DRAFT NISTIR 7665

Proceedings of the Privilege Management Workshop, September 1-3, 2009

Editors:
Sheldon A. Durrant
The MITRE Corporation

Tanya Brewer
Annie Sokol
NIST

January 2010

NIST National Institute of Standards and Technology • U.S. Department of Commerce

DRAFT NISTIR 7665

Proceedings of the Privilege Management Workshop, September 1-3, 2009

Editors:
Sheldon A Durrant
The MITRE Corporation

Tanya Brewer
Annie Sokol
NIST

January 2010

U. S. Department of Commerce
Gary Locke, Secretary

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

(This page intentionally left blank)

Table of Contents

Introduction

The first Privilege Management Workshop, co-sponsored by the National Institute of Standards and Technology (NIST) and the National Security Agency (NSA), was held at NIST's main campus in Gaithersburg, Maryland, September 1–3, 2009. The workshop was attended by approximately 120 people representing executive branch federal agencies, the private sector, and academia.

Like so many other areas of information assurance-related technical endeavor, privilege management is large and complex, often the source of heated debate and opinion, and fraught with widely understood, yet ill-defined terminology and concepts. Thus, a primary goal of this first workshop was to bring together a wide spectrum of individuals representing differing viewpoints, use cases, and organizational needs with the intent of reaching a common understanding of several facets of this important area. This includes reaching consensus on the definition of privilege management and other terminology; understanding and analyzing the strengths and weaknesses of current and proposed access control models; ascertaining the current state of the practice and future research directions in privilege management; and understanding and articulating the managerial, legal, and policy requirements associated with privilege management. To facilitate these objectives, the workshop was organized into four tracks:

- Standards, Definitions, and Terms
- Models and Frameworks
- State of the Technology and Research Agenda
- Policies and Requirements

This document is a synopsis of the major proceedings of the plenary and tracks.

One major point of consensus reached at the workshop was that the term "privilege management" describes a set of processes for managing the data, attributes and policies in particular that determine a user's access rights to a system. This definition aligns with the Federal Identity, Credential, and Access Management (FICAM) definition of privilege management, and firmly establishes privilege management as one component of the wider access management framework. A more detailed technical position on privilege management can be found in an upcoming NIST Interagency Report (NISTIR 7657) on privilege management.

Plenary Proceedings

The Plenary, a precursor to the track discussions, set the overall tone for what issues and topics attendees should address during the course of the workshop. The plenary session began with an outline of a potential framework and context within which to view the privilege management space. Subsequent presenters stressed the importance—and difficulty—of representing, creating, and enforcing policies, and used healthcare and the Health Insurance Portability and Accountability Act (HIPAA) to illustrate the shortfalls between legal requirements and what could actually be enforced with current privilege

management capabilities. Presenters briefed attendees on eXtensible Access Control Markup Language (XACML) and the differences between it and an implementation of a Policy Machine. Presenters also introduced the Risk-Adaptable Access Control (RAdAC) concept, and discussed what constitutes an authoritative attribute source, which provided the seed for further discussion in the various tracks.

During the plenary proceedings, attendees were invited to comment on the presenters' briefings and to pose questions. Details of these discussions can be found in Appendix A.

Track 1: Standards, Definitions, and Terms

Participants in Track 1 focused their efforts on one of the most important goals of the workshop—defining many of the key concepts associated with privilege management. After much debate, track participants agreed that "privilege management" is a component of the larger "access management" space. This is consistent with the FICAM framework supported by the Federal Chief Information Officers (CIO) Council and adopted by the Department of Defense (DoD).

Track 2: Models and Frameworks

A variety of models and frameworks exist for controlling access to enterprise resources. Track 2 participants sought to understand them, put them into context, and come to consensus on their advantages and limitations. The models discussed ranged from basic Access Control Lists (ACLs) and Role-Based Access Control (RBAC) to emerging models such as Attribute-Based Access Control (ABAC) to over-the-horizon models that include Policy-Based Access Control (PBAC) and Risk-Adaptable Access Control (RAdAC), a concept for integrating risk metrics, complex policies, machine learning, and a variety of attributes into making automated, dynamic access control decisions. One key conclusion reached in Track 2 discussions is that no access control model is better or worse than any other; they each have their place in the enterprise and should be adopted according to their suitability for a particular set of circumstances. Participants also concluded that Risk-Adaptable Access Control held promise, but had limitations that would preclude its use any time in the near future.

Track 3: Technology and Research Agenda

Track 3 attendees discussed the current state of the practice for privilege management and access control; analyzed gaps in current privilege management capabilities; discussed on going research in government, academia, and commercial spaces; and made suggestions on future research areas that need to be addressed. Participants agreed, with very few exceptions, that current privilege management implementations were limited to RBAC and that Attribute-Based Access Control was necessary, but needed to be more clearly defined and thoroughly researched. This research should include both stand-alone and federated ABAC models.

Track 4: Policies and Requirements

Track 4 tasked participants to discuss and analyze the relationship between privilege management and policy. Participants discussed what they thought would be needed to enable privilege management to address organizational policy, legal, and compliance requirements. Track attendees concluded that the trend in privilege management is toward using it to enforce organizational policy and compliance requirements, but that the greatly needed mechanisms to express digital policies were not yet mature. XACML was presented as a mechanism that has the potential to address some of these concerns, but it is not yet fully mature and has a variety of limitations.

Links and Resources

The following list includes the briefing slides from the workshop, as well as other materials and background information pertinent to the workshop proceedings. For further information, visit the workshop's Web site available at http://csrc.nist.gov/news_events/privilege-management-workshop/.

- Ferraiolo, David. "Policy Machine: Towards a Unifying Access Control Mechanism" http://csrc.nist.gov/news_events/privilege-management-workshop/presentations/Ferraiolo_Plenary.pdf. Privilege Management Workshop. September 1, 2009. Last Accessed 10/05/09.
- Ferraiolo, David, Serban Gavrila and Steve Quirolgico. "The Policy Machine: A Mechanism for the Specification and Enforcement of Arbitrary Attribute-Based Access Control Policies." http://csrc.nist.gov/groups/SNS/pm/PMInventionDescription.pdf. Last Accessed 10/05/09.
- LaPadula, Leonard J. "Privilege Management Framework." http://csrc.nist.gov/news_events/privilege-management-workshop/presentations/Len_LaPadula.pdf. Privilege Management Workshop. September 1, 2009. Last Accessed 10/05/09.
- Lafky, Deborah L. "Health Information Technology and Privilege Management: A Policy Agenda for Progress." http://csrc.nist.gov/news_events/privilege-management-workshop/presentations/Deborah_Lafky.pdf. Privilege Management Workshop. September 1, 2009. Last Accessed 10/05/09.
- McGraw, Robert W. "Risk-Adaptable Access Control (RAdAC)." http://csrc.nist.gov/news_events/privilege-management-workshop/radac-Paper0001.pdf . Last Accessed 10/05/09.
- Waterman, K. Krasnow. "Representing Policy for Enterprise Compliance." http://csrc.nist.gov/news_events/privilege-management-workshop/presentations/Krasnow_Waterman.pdf Privilege Management Workshop. September 1, 2009. Last Accessed 10/05/09.
- Westman, Roger. "What Constitutes an Authoritative Attribute Source?" http://csrc.nist.gov/news_events/privilege-management-workshop/presentations/Roger_Westman.pdf Privilege Management Workshop. September 2, 2009. Last Accessed 10/05/09.

Appendix A: Proceedings of the Plenary Question and Answer Session

Question/Comment	Answer/Response
Introductory Access Control Framework	
What are your thoughts on real-time risk management?	Stakeholders need to consider the risks and benefits of dynamically allowing access versus the risks and benefits of providing inappropriate access to a resource. These risks will be different depending on whether the organization in question is a profit-making company or a military organization. There is no single acceptable answer for everyone and every situation.
Where does federation fit into the framework that you presented and what about crossing authority boundaries?	I have no opinion on federation at this moment. The topic of federation and whether and how it fits into privilege management is something that the tracks will address during the workshop.
Policy Machine	
There is a need for standardized definitions of Policy Enforcement Point and Policy Decision Point. With regards to the Policy Machine, are there any papers that support the details of the work that the team can make available?	Yes. The policy machine team has papers, specifications, reference implementations, etc.
What is the difference between XACML (eXtensible Access Control Markup Language) and the Policy Machine?	XACML is based on a language. It does not deal with processes like cut-and-paste, etc. XACML is only a way to specify policy, not a way to enforce policy, which the policy machine is able to do.

FICAM

How extensible are Personal Identity Verification (PIV) credentials?	The goal is to minimize the number of PIVs and to have interoperable systems that include attribute management. Another goal is to be able to provision and validate identities across the board. PIVs need to be extensible to accomplish all of that.
Strong authentication is stable with Public Key Infrastructure (PKI). Does Authorization need commensurate levels?	There is a Defense Information Systems Agency's request for proposal (DISA RFP) coming for authorization, product evaluations, and procurement in FY11.
A federation component is going to be needed in order to access policies. Will there be rules in the framework on how to validate or verify access control policies like the Liberty Alliance uses with third-party assessors?	The Chief Information Officers (CIO) council has an ICAM subgroup that should focus on addressing those issues.
Has there been an uptake on Information Cards?	InfoCards have not taken off yet.

RAdAC

Is RAdAC a model or a set of requirements?	There are many ways to quantify risk and there are many ways to implement RAdAC. There are still pieces that need to be addressed. RAdAC is more of a concept than a model at this point.
RAdAC conceptually feels good but the devil is in the details. What has been done to socialize and quantify risk?	IBM has done some work in this area. For RAdAC to be successful the capabilities have to be put out in parallel and we need to develop the knowledge base so that appropriate risk models are put in place. Also, the infrastructure for Attribute-based Access Control (ABAC) needs to be in place.
Will there be a human in the loop in RAdAC deployments?	Initially, there will be a human in the loop, but the vision is to have a fully automated system.
Who is held accountable in a RAdAC environment if a machine makes all the access control decisions? Many people will be nervous about removing a human from the loop because of a perceived lack of accountability.	Who will be responsible is going to depend on who sets and manages the policy, whether that is the CIO, the CISO or the data owner, etc.
Could RAdAC allow only partial access to a resource in order to offset the risks of not having a human in the loop?	RAdAC does not have to be a yes or no answer. Provisions could be made to leverage “hints.”

Health IT

What are the privacy implications of the government mandate that all health records be digitized by 2014? Big Brother is not just the U.S government. Isn't there a risk that centralization in the private sector could pose privacy concerns?	From a U.S. government point of view, centralization is not an option because the public may perceive the government taking on the role of "big brother." The goal of making health records digitized and available online is to have a network of networks in which Health Insurance Portability and Accountability Act (HIPAA) will be the standard with which everyone must comply. This would take the pressure off the U.S. government and place responsibility onto the owners and stewards of the data.
60-80% of the medical practitioners are not from "technical" backgrounds. Does that shift concerns away from privacy and onto availability of the health data?	Yes. The more medical practitioners rely on online data, the more this reliance will drive the requirements for availability guarantees.