Network Working Group
Request for Comments: 4745
Category: Standards Track

H. Schulzrinne
Columbia U.
H. Tschofenig
Siemens Networks GmbH & Co KG
J. Morris
CDT
J. Cuellar
Siemens
J. Polk
J. Rosenberg
Cisco

February 2007

Common Policy: A Document Format for Expressing Privacy Preferences

Status of This Memo

Copyright Notice

Abstract

This document defines a framework for authorization policies controlling access to application-specific data. This framework combines common location- and presence-specific authorization aspects. An XML schema specifies the language in which common policy rules are represented. The common policy framework can be extended to other application domains.

Table of Contents

1. Introduction

This document defines a framework for creating authorization policies for access to application-specific data. This framework is the result of combining the common aspects of single authorization systems that more specifically control access to presence and location information and that previously had been developed separately. The benefit of combining these two authorization systems is two-fold. First, it allows building a system that enhances the value of presence with location information in a natural way and reuses the same underlying authorization mechanism. Second, it encourages a more generic authorization framework with mechanisms for extensibility. The applicability of the framework specified in this document is not limited to policies controlling access to presence and location information data, but can be extended to other application domains.

The general framework defined in this document is intended to be accompanied and enhanced by application-specific policies specified elsewhere. The common policy framework described here is enhanced by domain-specific policy documents, including presence [7] and location [8]. This relationship is shown in Figure 1.

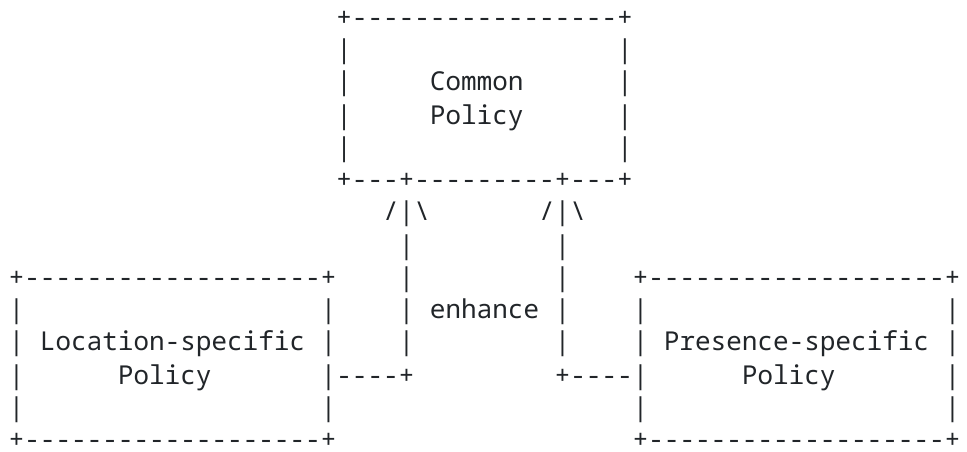

Figure 1: Common Policy Enhancements

This document starts with an introduction to the terminology in Section 2, an illustration of basic modes of operation in Section 3, a description of goals (see Section 4) and non-goals (see Section 5) of the policy framework, followed by the data model in Section 6. The structure of a rule, namely, conditions, actions, and transformations, is described in Sections 7, 8, and 9. The procedure for combining permissions is explained in Section 10 and used when conditions for more than one rule are satisfied. A short description

of meta policies is given in Section 11. An example is provided in Section 12. The XML schema will be discussed in Section 13. IANA considerations in Section 15 follow security considerations in Section 14.

2. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT","RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [1].

This document introduces the following terms:

PT - Presentity / Target: The PT is the entity about whom information has been requested.

RM - Rule Maker: The RM is an entity that creates the authorization rules that restrict access to data items.

PS - (Authorization) Policy Server: This entity has access to both the authorization policies and the data items. In location-specific applications, the entity PS is labeled as location server (LS).

WR - Watcher / Recipient: This entity requests access to data items of the PT. An access operation might be a read, a write, or any other operation.

A policy is given by a 'rule set' that contains an unordered list of 'rules'. A 'rule' has a 'conditions', an 'actions', and a 'transformations' part.

The term 'permission' indicates the action and transformation components of a 'rule'.

The term 'using protocol' is defined in [9]. It refers to the protocol used to request access to and to return privacy-sensitive data items.

3. Modes of Operation

The abstract sequence of operations can roughly be described as follows. The PS receives a query for data items for a particular PT, via the using protocol. The using protocol (or more precisely, the authentication protocol) provides the identity of the requestor, either at the time of the query or at the subscription time. The authenticated identity of the WR, together with other information provided by the using protocol or generally available to the server,

is then used for searching through the rule set. All matching rules are combined according to a permission combining algorithm described in Section 10. The combined rules are applied to the application data, resulting in the application of privacy based on the transformation policies. The resulting application data is returned to the WR.

Three different modes of operation can be distinguished:

3.1. Passive Request-Response - PS as Server (Responder)

In a passive request-response mode, the WR queries the PS for data items about the PT. Examples of protocols following this mode of operation include HTTP, FTP, LDAP, finger, and various remote procedure call (RPC) protocols, including Sun RPC, Distributed Computing Environment (DCE), Distributed Component Object Model (DCOM), common object request broker architecture (Corba), and Simple Object Access Protocol (SOAP). The PS uses the rule set to determine whether the WR is authorized to access the PT's information, refusing the request if necessary. Furthermore, the PS might filter information by removing elements or by reducing the resolution of elements.

3.2. Active Request-Response - PS as Client (Initiator)

Alternatively, the PS may contact the WR and convey data items. Examples include HTTP, SIP session setup (INVITE request), H.323 session setup or SMTP.

3.3. Event Notification

Event notification adds a subscription phase to the "Active Request-Response - PS as Client (Initiator)" mode of operation. A watcher or subscriber asks to be added to the notification list for a particular presentity or event. When the presentity changes state or the event occurs, the PS sends a message to the WR containing the updated state. (Presence is a special case of event notification; thus, we often use the term interchangeably.)

In addition, the subscriber may itself add a filter to the subscription, limiting the rate or content of the notifications. If an event, after filtering by the rule-maker-provided rules and by the subscriber-provided rules, only produces the same notification content that was sent previously, no event notification is sent.

DOCKET ALARM

Explore Litigation Insights

Docket Alarm provides insights to develop a more informed litigation strategy and the peace of mind of knowing you're on top of things.

Real-Time Litigation Alerts

Keep your litigation team up-to-date with **real-time alerts** and advanced team management tools built for the enterprise, all while greatly reducing PACER spend.

Our comprehensive service means we can handle Federal, State, and Administrative courts across the country.

Advanced Docket Research

With over 230 million records, Docket Alarm's cloud-native docket research platform finds what other services can't. Coverage includes Federal, State, plus PTAB, TTAB, ITC and NLRB decisions, all in one place.

Identify arguments that have been successful in the past with full text, pinpoint searching. Link to case law cited within any court document via Fastcase.

Analytics At Your Fingertips

Learn what happened the last time a particular judge, opposing counsel or company faced cases similar to yours.

Advanced out-of-the-box PTAB and TTAB analytics are always at your fingertips.

API

Docket Alarm offers a powerful API (application programming interface) to developers that want to integrate case filings into their apps.

LAW FIRMS
Build custom dashboards for your attorneys and clients with live data direct from the court.

Automate many repetitive legal tasks like conflict checks, document management, and marketing.

FINANCIAL INSTITUTIONS
Litigation and bankruptcy checks for companies and debtors.

E-DISCOVERY AND LEGAL VENDORS
Sync your system to PACER to automate legal marketing.