

Pluggable Authorization and Distributed Enforcement with `pam_xacml`

Andreas Klenk^{1,2}, Tobias Heide²,
Benoit Radier³, Mikael Salaun³, Georg Carle^{1,2}

¹TU München, Boltzmannstraße 3, 85748 München, Germany

²Wilhelm-Schickard-Institut, Sand 13, 72076 Tübingen, Germany

³France Télécom R&D, avenue Pierre Marzin 2, 22307 Lannion, France

Abstract. Access control is a critical functionality in distributed systems. Services and resources must be protected from unauthorized access. The prevalent practice is that service specific policies reside at the services and govern the access control. It is hard to keep distributed authorization policies consistent with the global security policy of an organization. A recent trend is to unify the different policies in one coherent authorization policy. XACML is a prominent XML standard for formulating authorization rules and for implementing different authorization models. Unifying authorization policies requires an integration of the authorization method with a large application base. The XACML standard does not provide a strategy for the integration of XACML with existing applications. We present `pam_xacml`, an authorization extension for the Pluggable Authentication Modules (PAM). We argue how existing applications can leverage XACML without modification and state the benefits of using our extended version of the authorization API for PAM. Our experimental results quantify the impact of security and connection establishment of using remote Policy Decision Points (PDP). Our approach provides a method for introducing XACML authorization into existing applications and is an important step towards unified authorization policies.

1 Introduction

Since the early days of the computers, the decision about permissible actions on a resource was critical. For a long time application specific configuration languages defined if and how the applications should act upon a request by specific subjects. These policies are sometimes hard to separate from the application configuration and are present at each instance of the application. The fact that the policies reside at the applications, distributed across the network, makes it difficult to gain a coherent view of the access control decisions and to change these policies.

In recent years general purpose policy description languages emerged. These languages provide a method for formulating application independent policies that can be evaluated by general purpose Policy Decision Points (PDP). Policy languages [1][4][12][13] enable unified and centralized security policies that allow for

a better insight into the authorization decisions and facilitate timely changes to the security policy. As an important benefit, unified policies can be tested for compliance with a high level security policy. Model checkers can automatically test policies for redundancy, constraints, and safety properties, and can perform a change-impact analysis [7]. The eXtensible Access Control Markup Language (XACML) [19] standard enjoys growing popularity for specifying authorization policies. However, the use of policy languages is limited if they lack integration with applications.

The Pluggable Authentication Module (PAM) [22] is well established for providing unified authentication to applications. As a consequence, these applications become independent of the underlying authentication method (e.g. Kerberos, LDAP, Radius, and Smart Cards). Unfortunately such a solution is missing for authorization. This paper introduces the novel `pam_xacml` module for authorization with XACML. `pam_xacml` can be used for most existing applications that support PAM without the need for code modifications. An extension to the PAM conversation function enables even richer authorization decisions for applications that implement the new interface, and allows to return obligations for permissible actions. Introducing unified policies comes at the cost of additional communication with the PDP. The different deployment options of XACML PDPs have a significant impact on the performance of authorization. Authorization is a function that may happen multiple times at different applications during a single service request. Our experiments provide performance figures on the effect of multi-point authorization in distributed systems. This paper presents the following contributions, thereby motivating wide-spread use of XACML in distributed heterogeneous environments:

1. A PAM module for authorization decisions with XACML, usable for existing applications without modifications.
2. An extension of the PAM conversation function to provide more input for the authorization decisions and to return obligations to the application.
3. A study on the communication cost of distributed authorization and advice on different deployment options.

The remainder of this paper is structured as follows. Sec. 2 reviews related work on authorization in distributed systems. Sec. 3 describes our approach for use by many PAM enabled applications. Sec. 4 presents results from our measurement study with the prototype implementation.

2 Related Work

Authorization languages were subject to research and standardization during the last years. The Authorisation Specification Language was introduced in [12], Ponder [4] is a declarative language for security and management policies, X-RBAC [1] uses XML to express Role Based Access Control (RBAC) policies. We focus on the eXtensible Access Control Markup Language (XACML) [19] OASIS standard for authorization policies. XACML can take authorization rules

in XML and is able to implement different authorization models, such as Discretionary Access Control, Mandatory Access Control (MAC) and Role Based Access Control (RBAC). Positive experiences of using of XACML for access control in heterogeneous distributed systems have been reported in [18].

The IETF laid down the nomenclature and architecture commonly found in modern authorization systems [23], but did not specify concrete interfaces or messages. The Generic AAA (GAAA) [9] is one implementation that follows the ideas of the IETF framework. Follow up work on GAAA is done by [5]. Gheorghiu et al. introduced the General Authorization and Access API (GA-API) [8]. Two IETF-Drafts have been written for this API, both of which are now expired. The Common Open Policy Service (COPS) [6] has a binary encoded policy interface that allows to contact a PDP, mainly for decisions on media level access control and QoS specifications. The Open Group published the Authorization (AZN) API standard in [21]. It describes a C interface for authorization based on the ISO-standard 10181-3. Applications can use the AZN-API for authorization requests. PERMIS is discussed in [3]. It uses a simplified version of the AZN-API introduced before. The authors in [11] propose an extension of the Java Authentication and Authorization Service for class-instance level access control with XACML, which is only available for Java programs.

The Security Assertion Markup Language (SAML) [2] defines an authorization decision query protocol. The CARDEA system [17] supports a PDP with SAML interface that can use XACML for its decisions. In case SAML serves as an interface to PDPs, it requires that the requesters of an authorization decision support the corresponding protocol binding, usually SOAP over HTTP. The public key cryptography for the verification of SAML documents also introduces complexity at the client. Our approach with `pam_xacml` shields the application from the communication interface with the PDP, by separating the PDP communication from the interface with the application.

3 Pluggable Authorization with XACML

This Section introduces `pam_xacml`, an authorization module based on the Linux Pluggable Authentication Modules (PAM) [22] for XACML [19] PDPs. Pluggable Authentication Modules were proposed by Sun Microsystems and were standardized by the Open Software Foundation in RFC 86.0 [22]. PAM allows for the development of authentication modules that can easily be plugged into the PAM library and hide the specifics of the authentication process from the applications. There also exist PAM modules that allow for basic authorization decisions. `pam_time` makes access dependent on the local time. `pam_ldap` retrieves user specific data (e.g. access levels) from the LDAP database, possibly using host name or time as selection criteria. All these methods are isolated and lack support for unified authorization policies.

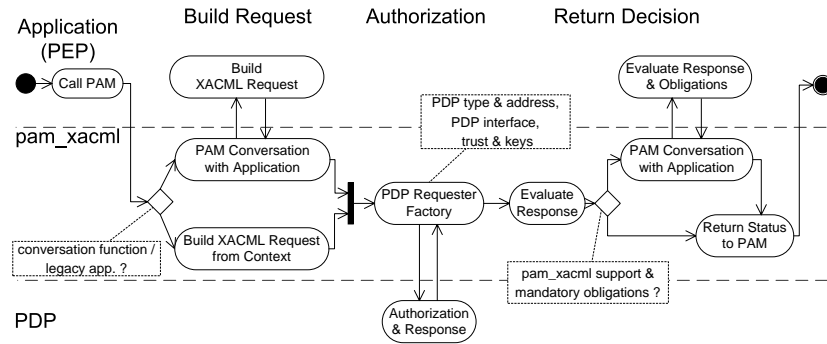


Fig. 1. Request Creation, Authorization at PDP and Response Evaluation

3.1 pam_xacml for unmodified Applications that support PAM

We introduce a PAM module to enable XACML authorization policies for arbitrary PAM enabled applications. `pam_xacml`¹ is an intermediary between application and XACML PDP (see Figure 1). The communication with the PDP works completely isolated from the application. The application is mainly concerned about the result of the authorization, hence it is sufficient for a PAM authorization module to fail upon negative decisions or succeed if all required modules returned positive results.

The authorization request can incorporate information from different sources. Existing PAM applications provide usually the user name and additional attributes which might not always be present, for instance, the address of the requester. Context information is accessible through the execution environment, for example, host name, system load, and local time. Services and databases can provide user related attributes, such as, account expiration date or account balance.

Any application that uses PAMs `pam_acct_mgmt` for account management or `pam_sm_authenticate` for authentication can simply add the `pam_xacml` module to the list of required PAM modules. PAM modules are stackable, that allows for rules that require all modules or a subset of the modules to succeed, for declaring a user to be authorized and be authenticated. We can take advantage of the existing authentication modules and demand a PAM configuration that enforces the user identities to be authenticated, before authorization is invoked. Using `pam_xacml` for authorization introduces a central instance in the system that allows to establish trust with PDPs, instead of configuring each application to authenticate and trust a given PDP. The keying material can be stored in one location of the OS, and thereby reduce the risk of the keys to become compromised. The module can monitor authorization requests and decisions and provide detailed log traces for the tracking of access requests and for the identification of possible attacks. The stacking of modules allows for the invocation

¹ available at <http://pamxacml.sourceforge.net>

of multiple PDPs for one authorization request by the application. An example of the stacking of `pam_xacml` is a scenario where the host has its own PDP for access control for local resources (e.g. files, applications), the PDP of the department decides upon access to the resources in the infrastructure (e.g. database, printer), whereas the corporate PDP decides if the organization wide business rules permit the request.

Another important aspect is the variety of communication methods that exist for querying remote XACML PDPs. SAML [2] allows for authorization requests and the transport of authorization decisions. SOAP [10] is a popular Web Service protocol that can be used to transport native XACML authorization requests. Data origin authentication, confidentiality, message integrity and replay protection introduce even more complexity for the communication. Viable options for secure authorization are HTTPS, port tunneling over SSH, or the use of WS-Security [15]. `pam_xacml` hides all these details from the application and provides great flexibility to introduce new access methods and additional PDPs to the infrastructure, with only minimal impact on the system configuration.

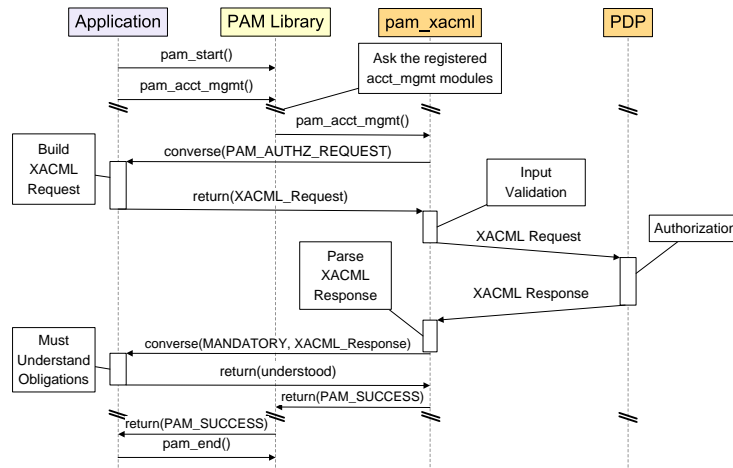


Fig. 2. Authorization with the PAM Conversation Mechanism

3.2 Applications with Authorization Interface

The prior Section assumed applications without any special support for `pam_xacml`. This limits the available information for the authorization. The application itself knows best about the parameters of the service request. The use of `pam_xacml` for existing applications can only regulate if access to an application as a whole should be possible. It has no means to limit access to certain actions on the application.

The basic idea is to give the application full control over the authorization by building the authorization request within the application. Authorization interfaces of applications are long lived, because they are linked to control flows deep within the programs. They must provide exhaustive information about the current situation and ensure effective enforcement of the access control decisions. We decided to go for the XACML policy language [19] for the interface with the application, because it has a set of common data types, well-defined semantic and has a straightforward XML language for authorization request and authorization response. The advantage of the XACML interface is, that it provides an migration path for applications that evolve towards using unified policies. These applications can start with `pam_xacml` to integrate support for XACML request and response XML documents and can be extended for native support of SAML XACML authorization later on. Alternative authorization interfaces with the application, such as AZN API [21] or GA-API [8] can be added in the future. The application builds the raw XACML request and lets `pam_xacml` handle the communication and security. The application does not necessarily need to understand the XACML response, because `pam_xacml` parses the response and signals an authorization failure to PAM if the PDP denied access. This approach allows for a lightweight integration of the authorization interface in the application. We use the PAM Conversation mechanism that allows applications to dynamically link against the PAM library, without the need to recompile the application when new PAM modules become available. The conversation mechanism is invoked to transport a message type and message content between application and PAM module. `pam_xacml` introduces new message types to enable an application to pass authorization request and to receive authorization response (see Figure 2). An application can supply detailed specifications about requested resources and intended actions. Positive decisions can carry Obligations to put constraints on permitted actions. An example for an Obligation is a bandwidth constraint on the communication, enforced by a middlebox. We implemented a template based XACML request builder with the XACML aware PAM conversation function that needs less than 60 lines of C code in the application and works without any XML library. Future work is to add support for non-XACML PDPs, transparent for the applications.

4 Experiments on Distributed Authorization

Policy enforcement is usually pessimistic as that it continues only with processing of a request, after a positive authorization decision was taken, hence it has a significant impact on session establishment. None of the publications we surveyed in Section 2 presented a detailed analysis on the performance of distributed authorization. We will focus on the impact of communication and cryptography, because no matter how fast a PDP policy engine implementation will become, the communication cost for making the request and receiving the decision will always be present.




Fig. 3. Consecutive Authorization Decisions

4.1 Reference Scenario with Consecutive Authorization Decisions

Access control is guided by the principle of minimal exposure to adversaries. Services, applications and infrastructure have their own policy enforcement code in place. One service access usually involves multiple components, each of which must authorize the access. This architecture leads to multiple lines of defense that an intruder must circumvent. Service interfaces, for instance, are usually complex, and subject to change, making it hard to guarantee correct enforcement of access control decisions and assure the absence of exploits. As opposed to firewalls that are verifiable reference monitors and can filter most malicious requests, but lack insight into the application logic.

The Figure 3 introduces the reference scenario with consecutive authorizations of a SOAP request, using `pam_xacml` at each step. The firewall has an integrated SOCKS[16] server that relays communication to the web server after successful authentication and authorization of the requester. We used the SS5² SOCKS server and extended it to support bandwidth restrictions specified by XACML Obligations. SOCKS knows about user identities and requested service addresses. Hence, it can only authorize that a user is allowed to access the web sever at all. The web server in turn receives a HTTP request specifying an URI which can now be checked. The Web Service can authorize at the finest granularity, by using parameters like, intended action on the resource, resource specification, current state, and context for its authorization request. The Web Service and the PDP were accessible via SOAP interfaces through a stack of Apache Axis, the Tomcat servlet container and the Apache web server³. All machines were installed with a standard configuration of Fedora Core 4 and had an Intel Xeon 2.80GHz CPU with 1GB of memory, except the client which had an Intel Pentium 4 2.26GHz CPU with 512MB memory. The computers were connected by a 100MBit/s LAN with Round Trip Times (RTT) smaller than 0.1 milliseconds. We used different methods for issuing authorization requests to the PDP: plain SOAP, SOAP over TLS with the OpenSSL library and SOAP over an SSH tunnel. We used 1024bit RSA certificates for TLS server and client to assure

² <http://ss5.sourceforge.net>

³ Apache, Tomcat, Axis available at <http://www.apache.org>

mutual authentication. Port forwarding with the SSH tunnel was established before each experiment and was authenticated by the server certificate and challenge response. Refer to [20] for a measurement study on the performance and throughput of different security protocols. We instrumented the client application and the `pam_xacml` module and used `tcpdump`⁴ to determine the exact time stamps of packets arriving at the host and packets that are being sent from the host.




Fig. 4. Breakdown of Service Access Time for 1/2/3 Authorization Requests

4.2 Measurements

We were interested in the delay of single requests, for this reason, we made 100 requests with a time of 3 seconds between each request. We performed measurements for one, two or three consecutive authorizations, for each protocol binding (see Figure 4). The processing time for obtaining a decision at the PDP was virtually the same for all experiments. The interpretation of the performance numbers of the PDP is limited, because we use Sun’s XACML⁵ implementation with a simple policy, not comparable with real world policies. We determined the cost for a PDP decision without communication and SOAP overhead as 9 ms. The PDP decision inclusive the handling by the Operating System, the server and the SOAP processing took 15ms for one authorization. The `pam_xacml` module introduced a delay of at most 5ms for the PAM execution, the communication with the application and the XML processing. The delay for the whole communication was below 2ms. The signaling from the application with SOCKS and

⁴ <http://www.tcpdump.org>

⁵ <http://sunxacml.sourceforge.net>

the service request/response took 22ms in total. TLS cryptography introduced a significant overhead of 35ms for the connection handshake and derivation of the session key. The already established secure link with ssh took 4ms for the cryptography, only little more than the unencrypted messages.

In case of consecutive authorizations, the delay of all components involved in the authorization grows linearly with the number of authorization decisions. The whole process with 3 authorization requests takes a client between 90ms and 180ms over a secure link.




Fig. 5. Delay Contributions to Time for Authorization

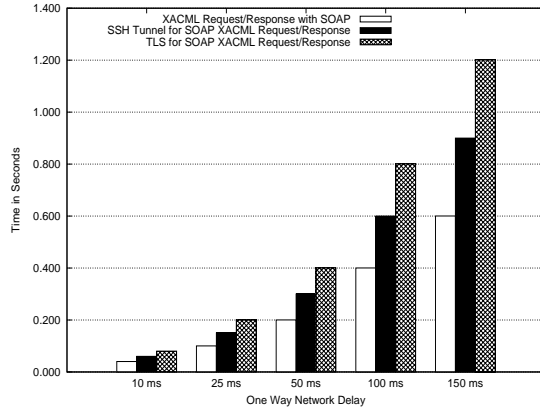


Fig. 6. Impact of Propagation Delay on Authorization for different Communication Protocols

4.3 Impact of Propagation Delay

The impact of the communication on the authorization was negligible in the testbed. As many organizations have global scope, we want to estimate the impact of the propagation delay on the authorization.

Many factors contribute to the end-to-end delay of authorization (see Figure 5): The processing time α_n in the Operating System and within the client and service programs. The delay β_n for sending the packets that belong to one message waiting in the queue of the Operating System. The β_n delay consists mainly of the serialization delay σ_i of the individual packets. The σ_i can be calculated from the bandwidth of the link and size of the data to be sent. The bandwidth is usually constrained by the access network of client and server. The network delay φ_n is influenced by factors such as queuing and forwarding at routers, packet loss, effects of multi-path routing and most important, the distance of the link. The distance of the link determines the propagation delay at roughly $5\mu s$ per km. The time for a queued message to arrive at its destination is therefore $\beta_n + \varphi_n$. We observed always the same sequence of messages in our setup, for instance,

only after a request has been fully received a response was sent. We know the exact times α_n of each message after it has been completely received at a host and the first packet of the reply is going to be sent. We can approximate the T_{authz} for a complete authorization exchange as $T_{authz} \approx \sum \beta_n + \varphi_n + \alpha_n$, where n is one message during the authorization. We can estimate how larger propagation delays impact the authorization performance, because the processing time α_n is independent of the network delay and longer routes in the network have small impact on β_n . The effects of longer routes through the network is mostly independent of the packet size and is therefor dominated by φ_n which is dependent on the number of messages sent through the network. We can substitute the network delay in the testbed with a more realistic fictive network delay φ'_n for each message (see Figure 6). Our numbers tend to underestimate the total delay, because we neither consider packet loss and fragmentation, nor the effects of multi path routing which all contribute to the delay variance in real networks. We can now calculate a lower bound on the communication delay, based on real world RTT, measured with ping from our network. Accessing a PDP in Europe (50ms RTT) would cost 200ms using TLS, 150ms with SSH and for a SOAP document 100ms. We confirmed our methodology by analyzing delays of tcpdump traces of HTTPS/TLS connections over large distances. As we argued before, the delay of consecutive authorization would grow linearly with the number of authorization decisions. Hence, our reference scenario would have an noticeable communication delay of 600ms using TLS plus the 180ms for the cryptographic operations as presented in Section 4.2. We confirmed with additional measurements, that the effect of WS-Security [14] encrypted and signed SOAP documents experiences comparable network delay contributions as with our SOAP only scenario.

The conclusion from these results is to locate the PDP close to the applications that need authorization and to use a persistent secure channel. One optimization could be, albeit not always applicable, to use optimistic authorization to start processing a request before the decision is taken and to discard and roll back changes to system upon a negative outcome.

5 Conclusion

XACML is a policy language for authorization that is gaining momentum in research and standardization. It can greatly simplify administration if it allows to regulate access to resources in distributed systems with one policy set. The success of XACML as enabler for unified policies depends on the extent existing applications are able to issue authorization requests and act upon authorization decisions. We harvest the power of XACML for applications by providing an authorization module for the PAM system. It is pluggable into existing applications and is stackable to authorize with a sequence of PDPs for one request. Applications that use PAM for authentication or account verification can directly benefit from the `pam_xacml` module without any changes to the code. XACML support can simply be enabled in the PAM configuration of the application. As

this mechanism is restricted in the information that it can obtain from the application, and as the enforcement decision boils down to permitting or denying access to an application in total, an application interface for authorization is desirable. We introduce new message types for the PAM conversation function, to allow an application to express all parameters relevant for authorization, such as the requesting subject, accessed resources or intended actions. The authorization response can contain obligations to restrict permissible actions. The proposed `pam_xacml` messaging interface allows for more powerful authorization compared to legacy applications and can open a migration path for application developers towards full XACML support, by first supporting the new PAM message types for authorization and adding communication interfaces for direct communication with the PDPs later on. A measurement study on the impact of consecutive authorization decisions showed the cost of using unified policies in distributed systems. We expect `pam_xacml` to have great potential for bringing XACML support to existing systems and for realizing unified policies.

References

1. Rafae Bhatti, James Joshi, Elisa Bertino, and Arif Ghafoor. Access Control in Dynamic XML-Based Web-Services with X-RBAC. In Liang-Jie Zhang, editor, *Proceedings of the International Conference on Web Services, ICWS '03, June 23 - 26, 2003, Las Vegas, Nevada, USA*, pages 243–249. CSREA Press, 2003.
2. Scott Cantor, John Kemp, Rob Philpott, and Eve Maler. Assertions and Protocols for the OASIS Security Assertion Markup Language (SAML) V2.0. Standard, OASIS, March 2005.
3. David W Chadwick and Alexander Otenko. The PERMIS X.509 Role Based Privilege Management Infrastructure. In *SACMAT'02*. ACM, June 3-4 2002.
4. N. Damianou, N. Dulay, E. C. Lupu, and M. Sloman. Ponder: a language for specifying security and management policies for distributed systems. *Imperial College Research Report DoC 2000/1*, 2000.
5. Yuri Demchenko, Leon Gommans, and Cees de Laat. Using SAML and XACML for Complex Authorisation Scenarios in Dynamic Resource Provisioning. In *ARES '07: Proceedings of the The Second International Conference on Availability, Reliability and Security*, pages 254–262, Washington, DC, USA, 2007. IEEE Computer Society.
6. David Durham, Jim Boyle, Ron Cohen, Shai Herzog, Raju Rajan, and Arun Sastry. *RFC 2748: The COPS (Common Open Policy Service) Protocol*. The Internet Society, January 2000.
7. Kathi Fisler, Shriram Krishnamurthi, Leo A. Meyerovich, and Michael Carl Tschantz. Verification and Change-Impact Analysis of Access-Control Policies. In *ICSE '05: Proceedings of the 27th international conference on Software engineering*, pages 196–205, New York, NY, USA, 2005. ACM.
8. Grig Gheorghiu, Tatyana Ryutov, and Clifford Neuman. Authorization for Meta-computing applications. In *Proceedings of the 7th IEEE International Symposium on High Performance Distributed Computing*, July 28-31 1998.
9. Leon Gommans, Cees de Laat, Bas van Oudenaarde, and Arie Taal. Authorization of a QoS path based on generic AAA. *Future Gener. Comput. Syst.*, 19(6):1009–1016, 2003.

10. Martin Gudgin, Marc Hadley, Noah Mendelsohn, Jean-Jacques Moreau, Henrik Frystyk Nielsen, Anish Karmarkar, and Yves Lafon. SOAP Version 1.2 Part 1: Messaging Framework (Second Edition). W3C Recommendation, April 2007.
11. Rajeev Gupta and Manish Bhide. A Generic XACML Based Declarative Authorization Scheme for Java. In Sabrina De Capitani di Vimercati, Paul F. Syverson, and Dieter Gollmann, editors, *ESORICS*, volume 3679 of *Lecture Notes in Computer Science*, pages 44–63. Springer, 2005.
12. Sushil Jajodia, Pierangela Samarati, and V. S. Subrahmanian. A Logical Language for Expressing Authorizations. In *SP '97: Proceedings of the 1997 IEEE Symposium on Security and Privacy*, page 31, Washington, DC, USA, 1997. IEEE Computer Society.
13. L. Kagal, T. Finin, and A. Joshi. A policy language for a pervasive computing environment, 2003.
14. Christiaan Lamprecht and Aad van Moorsel. Performance Measurement of Web Services Security Software. In *21st UK Performance Engineering Workshop*, 2005.
15. Kelvin Lawrence, Chris Kaler, Anthony Nadalin, Chris Kaler, Ronald Monzillo, and Phillip Hallam-Baker. Web Services Security: SOAP Message Security 1.1 (WS-Security 2004). W3C Recommendation, February 2006.
16. M. Leech, M. Ganis, Y. Lee, R. Kuris, D. Koblas, and L. Jones. RFC 1928, SOCKS Protocol Version 5, June 1996.
17. R. Lepro. Cardea: Dynamic access control in distributed systems. Technical Report NAS Technical Report NAS-03-020, NASA Advanced Supercomputing (NAS) Division, Moffett Field, CA 94035, November 2003.
18. Markus Lorch, Seth Proctor, Rebekah Lepro, Dennis Kafura, and Sumit Shah. First experiences using XACML for access control in distributed systems. In *XMLSEC '03: Proceedings of the 2003 ACM workshop on XML security*, pages 25–37, New York, NY, USA, 2003. ACM.
19. Tim Moses. eXtensible Access Control Markup Language (XACML) Version 2.0. OASIS Standard, February 2005.
20. Heiko Niedermayer, Andreas Klenk, and Georg Carle. The Networking Perspective of Security Performance - a Measurement Study. In *MMB 2006, Nürnberg, Germany*, March 2006.
21. The Open Group. *Authorization (AZN) API*. Jan. 2000. ISBN: 1-85912-266-3.
22. Vipin Samar and Roland J. Schemers. Unified Login with Pluggable Authentication Modules (PAM). Open Software Foundation: Request For Comments RFC 86.0, October 1995.
23. John R. Vollbrecht, Pat R. Calhoun, Stephen Farrell, Leon Gommans, George M. Gross, Betty de Bruijn, Cees T.A.M. de Laat, Matt Holdrege, and David W. Spence. *RFC 2904: AAA Authorization Framework*. The Internet Society, Aug. 2000.