NSIS
Internet Draft

H. Tschofenig
Siemens
M. Buechli
S. Van den Bosch
Alcatel
H. Schulzrinne
Columbia U.
T. Chen
TU Berlin

Document:
draft-tschofenig-nsis-qos-authz-issues-00.txt
Expires: December 2003 June 2003

QoS NSLP Authorization Issues
<draft-tschofenig-nsis-qos-authz-issues-00.txt>

Status of this Memo

Abstract

Various proposals for NSIS QoS NSLPs have been published recently. The design of a QoS NSLPs has to consider more than only exchanging QoS objects. Authorization has to be handled properly to make this protocol both useful and performant. Authorization in mobile environments, unfortunately, raises additional questions. This document provides an introduction to the topic.

Table of Contents

1. Introduction

Authorization is a necessary function in order to prevent theft-of-service and to enable charging. With regard to authorization a few issues still need to be resolved to specify the protocol interaction for a QoS NSLP with regard to authorization of resource requests.

[Her95] and [Her96] give some hints about policy handling and authorization in RSVP [RFC2205]. A number of papers have been published in the meanwhile proposing numerous different procedures for handling pricing, charging and even for including micro-payment protocols. None of these proposals, however, plays a role today. To avoid proposing many new alternative ways to handle authorization we would like to draw the attention of the working group to this topic in their effort to create a QoS NSLP.

With [TB+03] we tried to address the issues of authorization although due to terminology most NSIS working group members have probably not read the draft. Some others even think that these issues are independently of the NSIS NSLP protocol itself.

We think that the following questions should be addressed during the work on a QoS NSLP:

a) Which entities are involved in computing the authorization decision?

b) How long is the authorization decision valid?

c) What information is needed to compute the authorization decision?

We will provide more details to these questions in the subsequent sections.

It should be noted that the result of the authorization process might be a "yes" or "no" decision or even a complex policy. In some cases the latter might allow to answer further authorization requests locally.

2. Terminology

This draft uses terminology described in [TB+03].

3. Which entities are involved in computing the authorization decision?

At an abstract level we have two different cases with regard to NSIS signaling:

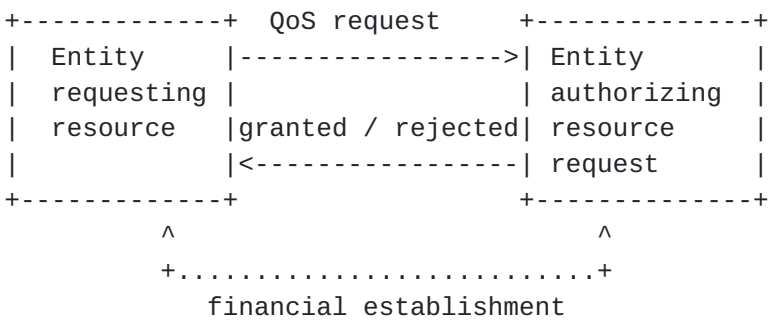

Figure 1: Two party approach

Figure 1 describes the simple and basic approach where
(a) the authorization decision is purely executed between the two entities only or
(b) where previous (out-of-band) mechanisms separated the signaling protocol from executing other entities during NSIS protocol execution.

The entity authorizing the resource request and the entity actually performing the QoS reservation are in the same administrative domain. Hence they are treated as a single logical entity.

Examples for this type of model can be found between two neighboring networks (inter-domain signaling) where a long-term contract (or other out-of-band mechanisms) exists and allows both networks to know

- how much a resource reservation costs

- how to charge the other entity (i.e. how the authorizing entity finally gets paid for the consumed resources) and

- how to authorize the resource requesting entity (i.e. associating the identifier of the protected signaling message to the identity

used in the authentication and key exchange protocol run and finally this identity to the user identity of the contract for the purpose of charging).

The consequence for an NSIS QoS NSLP protocol in this case is:
- No additional message signaling for authorization is required

- It might be necessary to include only some new objects.

- Triggering other protocols (such as credit control) might be necessary but has no impact on the NSIS signaling machinery

It might also be possible to count micro-payment protocol approaches to the two party case since it is a pure two party protocol. Fully integrating a payment protocol into NSIS, however, requires modifications to the NSIS protocol machinery itself since the message flows of NSIS and the message flows of the payment protocol might not be compatible.

Next a three party approach is presented which has two facets whereby the first variant is shown in Figure 2 and the alternative approach in Figure 3:

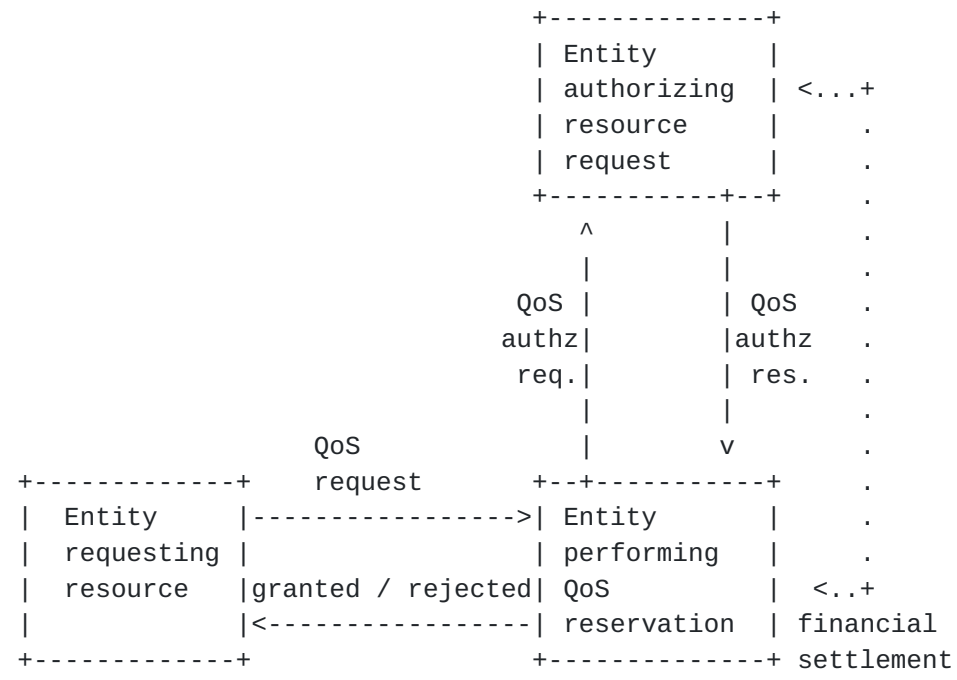

Figure 2: Three party approach

The three party approach is considerably more complex since an NSIS protocol has to enable the corresponding mechanisms to contact a third party which executes the authorization request and (if

successful) establishes a financial settlement to the entity providing the QoS reservation. The requesting entity is involved in this three party approach since it has to authenticate itself to the entity authorizing the request.

This type of configuration is common in mobility environments with per-session authorization. Section 6 gives an example of per-session authorization (per-session or even more often). Thereby a resource request by the end host is send to an RSVP node in the local network and then forwarded to the local PDP (via COPS). Since the local domain is unable to verify the request it has to be forwarded to the user's home network where authorization is provided. The response is then returned and resources are granted (in case of a successful authorization decision). The interaction between the visited network and the home network establishes the necessary financial infrastructure to latter charge the user for the consumed resources.

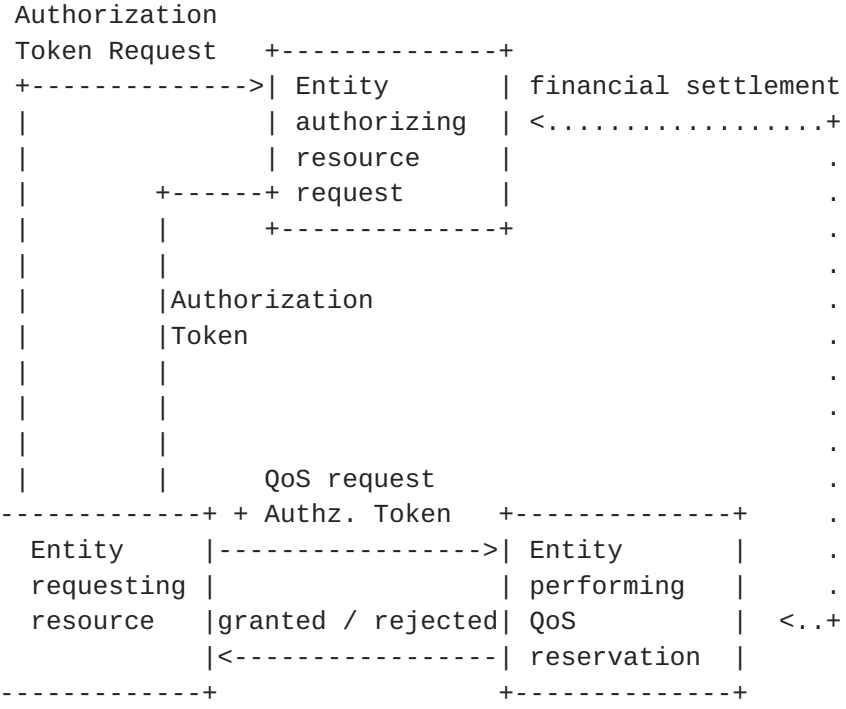

Figure 3: Token based Three party approach

The token based three party approach is applicable in environments where a previous protocol interaction is used to request authorization tokens (or something similar) to assist the authorization process at the entity performing the QoS reservation. This approach can be found in solutions where OSP Tokens [OSP] are or Authorization Tokens as used as described in [RFC3520] and in [RFC3521].

Additionally to consider are the following questions:

- A resource request might be necessary between neighboring entities only or between non-neighboring entities.

- Additionally of interest is whether authorization should be provided to more than one entity along the path.

Both issues refer to the difference between the New Jersey Turnpike and the New Jersey Parkway Model. See [TB+03] for a description of the different trust models. Furthermore Section 6 of [TB+03] is of interest when it comes to the question whether the sender or the receiver should authorize a QoS request.

4. How long is the authorization decision valid?

For the NSIS QoS NSLP protocol machinery it is important to consider at what frequency authorization decisions are made. Some possible options are:

- Per request
  (e.g. a request for more QoS resources than previously requested)

- Per session
  (e.g. only during the initial setup of a QoS resource)

- Periodically
  (authorization decision is repeated after a certain time interval)

- Event triggered
  (as soon as something changes e.g. price changes due to mobility which requires reauthorization)

The concept of a per-channel authorization (and financial establishment) is introduced in [TB+03] and tries to move a three party to a two party scenario by establishing the necessary infrastructure outside NSIS and to thereby make it simpler. Thereby it is possible to authorize QoS resource requests locally. The feasibly of this approach heavily depends on assumptions.

If authorization is, however, provided based on the three party approach then for example a periodically triggered re-authorization request requires that the third party is contacted with every authorization request. This might place a considerable burden onto the QoS signaling protocol in a mobile environment.

5. What information is needed to compute the authorization decision?

Whenever an authorization decision has to be made then there is the question which information serves as an input to the authorizing

entity. The following information items have been mentioned in the past for computing the authorization decision (in addition to the authenticated identity):

- Price

- QoS objects

- Policy rules

Policy rules include attributes like time of day, subscription to certain services, membership, etc. into consideration when computing an authorization decision.

Some of these information items are only available at certain places. By, for example, relying on policy rules it is likely that an authorization decision has to be made in the user's home network since the network administrator might not be willing to disclose the policies to other networks in order to offload the authorization decision.

In case of QoS objects it might be fairly difficult for an authorizing entity to grant a QoS authorization request since the objects by themselves might not always allow inferring to the price of the reservation. Hence in most cases it might be desirable to provide additionally some price information (if not agreed between the networks in advance).

6. Authorization example based on RSVP

This section illustrates a simple message flow based on the features offered by RSVP. We assume a mobile environment where an end host is attached to a network which is not his own home network. We do not distinguish the case where the user has no home network and where alternative means of access are used to authorize network access and other resources. A short description of the two principal network access authentication scenarios can be found in [Tsc03]. They are also applicable for this discussion.

The description in [RFC3182], in [Her95] and in [Her96] gave us the impression that RSVP aims to target authorization on the basis of an individual RSVP message. Furthermore it seems that the New Jersey Turnpike Model is the favorite model (although not directly mentioned).

Figure 4 shows a typically message flow whereby the end host starts with network access authentication before address configuration occurs. Subsequently QoS signaling with RSVP starts with a PATH message. The RSVP PATH message contains a Policy Object with a

digital signature (and the corresponding certificate) as proposed in [RFC3182].

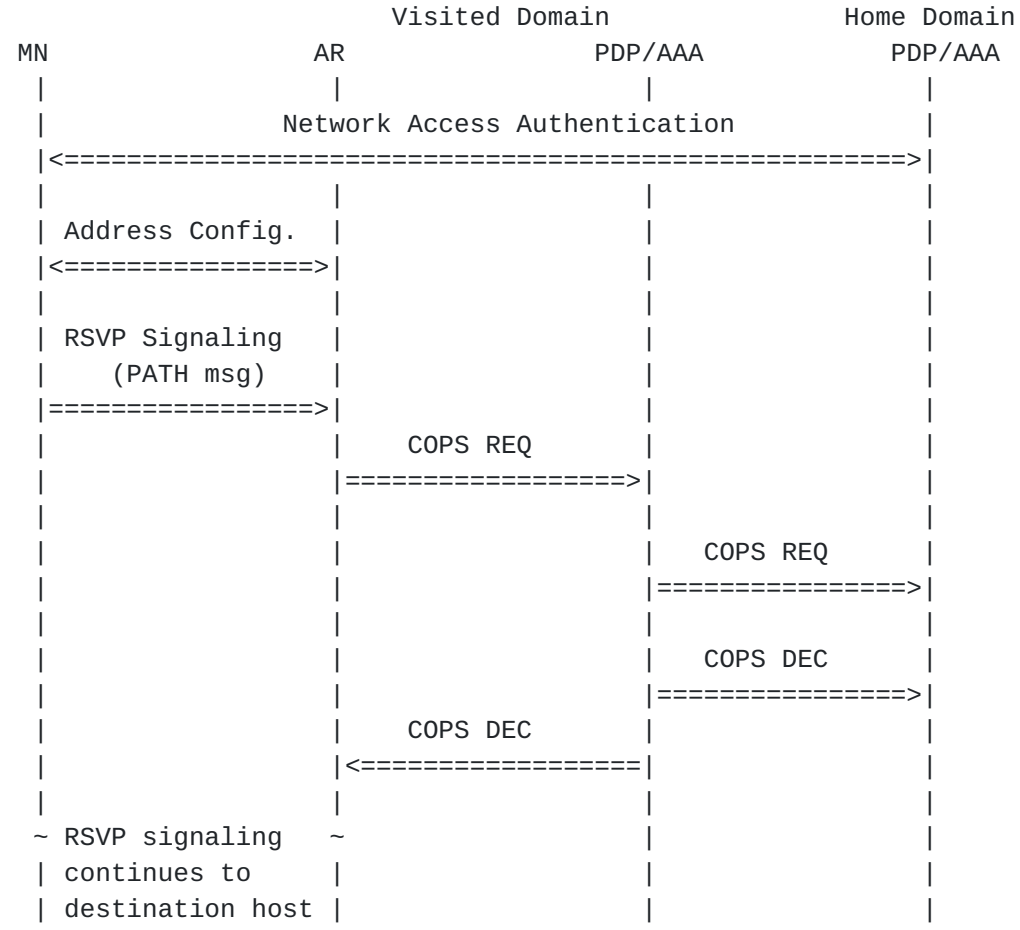

Figure 4: RSVP Signaling Message Exchange with PDP Interaction

In [Tho02] it is suggested to delegate the authorization decision to the local PDP and subsequently to the user's home PDP. This seems to be necessary if an authorization decision has to be provided for each individual session or even for each individual RSVP signaling message. Verification of the digital signature might not help with authorization in most environments.

The digital signature allows authentication of the client to the PDP at the home network. Mutual authentication is not offered and replay protection is most likely based on timestamps (although not mentioned in [RFC3182]). In addition to the Policy Object it is also necessary to forward information about the requested resources otherwise an authorization decision by the user's home PDP is worthless. Even then it is difficult for the PDP in the user's home network to perform an authorization decision since the costs of the reservation are most likely not known at this time. Since the

duration of the QoS reservation during reservation setup, the

authorization request/response scheme would have to be repeated periodically. In this sender-authorizing scheme it is difficult to determine how much resources will be actually reserved due to the nature of the RSVP PATH message with its ADSPEC object and the ability of the receiver to change the QoS reservation.

Public key based authentication between the user and his home network would typically be used because

(a) user identity confidentiality is desired or

(b) if the user authenticates itself to the local network (and not to the home network)

Since the public key based authentication proposed in [RFC3182] does not provide (a) and scenarios for (b) do not require client based public key based authentication it seems to be difficult to find a motivation for using a performance intensive mechanism without an additional benefit.

Clients today use a number of different authentication protocols such as SRP, UTMS-AKA, etc. which offer different cryptographic properties. In a mobile environment RSVP, together with COPS, simulates functionality known from Radius and Diameter. It seems to be unlikely that network operations add COPS for inter-domain signaling only although Radius and Diameter already offers the same functionality.

[RFC3182] also offers authentication based on a shared secret. For entity authentication between the end host and the user's home network this seems to be the most efficient approach although the sequence number handling might not be the best replay protection approach.

As with pk-based authentication and authentication based on symmetric keys, Kerberos authentication to the PDP in the user's home network does not provide session key distribution to the first RSVP node in the visited network. To protect signaling messages a session key for the RSVP Integrity object should be available. From a performance point of view it is highly recommended to execute this cross-realm authentication procedure only as frequently as absolutely necessary due the high overhead. If Kerberos should be additionally used to authenticate the user to the first RSVP node then additional problems occur. Kerberos cross-realm authentication does not match to AAA inter-domain handling. Several roundtrips might be required to obtain the Ticket Granting Ticket of the visited domain and finally the service ticket for either the PDP or the first policy aware RSVP router.

In case of the New Jersey Turnpike Model authorization is only provided between neighboring entities. For signaling messages which are exchanged between neighboring domains it is not necessary to perform per-session authorization by including a Policy Object. Since the neighboring domains have long-term contracts and security associations can easily be established data origin authentication is provided by the RSVP Integrity Object. The identifier used to select the key for the Integrity object can be associated with the identity which allows authorizing the QoS request. Hence we argue that the Policy Object should not be used for entity authentication between neighboring networks due to the performance restrictions and the presence of the RSVP Integrity object.

It should be noted that the policy control and the admission control procedure perform different functions although they use similar information. Both procedures might require information about the requested resources (i.e. QoS objects). The admission control procedure does not need to use user identity information or other complex policy rules for deciding whether to grant a request or not. The two entities executing the policy control and the admission control procedure do not need to be co-located or even in the same network. In the mobile scenario case it seems that admission control is executed at the local network whereas policy control is provided at the user's home network as part of the authorization procedure. Most important for determining an authorization decision at the user's home network is most likely a monetary amount - and not a QoS object. In some cases it might be, however, possible for the PDP in the user's home network to associate the cost of a QoS reservation with the provided IntServ parameter.

7. Security Considerations

This document address authorization for QoS NSLPs and tries to raise some questions about the expected functionality of this specific application signaling protocol.

A definition of authorization in the QoS environment should be created as part of a working group discussion to allow an NSLP protocol to address the corresponding security requirements.

8. Acknowledgments

We would like to thank Allison Mankin for their comments to the NSIS AAA draft. Her comments gave us the impression that we should work on a shorter draft which raises the most important open issues.

9. References

[RFC2205] R. Braden, Ed., L. Zhang, S. Berson, S. Herzog, and S. Jamin, "Resource ReSerVation protocol (RSVP) -- version 1 functional specification," RFC 2205, Internet Engineering Task Force, Sept. 1997.

[TB+03] H. Tschofenig, M. Buechli, S. Van den Bosch and H. Schulzrinne: "NSIS Authentication, Authorization and Accounting Issues", Internet Draft Internet Engineering Task Force, (work in progress), March 2003.

[Her95] Herzog, S.: "Accounting and Access Control in RSVP", Internet Draft Internet Engineering Task Force, (expired), November, 1995.

[RFC3182] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T., Herzog, S., Hess, R.: "Identity Representation for RSVP", RFC 3182, October, 2001.

[Tho02] M. Thomas: "Analysis of Mobile IP and RSVP Interactions", Internet Draft Internet Engineering Task Force, (work in progress), October 2002.

[RFC3182] Yadav, S., Yavatkar, R., Pabbati, R., Ford, P., Moore, T., Herzog, S., Hess, R.: "Identity Representation for RSVP", RFC 3182, October, 2001.

[Her96] S. Herzog: "Accounting and Access Control for Multicast Distributions: Models and Mechanisms", PhD Dissertation, University of Southern California, June 1996, available at: http://www.policyconsulting.com/herzog/cv.html.

[Tsch03] H. Tschofenig: "PANA Framework Issues", Internet Draft Internet Engineering Task Force, January 2003.

[OSP] E. T. S. Institute, "Telecommunications and internet protocol harmonization over networks (tiphon); open settlement protocol (osp) for inter- domain pricing, authorization, and usage exchange, technical specification 101 321. version 2.1.0."

[RFC3521] L. Hamer, B. Gage, and H. Shieh, "Framework for session set-up with media authorization," RFC 3521, Internet Engineering Task Force, April 2003.

[RFC3520] L. Hamer, B. Gage, B. Kosinski, and H. Shieh, "Session Authorization Policy Element", RFC 3520, Internet Engineering Task Force, April 2003.

Author's Addresses

Hannes Tschofenig
Siemens AG
Otto-Hahn-Ring 6
81739 Munich
Germany
EMail: Hannes.Tschofenig@siemens.com

Henning Schulzrinne
Dept. of Computer Science
Columbia University
1214 Amsterdam Avenue
New York, NY 10027
USA
EMail: schulzrinne@cs.columbia.edu

Sven Van den Bosch
Alcatel
Francis Wellesplein 1
B-2018
Antwerpen
Phone: 32-3-240-8103
EMail: sven.van_den_bosch@alcatel.be

Maarten B chli
Alcatel
Francis Wellesplein 1
B-2018
Antwerpen
EMail: maarten.buchli@alcatel.be

Tianwei Chen
Technical University of Berlin
Sekr. FT 5-2, Einsteinufer 25
Berlin 10587
Germany
EMail: chen@ee.tu-berlin.de

Intellectual Property Statement

claims of rights made available for publication and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementors or users of this specification can be obtained from the IETF Secretariat.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights which may cover technology that may be required to practice this standard. Please address the information to the IETF Executive Director.

Full Copyright Statement

Acknowledgement

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