

Network Working Group
Request for Comments: 3521
Category: Informational

L-N. Hamer
B. Gage
Nortel Networks
H. Shieh
AT&T Wireless
April 2003

Framework for Session Set-up with Media Authorization

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2003). All Rights Reserved.

Abstract

Establishing multimedia streams must take into account requirements for end-to-end QoS, authorization of network resource usage and accurate accounting for resources used. During session set up, policies may be enforced to ensure that the media streams being requested lie within the bounds of the service profile established for the requesting host. Similarly, when a host requests resources to provide a certain QoS for a packet flow, policies may be enforced to ensure that the required resources lie within the bounds of the resource profile established for the requesting host.

To prevent fraud and to ensure accurate billing, this document describes various scenarios and mechanisms that provide the linkage required to verify that the resources being used to provide a requested QoS are in-line with the media streams requested (and authorized) for the session.

Table of Contents

1.	Introduction.....	2
2.	Conventions used in this document.....	3
3.	Definition of terms.....	4
4.	The Coupled Model.....	5
4.1	Coupled Model Message Flows.....	6
4.2	Coupled Model Authorization Token.....	8
4.3	Coupled Model Protocol Impacts.....	8
5.	The Associated Model <<using One Policy Server>>.....	8
5.1	Associated Model Message Flows <<using One Policy Server>>.....	9
5.2	Associated Model Authorization Token <<using One Policy Server>>.....	11
5.3	Associated Model Protocol Impacts <<using One Policy Server>>.....	11
5.4	Associated Model Network Impacts <<using One Policy Server>>.....	12
6.	The Associated Model <<using Two Policy Servers>>.....	12
6.1	Associated Model Message Flows <<using Two Policy Servers>>.....	13
6.2	Associated Model Authorization Token <<using Two Policy Servers>>.....	15
6.3	Associated Model Protocol Impacts <<using Two Policy Servers>>.....	16
7.	The Non-Associated Model.....	16
7.1	Non-Associated Model Message Flow.....	17
7.2	Non-Associated Model Authorization Token.....	19
7.3	Non-Associated Model Protocol Impacts.....	19
8.	Conclusions.....	20
9.	Security Considerations.....	21
10.	Normative References.....	22
11.	Informative References.....	23
12.	Acknowledgments.....	23
13.	Authors' Addresses.....	24
14.	Full Copyright Statement.....	25

1. Introduction

Various mechanisms have been defined through which end hosts can use a session management protocol (e.g., SIP [6]) to indicate that QoS requirements must be met in order to successfully set up a session. However, a separate protocol (e.g., RSVP [7]) is used to request the resources required to meet the end-to-end QoS of the media stream. To prevent fraud and to ensure accurate billing, some linkage is

required to verify that the resources being used to provide the requested QoS are in-line with the media streams requested (and authorized) for the session.

This document describes such a linkage through use of a "token" that provides capabilities similar to that of a gate in [12] and of a ticket in the push model of [10]. The token is generated by a policy server (or a session management server) and is transparently relayed through the end host to the edge router where it is used as part of the policy-controlled flow admission process.

In some environments, authorization of media streams can exploit the fact that pre-established relationships exist between elements of the network (e.g., session management servers, edge routers, policy servers and end hosts). Pre-established relationships assume that the different network elements are configured with the identities of the other network elements and, if necessary, are configured with security keys, etc. required to establish a trust relationship. In other environments, however, such pre-established relationships may not exist either due to the complexity of creating these associations a priori (e.g., in a network with many elements), or due to the different business entities involved (e.g., service provider and access provider), or due to the dynamic nature of these associations (e.g., in a mobile environment).

In this document, we describe these various scenarios and the mechanisms used for exchanging information between network elements in order to authorize the use of resources for a service and to coordinate actions between the session and resource management entities. Specific extensions to session management protocols (e.g., SIP [6], H.323), to resource reservation protocols (e.g., RSVP [4], YESSIR) and to policy management protocols (e.g., COPS-PR [9], COPS-RSVP [3]) required to realize these scenarios and mechanisms are beyond the scope of this document.

For clarity, this document will illustrate the media authorization concepts using SIP for session signalling, RSVP for resource reservation and COPS for interaction with the policy servers. Note, however, that the framework could be applied to a multimedia services scenario using different signalling protocols.

2. Conventions used in this document

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 BCP 14, RFC 2119 [1].

3. Definition of terms

Figure 1 introduces a generic model for session establishment, QoS and policy enforcement.

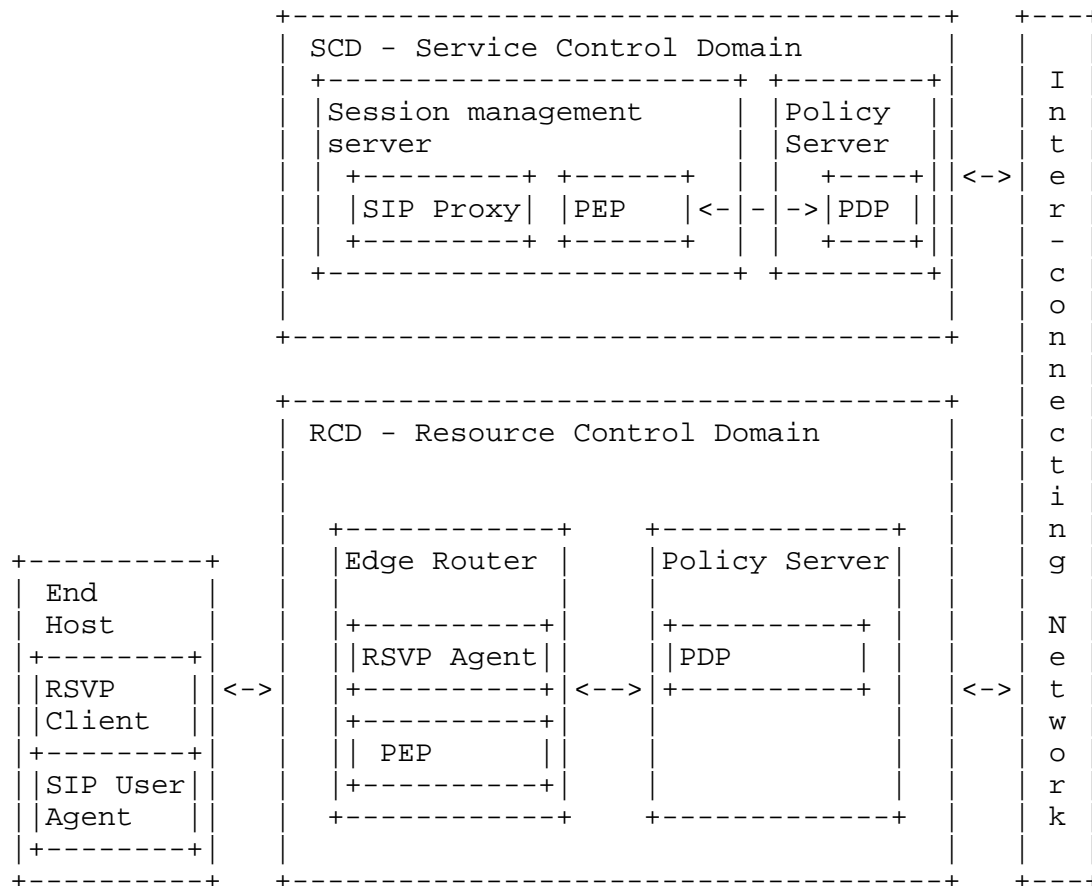


Figure 1: Generic media authorization network model

EH - End Host: The End Host is a device used by a subscriber to access network services. The End Host includes a client for requesting network services (e.g., through SIP) and a client for requesting network resources (e.g., through RSVP).

ER - Edge Router: The Edge Router is a network element connecting the end host to the rest of the Resource Control Domain. The Edge Router contains a PEP to enforce policies related to resource usage in the Resource Control Domain by the End Host. It also contains a signalling agent (e.g., for RSVP) for handling resource reservation requests from the End Host.

PDP - Policy Decision Point: The PDP is a logical entity located in the Policy Server that is responsible for authorizing or denying access to services and/or resources.

PEP - Policy Enforcement Point: The PEP is a logical entity that enforces policy decisions made by the PDP. Note that other PEPs may reside in other network elements not shown in the model of Figure 1, however they will not be discussed in this document.

PS - Policy Server: The Policy Server is a network element that includes a PDP. Note that there may be a PS in the Service Control Domain to control use of services and there may be a separate PS in the Resource Control Domain to control use of resources along the packet forwarding path. Note also that network topology may require multiple Policy Servers within either Domain, however they provide consistent policy decisions to offer the appearance of a single PDP in each Domain.

RCD - Resource Control Domain: The Resource Control Domain is a logical grouping of elements that provide connectivity along the packet forwarding paths to and from an End Host. The RCD contains ER and PS entities whose responsibilities include management of resources along the packet forwarding paths. Note that there may be one or more RCDs within an autonomous domain.

SCD - Service Control Domain: The Service Control Domain is a logical grouping of elements that offer applications and content to subscribers of their services. The Session Management Server resides in the SCD along with a PS. Note that there may be one or more SCDs within an autonomous domain.

SMS - Session Management Server: The Session Management Server is a network element providing session management services (e.g., telephony call control). The Session Management Server contains a PEP to enforce policies related to use of services by the End Host. It also contains a signalling agent or proxy (e.g., for SIP) for handling service requests from the End Host.

4. The Coupled Model

In some environments, a pre-established trust relationship exists between elements of the network (e.g., session management servers, edge routers, policy servers and end hosts). We refer to this as the "coupled model", indicating the tight relationship between entities that is presumed. The key aspects of this scenario are the following:

- Policy decisions, including media authorization, are made by a single Policy Server.
- The Edge Router, Session Management Servers and Policy Server involved in establishing the session are known a priori. For example, the End Host may be configured to use a Session Management Server associated with the Edge Router to which the EH is connected.
- There are pre-defined trust relationships between the SMS and the PS and between the ER and the PS.




Figure 2: The Coupled Model

4.1 Coupled Model Message Flows

In this model, it is assumed that there is one Policy Server serving both the Service Control and Resource Control Domains and that there are pre-defined trust relationships between the PS and SMS and between the PS and ER. Communications between these entities are then possible as described below. Only the originating side flows are described for simplicity. The same concepts apply to the terminating side.

1. The End Host issues a session set-up request (e.g., SIP INVITE) to the Session Management Server indicating, among other things, the media streams to be used in the session. As part of this step, the End Host may authenticate itself to the Session Management Server.

2. The Session Management Server, possibly after waiting for negotiation of the media streams to be completed, sends a policy decision request (e.g., COPS REQ) to the Policy Server in order to determine if the session set-up request should be allowed to proceed.
3. The Policy Server sends a decision (e.g., COPS DEC) to the Session Management Server, possibly after modifying the parameters of the media to be used. Included in this response is a "token" that can subsequently be used by the Policy Server to identify the session and the media it has authorized.
4. The Session Management Server sends a response to the End Host (e.g., SIP 200 or 183) indicating that session set-up is complete or is progressing. Included in this response is a description of the negotiated media along with the token from the Policy Server.
5. The End Host issues a request (e.g., RSVP PATH) to reserve the resources necessary to provide the required QoS for the media stream. Included in this request is the token from the Policy Server provided via the Session Management Server.
6. The Edge Router intercepts the reservation request and sends a policy decision request (e.g., COPS REQ) to the Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the token from the Policy Server provided by the End Host. The Policy Server uses this token to correlate the request for resources with the media authorization previously provided to the Session Management Server.
7. The Policy Server sends a decision (e.g., COPS DEC) to the Edge Router, possibly after modifying the parameters of the resources to be reserved.
8. The Edge Router, possibly after waiting for end-to-end negotiation for resources to be completed, sends a response to the End Host (e.g., RSVP RESV) indicating that resource reservation is complete or is progressing.

4.2 Coupled Model Authorization Token

In the Coupled Model, the Policy Server is the only network entity that needs to interpret the contents of the token. Therefore, in this model, the contents of the token are implementation dependent. Since the End Host is assumed to be untrusted, the Policy Server SHOULD take measures to ensure that the integrity of the token is preserved in transit; the exact mechanisms to be used are also implementation dependent.

4.3 Coupled Model Protocol Impacts

The use of a media authorization token in the Coupled Model requires the addition of new fields to several protocols:

- Resource reservation protocol. A new protocol field or object MUST be added to the resource reservation protocol to transparently transport the token from the End Host to the Edge Router. The content and internal structure (if any) of this object SHOULD be opaque to the resource reservation protocol. For example, this is achieved in RSVP with the Policy Data object defined in [8].
- Policy management protocol. A new protocol field or object MUST be added to the policy management protocol to transparently transport the token from the Policy Server to the Session Management Server and from the Edge Router to the Policy Server. The content and internal structure (if any) of this object SHOULD be opaque to the policy management protocol. For example, this is achieved in COPS-RSVP with the Policy Data object defined in [8].
- Session management protocol. A new protocol field or object MUST be added to the session management protocol to transparently transport the media authorization token from the Session Management Server to the End Host. The content and internal structure (if any) of this object SHOULD be opaque to the session management protocol (e.g., SIP [6]).

5. The Associated Model <<using One Policy Server>>

In this scenario, there are multiple instances of the Session Management Servers, Edge Routers and Policy Servers. This leads to a network of sufficient complexity that it precludes distributing knowledge of network topology to all network entities. The key aspects of this scenario are the following:

- Policy decisions, including media authorization, are made by the same Policy Server for both the Session Management Server and the Edge Router. However, the Policy Server may change on a per-transaction basis, i.e., on a per policy request basis.
- The Edge Router, Session Management Server and Policy Server involved in establishing the session are not known a priori. For example, the End Host may be dynamically configured to use one of a pool of Session Management Servers and each of the Session Management Servers may be statically configured to use one of a pool of Policy Servers.

In another example, the End Host may be mobile and continually changing the Edge Router that its point of attachment uses to communicate with the rest of the network.

- There are pre-defined trust relationships between the SMS and the PS and between the ER and the PS.




Figure 3: The Associated Model using One Policy Server

5.1 Associated Model Message Flows <<using One Policy Server>>

In this model, it is assumed that a Policy Server can make decisions for both the Service Control and Resource Control Domains and that there are pre-defined trust relationships between the PS and SMS and between the PS and ER. Communications between these entities are then possible as described below. Only the originating side flows are described for simplicity. The same concepts apply to the terminating side.

1. The End Host issues a session set-up request (e.g., SIP INVITE) to the Session Management Server indicating, among other things, the media streams to be used in the session. As part of this step, the End Host may authenticate itself to the Session Management Server.
2. The Session Management Server, possibly after waiting for negotiation of the media streams to be completed, sends a policy decision request (e.g., COPS REQ) to the Policy Server in order to determine if the session set-up request should be allowed to proceed.
3. The Policy Server sends a decision (e.g., COPS DEC) to the Session Management Server, possibly after modifying the parameters of the media to be used. Included in this response is a "token" that can subsequently be used by the Policy Server to identify the session and the media it has authorized.
4. The Session Management Server sends a response to the End Host (e.g., SIP 200 or 183) indicating that session set-up is complete or is progressing. Included in this response is a description of the negotiated media along with the token from the Policy Server.
5. The End Host issues a request (e.g., RSVP PATH) to reserve the resources necessary to provide the required QoS for the media stream. Included in this request is the token from the Policy Server provided via the Session Management Server.
6. The Edge Router intercepts the reservation request and inspects the token to learn which Policy Server authorized the media. It then sends a policy decision request to that Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the token from the Policy Server provided by the End Host. The Policy Server uses this token to correlate the request for resources with the media authorization previously provided to the Session Management Server.
7. The Policy Server sends a decision to the Edge Router, possibly after modifying the parameters of the resources to be reserved.
8. The Edge Router, possibly after waiting for end-to-end negotiation for resources to be completed, sends a response to the End Host (e.g., RSVP RESV) indicating that resource reservation is complete or is progressing.

5.2 Associated Model Authorization Token <<using One Policy Server>>

Since the ER does not know which SMS and PS are involved in session establishment, the token MUST include:

- A correlation identifier. This is information that the Policy Server can use to correlate the resource reservation request with the media authorized during session set up. The Policy Server is the only network entity that needs to interpret the contents of the correlation identifier therefore, in this model, the contents of the correlation identifier are implementation dependent. Since the End Host is assumed to be untrusted, the Policy Server SHOULD take measures to ensure that the integrity of the correlation identifier is preserved in transit; the exact mechanisms to be used are also implementation dependent.
- The identity of the authorizing entity. This information is used by the Edge Router to determine which Policy Server should be used to solicit resource policy decisions.

In some environments, an Edge Router may have no means for determining if the identity refers to a legitimate Policy Server within its domain. In order to protect against redirection of authorization requests to a bogus authorizing entity, the token SHOULD also include:

- Authentication data. This authentication data is calculated over all other fields of the token using an agreed mechanism. The mechanism used by the Edge Router is beyond the scope of this document.

The detailed semantics of an authorization token are defined in [4].

5.3 Associated Model Protocol Impacts <<using One Policy Server>>

The use of a media authorization token in this version of the Associated Model requires the addition of new fields to several protocols:

- Resource reservation protocol. A new protocol field or object MUST be added to the resource reservation protocol to transparently transport the token from the End Host to the Edge Router. The content and internal structure of this object MUST be specified so that the Edge Router can distinguish between the elements of the token described in Section 5.2. For example, this is achieved in RSVP with the Policy Data object defined in [8].

- Policy management protocol. A new protocol field or object MUST be added to the policy management protocol to transparently transport the token -- or at least the correlation identifier -- from the Edge Router to the Policy Server. The content and internal structure of this object SHOULD be opaque to the policy management protocol. For example, this is achieved in COPS-RSVP with the Policy Data object defined in [8].
- Session management protocol. A new protocol field or object MUST be added to the session management protocol to transparently transport the media authorization token from the Session Management Server to the End Host. The content and internal structure of this object SHOULD be opaque to the session management protocol (e.g., SIP [6]).

5.4 Associated Model Network Impacts <<using One Policy Server>>

The use of a media authorization token in this version of the Associated Model requires that the Edge Router inspect the token to learn which Policy Server authorized the media. In some environments, it may not be possible for the Edge Router to perform this function; in these cases, an Associated Model using Two Policy Servers (section 6) is required.

This version of the Associated Model also requires that the Edge Router interact with multiple Policy Servers. Policy decisions are made by the same Policy Server for both the Session Management Server and the Edge Router, however the Policy Server may change on per-transaction basis. Note that the COPS framework does not currently allow PEPs to change PDP on a per-transaction basis. To use this model, a new framework must be defined for policy decision outsourcing. This model also implies that the Policy Servers are able to interact and/or make decisions for the Edge Router in a consistent manner (e.g., as though there is only a single RCD Policy Server). How this is accomplished is beyond the scope of this document.

6. The Associated Model <<using Two Policy Servers>>

In this scenario, there are multiple instances of the Session Management Servers, Edge Routers and Policy Servers. This leads to a network of sufficient complexity that it precludes distributing knowledge of network topology to all network entities. The key aspects of this scenario are the following:

- Policy decisions, including media authorization, are made by Policy Servers.

- There is a PS in the Resource Control Domain that is separate from the PS in the Service Control Domain.
- The Edge Router, Session Management Server and Policy Servers involved in establishing the session are not known a priori. For example, the End Host may be dynamically configured to use one of a pool of Session Management Servers or the End Host may be mobile and continually changing the Edge Router that it uses to communicate with the rest of the network.
- There is a pre-defined trust relationship between the SMS and the SCD PS.
- There is a pre-defined trust relationship between the ER and the RCD PS.
- There is a pre-defined trust relationship between the RCD and SCD Policy Servers.




Figure 4: The Associated Model using Two Policy Servers

6.1 Associated Model Message Flows <<using Two Policy Servers>>

In this model, it is assumed that there is one Policy Server for the Service Control Domain and a different Policy Server for the Resource Control Domain. There are pre-defined trust relationships between the SCD PS and SMS, between the RCD PS and ER and between the RCD and SCD Policy Servers. Communications between these entities are then possible as described below. Only the originating side flows are described for simplicity. The same concepts apply to the terminating side.

1. The End Host issues a session set-up request (e.g., SIP INVITE) to the Session Management Server indicating, among other things, the media streams to be used in the session. As part of this step, the End Host may authenticate itself to the Session Management Server.
2. The Session Management Server, possibly after waiting for negotiation of the media streams to be completed, sends a policy decision request (e.g., COPS REQ) to the SCD Policy Server in order to determine if the session set-up request should be allowed to proceed.
3. The SCD Policy Server sends a decision (e.g., COPS DEC) to the Session Management Server, possibly after modifying the parameters of the media to be used. Included in this response is a "token" that can subsequently be used by the SCD Policy Server to identify the session and the media it has authorized.
4. The Session Management Server sends a response to the End Host (e.g., SIP 200 or 183) indicating that session set-up is complete or is progressing. Included in this response is a description of the negotiated media along with the token from the SCD Policy Server.
5. The End Host issues a request (e.g., RSVP PATH) to reserve the resources necessary to provide the required QoS for the media stream. Included in this request is the token from the SCD Policy Server provided via the Session Management Server.
6. The Edge Router intercepts the reservation request and sends a policy decision request (e.g., COPS REQ) to the RCD Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the token from the SCD Policy Server provided by the End Host.
7. The RCD Policy Server uses this token to learn which SCD Policy Server authorized the media. It then sends an authorization request [11] to that SCD Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the token from the SCD Policy Server provided by the End Host.
8. The SCD Policy Server uses this token to correlate the request for resources with the media authorization previously provided to the Session Management Server. The SCD Policy Server sends a decision [11] to the RCD Policy Server on whether the requested resources are within the bounds authorized by the SCD Policy Server.

9. The RCD Policy Server sends a decision (e.g., COPS DEC) to the Edge Router, possibly after modifying the parameters of the resources to be reserved.
10. The Edge Router, possibly after waiting for end-to-end negotiation for resources to be completed, sends a response to the End Host (e.g., RSVP RESV) indicating that resource reservation is complete or is progressing

6.2 Associated Model Authorization Token <<using Two Policy Servers>>

Since the RCD Policy Server does not know which SMS and SCD PS are involved in session establishment, the token MUST include:

- A correlation identifier. This is information that the SCD Policy Server can use to correlate the resource reservation request with the media authorized during session set up. The SCD Policy Server is the only network entity that needs to interpret the contents of the correlation identifier therefore, in this model, the contents of the correlation identifier are implementation dependent. Since the End Host is assumed to be untrusted, the SCD Policy Server SHOULD take measures to ensure that the integrity of the correlation identifier is preserved in transit; the exact mechanisms to be used are also implementation dependent.
- The identity of the authorizing entity. This information is used by the RCD Policy Server to determine which SCD Policy Server should be used to verify the contents of the resource reservation request.

In some environments, an RCD Policy Server may have no means for determining if the identity refers to a legitimate SCD Policy Server. In order to protect against redirection of authorization requests to a bogus authorizing entity, the token SHOULD include:

- Authentication data. This authentication data is calculated over all other fields of the token using an agreed mechanism. The mechanism used by the RCD Policy Server is beyond the scope of this document.

Note that the information in this token is the same as that in Section 5.2 for the "One Policy Server" scenario.

The detailed semantics of an authorization token are defined in [4].

6.3 Associated Model Protocol Impacts <<using Two Policy Servers>>

The use of a media authorization token in this version of the Associated Model requires the addition of new fields to several protocols:

- Resource reservation protocol. A new protocol field or object MUST be added to the resource reservation protocol to transparently transport the token from the End Host to the Edge Router. The content and internal structure of this object SHOULD be opaque to the resource reservation protocol. For example, this is achieved in RSVP with the Policy Data object defined in [8].
- Policy management protocol. A new protocol field or object MUST be added to the policy management protocol to transport the token from the SCD Policy Server to the Session Management Server and from the Edge Router to the RCD Policy Server. The content and internal structure of this object MUST be specified so that the Policy Servers can distinguish between the elements of the token described in Section 6.2. For example, this is achieved in COPS-RSVP with the Policy Data object defined in [8].
- Session management protocol. A new protocol field or object MUST be added to the session management protocol to transparently transport the media authorization token from the Session Management Server to the End Host. The content and internal structure of this object SHOULD be opaque to the session management protocol (e.g., SIP [6]).

Note that these impacts are the same as those discussed in Section 5.3 for the "One Policy Server" scenario. However the use of two Policy Servers has one additional impact:

- Authorization protocol. A new protocol field or object MUST be added to the authorization protocol to transport the token from the RCD Policy Server to the SCD Policy Server. The content and internal structure of this object MUST be specified so that the Policy Servers can distinguish between the elements of the token described in Section 6.2.

7. The Non-Associated Model

In this scenario, the Session Management Servers and Edge Routers are associated with different Policy Servers, the network entities do not have a priori knowledge of the topology of the network and there are no pre-established trust relationships between entities in the Resource Control Domain and entities in the Service Control Domain. The key aspects of this scenario are the following:

- Policy decisions, including media authorization, are made by Policy Servers.
- The PS in the Resource Control Domain is separate from the PS in the Service Control Domain.
- There is a pre-defined trust relationship between the SMS and the SCD PS.
- There is a pre-defined trust relationship between the ER and the RCD PS.
- There are no pre-defined trust relationships between the ER and SMS or between the RCD and SCD Policy Servers.




Figure 5: The Non-Associated Model

7.1 Non-Associated Model Message Flow

In this model it is assumed that the policy servers make independent decisions for their respective domains, obviating the need for information exchange between policy servers. This model also enables session authorization when communication between policy servers is not possible for various reasons. It may also be used as a means to speed up session setup and still ensure proper authorization is performed.

This model does not preclude the possibility that the policy servers may communicate at other times for other purposes (e.g., exchange of accounting information).

Communications between network entities in this model is described below. Only the originating side flows are described for simplicity. The same concepts apply to the terminating side.

1. The End Host issues a session set-up request (e.g., SIP INVITE) to the Session Management Server indicating, among other things, the media streams to be used in the session. As part of this step, the End Host may authenticate itself to the Session Management Server.
2. The Session Management Server, possibly after waiting for negotiation of the media streams to be completed, sends a policy decision request (e.g., COPS REQ) to the SCD Policy Server in order to determine if the session set-up request should be allowed to proceed.
3. The SCD Policy Server sends a decision (e.g., COPS DEC) to the Session Management Server, possibly after modifying the parameters of the media to be used. Included in this response is a "token" that can subsequently be used by the RCD Policy Server to determine what media has been authorized.
4. The Session Management Server sends a response to the End Host (e.g., SIP 200 or 183) indicating that session set-up is complete or is progressing. Included in this response is a description of the negotiated media along with the token from the SCD Policy Server.
5. The End Host issues a request (e.g., RSVP PATH) to reserve the resources necessary to provide the required QoS for the media stream. Included in this request is the token from the SCD Policy Server provided via the Session Management Server.
6. The Edge Router intercepts the reservation request and sends a policy decision request (e.g., COPS REQ) to the RCD Policy Server in order to determine if the resource reservation request should be allowed to proceed. Included in this request is the token from the SCD Policy Server provided by the End Host.
7. The RCD Policy Server uses this token to extract information about the media that was authorized by the SCD Policy Server. The RCD Policy Server uses this information in making its decision on whether the resource reservation should be allowed to proceed.

The Policy Server sends a decision (e.g., COPS DEC) to the Edge Router, possibly after modifying the parameters of the resources to be reserved.

8. The Edge Router, possibly after waiting for end-to-end negotiation for resources to be completed, sends a response to the End Host (e.g., RSVP RESV) indicating that resource reservation is complete or is progressing

7.2 Non-Associated Model Authorization Token

In this model, the token MUST contain sufficient information to allow the RCD Policy Server to make resource policy decisions autonomously from the SCD Policy Server. The token is created using information about the session received by the SMS. The information in the token MUST include:

- Calling party name or IP address (e.g., from SDP "c=" parameter).
- Called party name or IP address (e.g., from SDP "c=" parameter).
- The characteristics of (each of) the media stream(s) authorized for this session (e.g., codecs, maximum bandwidth from SDP "m=" and/or "b=" parameters).
- The authorization lifetime. To protect against replay attacks, the token should be valid for only a few seconds after the start time of the session.
- The identity of the authorizing entity to allow for validation of the token.
- Authentication data used to prevent tampering with the token. This authentication data is calculated over all other fields of the token using an agreed mechanism. The mechanism used by the RCD Policy Server is beyond the scope of this document.

Furthermore, the token MAY include:

- The lifetime of (each of) the media stream(s) (e.g., from SDP "t=" parameter). This field may be useful in pre-paid scenarios in order to limit the lifetime of the session.
- The Calling and called party port numbers (e.g., from the "m=" parameter).

The detailed semantics of an authorization token are defined in [4].

7.3 Non-Associated Model Protocol Impacts

The use of a media authorization token in the Non-Associated Model requires the addition of new fields to several protocols:

- Resource reservation protocol. A new protocol field or object MUST be added to the resource reservation protocol to transparently transport the token from the End Host to the Edge Router. The content and internal structure of this object SHOULD be opaque to the resource reservation protocol. For example, this is achieved in RSVP with the Policy Data object defined in [8].
- Policy management protocol. A new protocol field or object MUST be added to the policy management protocol to transport the token from the SCD Policy Server to the Session Management Server and from the Edge Router to the RCD Policy Server. The content and internal structure of this object MUST be specified so that the Policy Servers can distinguish between the elements of the token described in Section 7.2. For example, this is achieved in COPS-RSVP with the Policy Data object defined in [8].
- Session management protocol. A new protocol field or object MUST be added to the session management protocol to transparently transport the media authorization token from the Session Management Server to the End Host. The content and internal structure of this object SHOULD be opaque to the session management protocol (e.g., SIP [6]).

8. Conclusions

This document defines three models for session set-up with media authorization:

- The Coupled Model which assumes a priori knowledge of network topology and where pre-established trust relationships exist between network entities.
- The Associated Model where there are common or trusted policy servers but knowledge of the network topology is not known a priori.
- The Non-Associated Model where knowledge of the network topology is not known a priori, where there are different policy servers involved and where a trust relationship does not exist between the policy servers.

The Associated Model is applicable to environments where the network elements involved in establishing a session have a pre-determined trust relationship but where their identities must be determined dynamically during session set up. The Non-Associated Model is applicable to environments where there is a complex network topology and/or where trust relationships between domains do not exist (e.g., when they are different business entities).

In any given network, one or more of these models may be applicable. Indeed, the model to be used may be chosen dynamically during session establishment based on knowledge of the end points involved in the call. In all cases, however, there is no need for the End Host or the Session Management Server to understand or interpret the authorization token - to them it is an opaque protocol element that is simply copied from one container protocol to another.

Finally, the framework defined in this document is extensible to any kind of session management protocol coupled to any one of a number of resource reservation and/or policy management protocols.

9. Security Considerations

The purpose of this document is to describe a mechanism for media authorization to prevent theft of service.

For the authorization token to be effective, its integrity MUST be guaranteed as it passes through untrusted network entities such as the End Host. This can be achieved by using authentication data. There is no requirement for encryption of the token since it does not contain confidential information that may be used by malicious users.

This document assumes that trust relationships exist between various network entities, as described in each of the models. The means for establishing these relationships are beyond the scope of this document.

The different interfaces between the network entities described in this document have different natures requiring different security characteristics:

- The edge router and RCD policy server MUST have a trust relationship. If necessary, this relationship can be enforced through a formal security association [14].
- The network policies exchanged over the interface between edge router and RCD policy server SHOULD be integrity protected. This can be accomplished using integrity mechanisms built into the policy control protocol (e.g., the Integrity object in COPS [2]) or through generic IP security mechanisms [14].
- The SCD and RCD policy servers MUST have a trust relationship in the associated model. If necessary, this relationship can be enforced through a formal security association [14].

- The information exchanged over the interface between policy servers SHOULD be integrity protected. This can be accomplished using integrity mechanisms built into the policy exchange protocol [2] or through generic IP security mechanisms [14].
- The end host SHOULD be authenticated by the RCD to protect against identity theft. The network resource request/responses should be protected against corruption and spoofing. Thus, the interface between host and edge router SHOULD provide integrity and authentication of messages. For example, [13] provides integrity and authentication of RSVP messages.
- The end host SHOULD be authenticated by the SCD to protect against identity theft. The session setup request/response should be protected against corruption and spoofing. Thus, the interface between host and SMS SHOULD provide integrity and authentication of messages.
- The SMS and the SCD policy server MUST have a trust relationship. If necessary, this relationship can be enforced through a formal security association [14].
- The network policies exchanged over the interface between the SMS and SCD policy server SHOULD be integrity protected. This can be accomplished using integrity mechanisms built into the policy control protocol (e.g., the Integrity object in COPS [2]) or through generic IP security mechanisms [14].

10. Normative References

- [1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [2] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R. and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.
- [3] Herzog, S., Boyle, J., Cohen, R., Durham, D., Rajan, R. and A. Sastry, "COPS usage for RSVP", RFC 2749, January 2000.
- [4] Hamer, L-N., Gage, B., Kosinski, B. and H. Shieh, "Session Authorization Policy Element", RFC 3520, April 2003.
- [5] Handley, M. and V. Jacobson, "SDP: session description protocol," RFC 2327, April 1998.

- [6] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M. and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.
- [7] Braden, R., Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReSerVation protocol (RSVP) -- version 1 functional specification," RFC 2205, September 1997.
- [8] Herzog, S., "RSVP Extensions for Policy Control", RFC 2750, January 2000.
- [9] Chan, K., Seligson, J., Durham, D., Gai, S., McCloghrie, K., Herzog, S., Reichmeyer, F., Yavatkar, R. and A. Smith, "COPS Usage for Policy Provisioning (COPS-PR)", RFC 3084, March 2001.

11. Informative References

- [10] Vollbrecht, J., Calhoun, P., Farrell, S., Gommans, L., Gross, G., de Bruijn, B., de Laat, C., Holdrege, M. and P. Spence, "AAA Authorization Framework", RFC 2904, August 2000.
- [11] de Laat, C., Gross, G., Gommans, L., Vollbrecht, J. and D. Spence, "Generic AAA Architecture", RFC 2903, August 2000.
- [12] "PacketCable Dynamic Quality of Service Specification", CableLabs, December 1999.
- [13] Baker, F., Lindell, B. and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.
- [14] Kent, S. and R. Atkinson, "Security Architecture for the Internet Protocol", RFC 2401, November 1998.

12. Acknowledgments

The authors would like to thank to following people for their useful comments and suggestions related to this document: Kwok Ho Chan, Doug Reeves, Sam Christie, Matt Broda, Yajun Liu, Brett Kosinski, Francois Audet, Bill Marshall, Diana Rawlins and many others.

13. Authors' Addresses

Louis-Nicolas Hamer
Nortel Networks
PO Box 3511 Station C
Ottawa, ON
CANADA K1Y 4H7

Phone: +1 613.768.3409
EMail: nhamer@nortelnetworks.com

Bill Gage
Nortel Networks
PO Box 3511 Station C
Ottawa, ON
CANADA K1Y 4H7

Phone: +1 613.763.4400
EMail: gageb@nortelnetworks.com

Hugh Shieh
AT&T Wireless
7277 164th Avenue NE
Redmond, WA
USA 98073-9761

Phone: +1 425.580.6898
EMail: hugh.shieh@attws.com

14. Full Copyright Statement

Copyright (C) The Internet Society (2003). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

