attachment to this document.

The Authorization and Attributes Glossary references the following sources:

- AASC – Authorization and Attribute Services Committee
- CNSSI-4009 – CNSSI 4009, The National Information Assurance Glossary http://www.cnss.gov/Assets/pdf/cnssi_4009.pdf
- DoD-DS – DoD Net-Centric Data Strategy <http://cio-nii.defense.gov/docs/Net-Centric-Data-Strategy-2003-05-092.pdf>
- DoD-EDS – DoD Enterprise Directory Services Capability: Contact Attributes Specification, Version 2.0, July 14, 2009 <http://cio-nii.defense.gov/docs/Signed%20DoD%20Enterprises%20Services%20Specifications.pdf>
- DoD-JSAP – Department of Defense (DoD) Joint Staff Action Processing review of AATT deliverables, not publicly available
- ESM – Enterprise Security Management terms extracted from the GIG IA Architecture (restricted access document), and map back to the DoD Joint Capabilities Documents.
- FEA – The Federal Enterprise Architecture - Data Reference Model (FEA-DRM) Version 2.0, dated November 17, 2005 - http://www.whitehouse.gov/omb/assets/egov_docs/DRM_2_0_Final.pdf
- ICD 501 – Intelligence Community Directive Number 501: Discovery and Dissemination or Retrieval of Information Within the Intelligence Community, January 21, 2009 http://www.dni.gov/electronic_reading_room/ICD_501.pdf
- ICIA-Glossary – Intelligence Community Information Assurance Glossary Version 1.0, not publicly available
- IEEE-PP - IEEE Standard for a Protection Profile in Operational Environment A, IEEE Std 2600.1™-2009 <http://standards.ieee.org/getieee/2600/download/2600.1-2009.pdf>
- Intelink-U – Intellipedia, U.S. Government Unclassified Wiki, <https://www.intelink.gov/wiki/>, login account required
- OASIS XACML – Core and Hierarchical Role Based Access Control (RBAC) Profile of XACML v2.0, OASIS Standard, February 1, 2005 http://docs.oasis-open.org/xacml/2.0/access_control-xacml-2.0-rbac-profile1-spec-os.pdf
- PP – Protection Profile http://www.niap-ccevs.org/pp/pp_authsrv_br_v1.1/
- RFC 4949 – Internet Engineering Task Force (IETF) Request for Comments (RFC) 4949, Internet Security Glossary, Version 2, August 2007 - <http://www.ietf.org/rfc/rfc4949.txt>

- SAML – Glossary for OASIS Security Assertion Markup Language (SAML) Glossary Version 2.0, March 15, 2005: <http://docs.oasis-open.org/security/saml/v2.0/saml-glossary-2.0-os.pdf>
- WEB – Webster’s Online Dictionary - <http://www.merriam-webster.com/dictionary>
- X.812 - ITU-T Recommendation X.812 – Security Frameworks for Open Systems: Access Control Framework <http://www.itu.int/rec/T-REC-X.812/>
- 5 U.S.C. § 552a – The Privacy Act of 1974: <http://www.justice.gov/opcl/privacyact1974.htm>

Annex B: A Survey of Access Control Methods

The attached survey document was developed as an input to the Privilege (Access) Management Workshop as an aid to stimulate discussion. The survey document is an attachment to this document.

Annex C: Authoritative Attribute Source and Attribute Service Guidelines

This topic (authoritative attribute source and attribute service guidelines) was briefed at the workshop and documented as a draft special publication. While not representing an official position at this point, the draft captures technologies and reference data that were addressed during the course of the workshop. As such, it provides a good starting point for further investigation.

The guidelines document is an attachment to this document.

Annex D: Advanced Capabilities for Privilege Management

This annex lists advanced capabilities to consider in relation to enterprise privilege management. Analysis of each capability in the context of a particular enterprise is necessary to determine the viability, practicality, and impact of applying resources to gain the capability.

- Design and implementation guidance
- Risk-adaptable mechanisms and policies
- Measures of confidence guidance
- Brokered trust solutions
- Decision mechanisms for fine-grained policy evaluation
- Support for complex data types
- Support for multiple authentication methods
- Privilege federation architecture
- Federated data access capabilities suitable for access to attributes
- Enterprise authorization services
- Variable time policies
- Ability to accommodate near real-time policy and attribute changes
- Policy validation capabilities
- Ability to accommodate dynamic communities of interest
- Ability to accommodate disconnected/disadvantaged users
- Ability to assert the role in which a user is operating
- Controls to protect “protected” attributes (for example, cover identification)

Annex E: The Policy Machine

This annex provides a synopsis of The Policy Machine. Discussion of details and application can be found in [HU-01], [FERR-05], and [FERR-07b].

The Policy Machine (PM) is an access control framework that significantly extends the span of enforceable access control policies. It retains many of the advantages and eliminates many of the disadvantages of existing approaches. Although it exhibits features similar to those of other access control frameworks, the PM is not an extension or adaptation of any existing access control model or mechanism. Instead, the PM is meant to be a redefinition of access control in terms of a standardized and generic set of relations and functions that are reusable in the expression and enforcement of policies. Its objective is to provide a unifying framework to support a wide range of attribute-based policies or policy combinations through a single mechanism that requires changes only in its data configuration.

The PM supports two types of applications. The first type comprises those applications that provide services that are independent to access control. These applications include, for example, text editors, spreadsheets, and drawing packages. The second type comprises those applications that provide services effectively through access control. For example, email applications provide services through the discretionary distribution of messages and attachments, and workflow management applications provide services through the distribution of capabilities (read and write operations to an object [work item] to a prescribed sequence of users.

To demonstrate the PM's viability, NIST has developed a reference implementation. It can be shown, through this reference implementation, how to configure PM for the expression and enforcement of a diverse set of policies. The policies include instances, combinations, and hybrids of DAC, MAC, RBAC, Chinese wall, ORCON, history-based separation of duty, and so on. Not only can enforcement of these policies be shown on files but also enforcement can be demonstrated within and across a rich user environment that includes the Open Office suite of applications, email, workflow management, and records and forms management.

Annex F: Security Framework for Privilege Management

This annex presents comments submitted by Anil Ramcharan¹⁶ during the public review period. Although the comments do not exactly follow the guidelines posted with the NISTIR, they collectively form a cohesive view of privilege management that is interesting and informative. This view differs somewhat from the view taken in this NISTIR, as noted in the accompanying commentary by the editor of this NISTIR. This alternate view is well-considered and provides an additional springboard for further exploration of the area covered by this report.

Each of Anil Ramcharan's comments begins with a page and figure reference, together with a category of comment. Following the comment's text and figure (if any) in quotation marks, are comments by the editor, bounded by square brackets. Also included in square brackets, for the convenience of the reader, are the figures in the NISTIR, if any, to which Anil Ramcharan's comment refers.

Page 3, Figure 2 — Suggested Modification

“The suggested modification is to change the terms Policy Enforcement Point and Policy Decision Point to Access Control Decision and Access Control Enforcement. The rationale behind dropping the “point” at the end of the term is to remain consistent with the rest of the terminology on the diagram (i.e., Attribute and Policy Information as opposed to Information Point). The rationale behind changing the Policy to Access Control is to add specificity to the kind of policy decision being made. Policy has a definition that can extend well beyond access controls such as governance or a course of action. Policies can be evaluated and enforced that are unrelated to the access control scenario being presented. Since the decision being calculated and enforced relates directly to access control for the access request provided, the terms should be relabeled as access control enforcement and access control decision.”

[Editor: The terms “Policy Enforcement Point” and “Policy Decision Point” are the terms adopted by the workshop principals (organizers and presenters) and participants (attendees), based on definitions found in authoritative glossaries and some public literature. However, the terms Access Control Enforcement and Access Control Decision are eminently sensible and might be the preferred terms in some contexts. As Anil Ramcharan has pointed out, “Policy has a definition that can extend well beyond access controls . . .” Recognizing this, we have attempted to bind the scope of “policy” as generally used in this document to the area of access management.]

¹⁶ Senior Consultant , Deloitte Consulting LLP, 1676 International Dr, Mclean, VA, 22102

[Editor: The figure and caption in the body of the NISTIR that Anil Ramcharan has referred to are as follows (reproduced here for the convenience of the reader)—please note that this figure is from an earlier version of this NISTIR and differs from the corresponding figure in this final report:

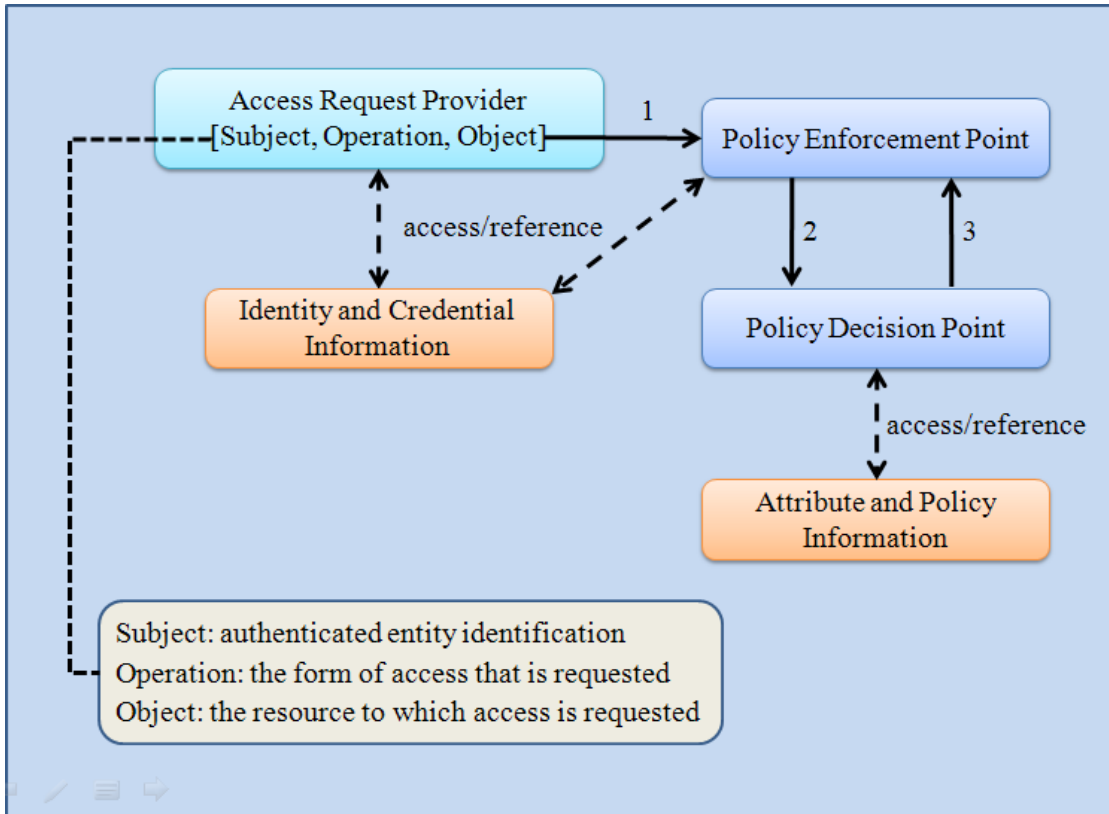
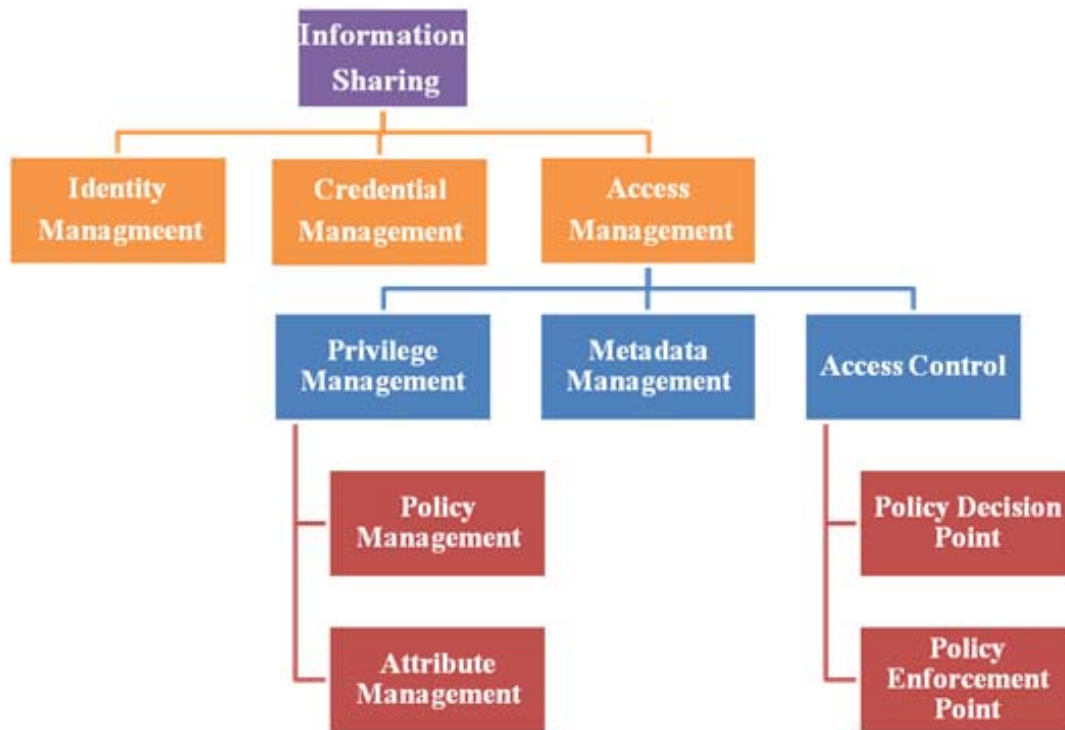


Figure 2. High-Level View of Real-Time Access Control

—end of Editor comment.]

Page 4, Figure 3 — Suggested Modification/Correction

“To answer the question, ‘Where does privilege management fit in this view?’ the suggested modification/correction is to include a taxonomy in place of Figure 3 to show the relationships between identity, credential, access, privilege, attribute, policy, and metadata management and access control.



Privilege management should be used to refer to the management of entity permissions that correlate to attributes and the applicable policies. (More info on this in the next comment.) Metadata management would include the characteristics of resources (i.e., data, files, applications) to which access is controlled. Access control would be the calculations and execution of permissions defined in privilege management.”

[Editor: The proffered figure is a good representation of logical relationships, while Figure 3 in this document shows functional relationships. As an aside, the labels “Policy Decision Point” and “Policy Enforcement Point” should probably be “Access Control Decision” and “Access Control Enforcement” to be consistent with the earlier comment. In any case, the proffered figure is an excellent adjunct to those in the body of this document. The statements following “Privilege should be used to refer to the management of entity permissions . . .” represent a departure from the view of privilege management taken in this document. As Anil Ramcharan has suggested in his comment, there is more information on this in the next comment, where we will address the issue.

[Editor: The figure and caption in the body of the NISTIR that Anil Ramcharan has referred to are as follows (reproduced here for the convenience of the reader):

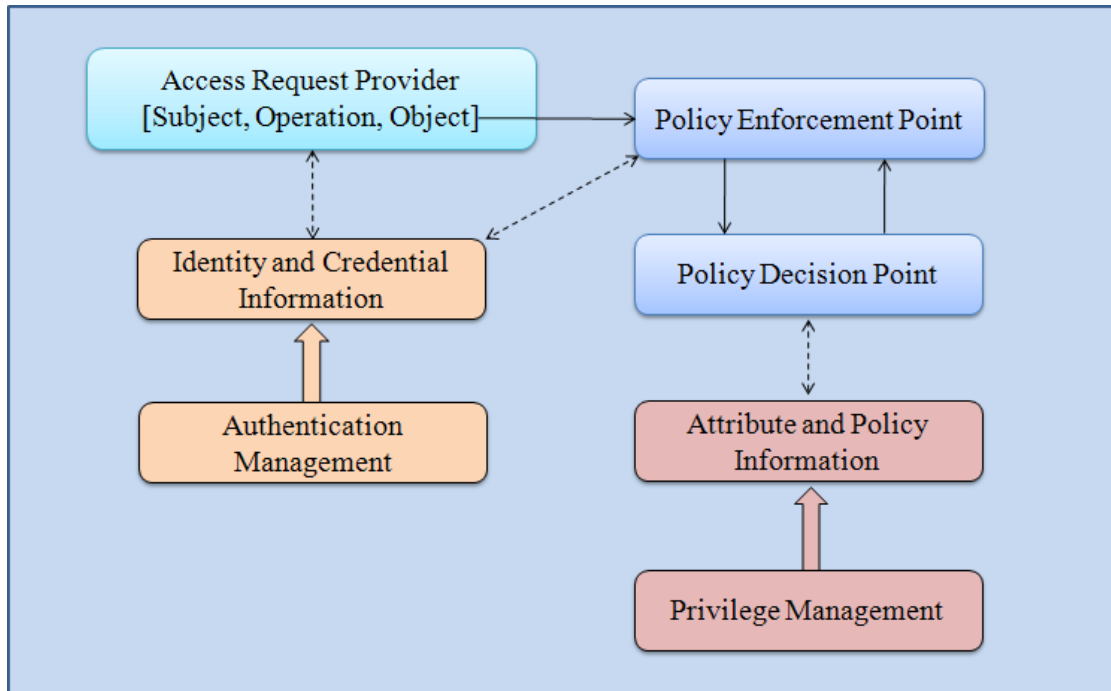


Figure 3. Authentication Management and Privilege Management

—end of Editor comment.]

Page 5, Figure 4 — Suggested Modification

“In the privilege management diagram, I would recommend changing the third level below ‘attribute management’ and ‘policy management’ to cover their supporting processes as opposed to listing entity, resource, and environment. The interpretation of privilege management could be the maintenance of permissions associated with entities or subjects requesting access to a resource. It would be incongruent to say that the management of a document classification level is the same of managing that document’s privilege. To manage an entity’s privileges is to manage what an entity is able to do, which includes the entity characteristics and the access control policy. In an ABAC or PBAC model, the resources may have metadata associated with them such as Community of Interest (COI) restrictions that are an input to the policy using the process of evaluating an access control decision, but this is not the same as saying that the resource metadata is a privilege. Nonetheless, both privileges and resource metadata (as well as environmental attributes) are part of access management in the respect that they play a part in defining security and transaction context. In this light, resource attributes and environment attributes would be separate categories on the same level as privilege management and access control (see the diagram above).”

[Editor: We have taken the view in this document that privilege management has a broader scope than suggested by Anil Ramcharan. As used in this document, privilege management creates attributes and assigns attribute values to entities. Thus, in our view, privilege management does deal with a document’s classification as well as that document’s “privilege,”

as expressed in a digital policy. Further, metadata is encompassed by the term “attributes.” Thus, a COI restriction associated with a resource is expressed as an attribute; the COI restriction is then taken into account by a digital policy, which is also created and managed by privilege management.]

[Editor: The figure and caption in the body of the NISTIR that Anil Ramcharan has referred to are as follows (reproduced here for the convenience of the reader):

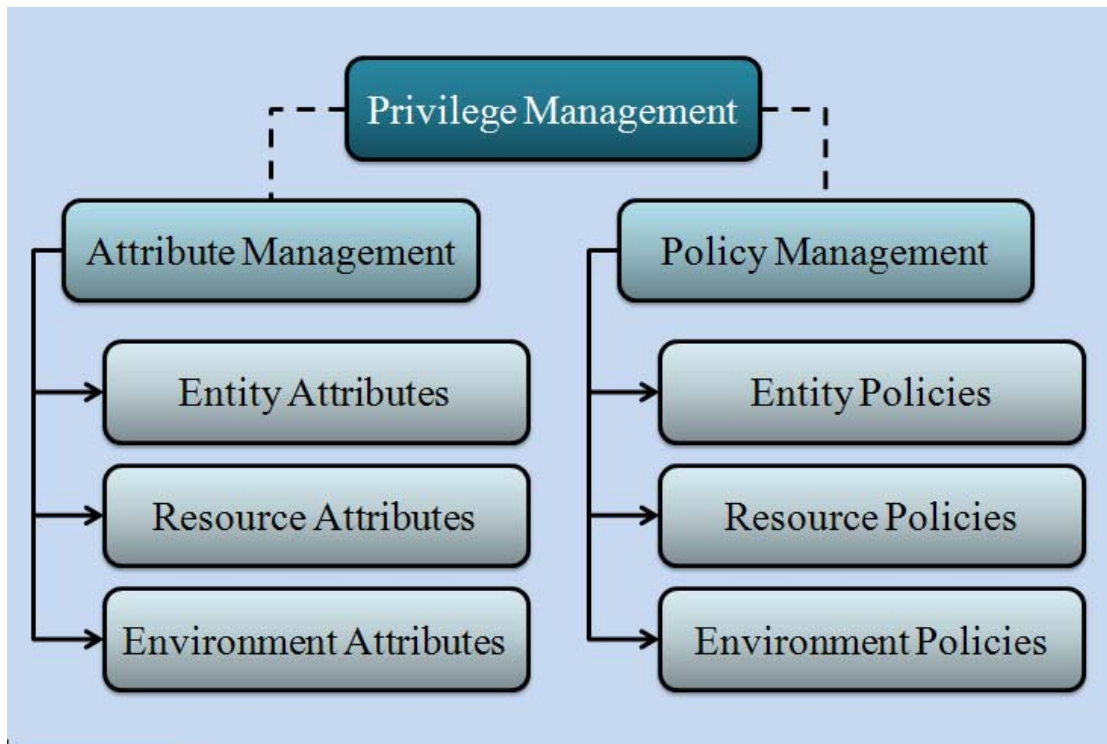


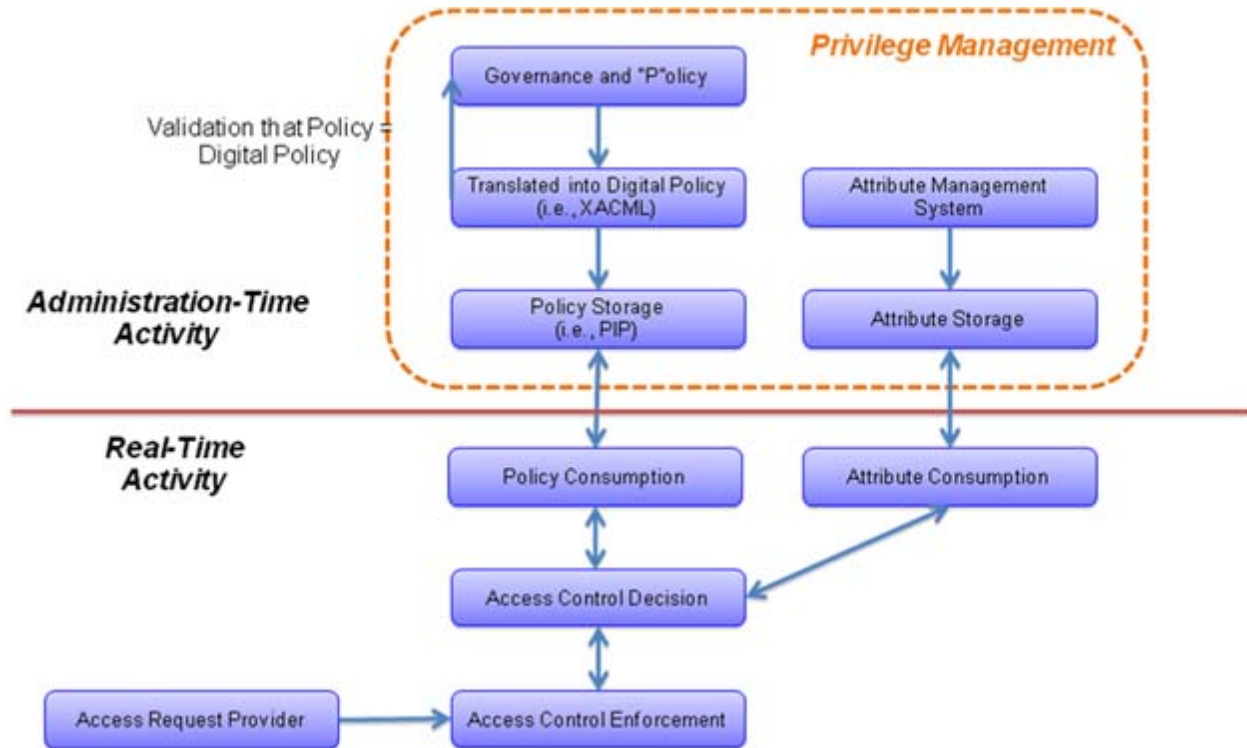
Figure 4. Information Managed by Privilege Management

—end of Editor comment.]

Pages 7 and 8, Figures 5 and 6 — Suggested Modification/Correction

“The suggestion is to replace the diagram in Figure 5 with the diagram below. The diagram shown below provides greater detail relating to the administration-time policy and attribute management activities. The existing diagram 5 is nondescript in terms of what are the activities related to privilege management. The suggested diagram illustrated the development of a digital access control policy starting with a governance policy. This “P”olicy is then translated into a digital policy in a format such as XACML. There is a step to validate that the digital representation of the governance policy as an access control policy is valid, and upon successful validation, the policy is entered into policy storage where it is accessible for evaluating access control decisions. With respect to attributes, there is an attribute management life cycle that includes a vetting and binding process. The resulting attributes are then housed in an attribute store where it is accessible by the access control decision mechanism in order to support the evaluation of a digital policy. Also depicted in the diagram below are policy and attribute consumption. These interfaces are intended to provide for the discovery and retrieval of policy

and attributes. This allows for the access control decision mechanism to operate without knowledge of the policy or attribute stores. This componentization protects the access control decision mechanism from being impacted to changes made by the privilege management functions.”



[Editor: The proffered diagram reflects the architecture displayed in Figure 5 of this document and provides added detail—the functional boxes “Policy Consumption,” “Attribute Consumption”, and “Translated into Digital Policy.” It takes Figure 5 a step closer to a possible implementation. In addition, it expands the scope of privilege management to include the box labeled “Governance and ‘P’olicy. This addition seems reasonable and leads to a number of new questions that one could address—for example, what kind of automated assistance makes sense and how would it interface to “Translated into Digital Policy.” Also, the overall architecture displayed in the proffered diagram leads to the question of how the details and increased scope can be dealt with in the alternate architecture described in Figure 6 of this document. Thus, there are many considerations here that could inform an approach for a follow-on workshop on privilege management.]

[Editor: The figure and caption in question are as follows:

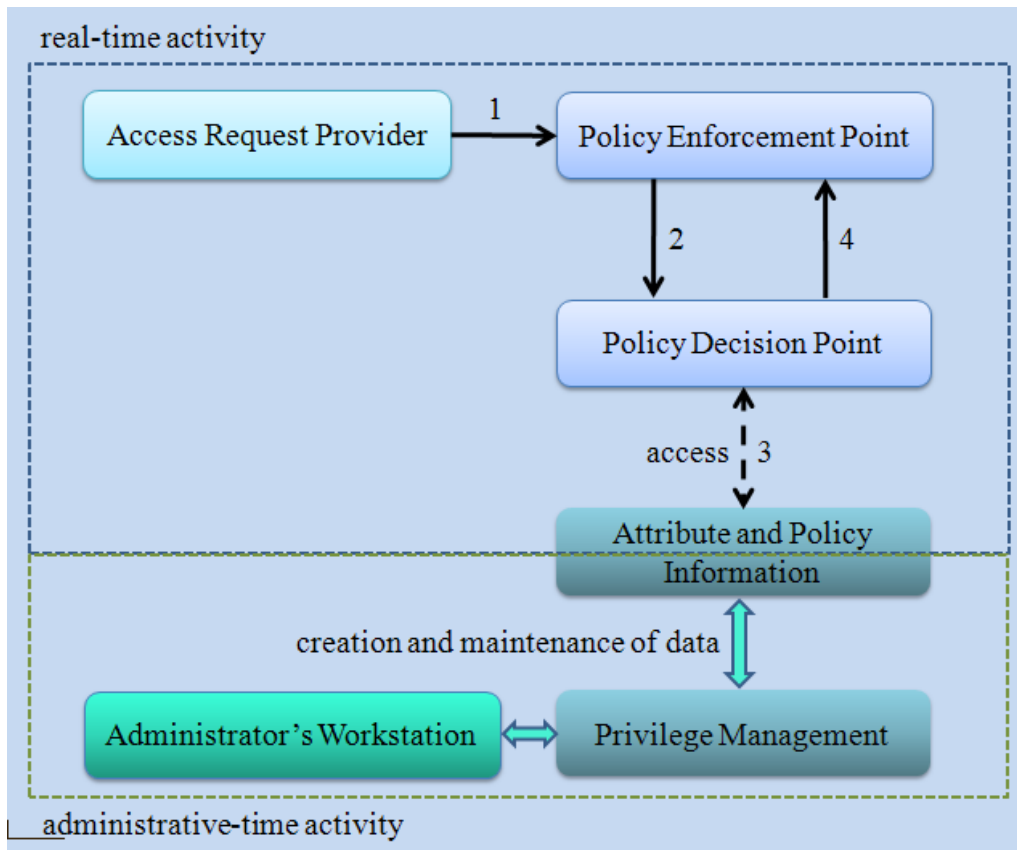


Figure 5. Interfaces of Privilege Management – View One

—end of Editor comment.]