Internet Engineering Task Force NSIS
Internet Draft
H. Tschofenig,M. Buechli
S. Van den Bosch, H. Schulzrinne
Siemens/Alcatel/Alcatel/Columbia
draft-tschofenig-nsis-aaa-issues-00.txt
5 February 2003
Expires: August 2003

NSIS Authentication, Authorization and Accounting Issues

STATUS OF THIS MEMO

Abstract

This document describes the implications of authentication, authorization and accounting for an NSIS QoS signaling protocol. We try to show that authorization and charging are very important for the internal machinery of a signaling protocol and for the security and trust model behind it. This document only addresses charging aspects for unicast data traffic.

1 Introduction

When RSVP [1] was designed a few assumptions had to be made. These assumptions are, however, not described in too much detail. Especially with regard to accounting and charging a number of white spots have been left open which make it difficult for providers to create a solution without considerable effort. This document tries to highlight some of

these issues and explain why NSIS should consider them during the design phase. This document does not try to introduce a new charging or accounting infrastructure and does not aim to provide a literature review of pricing mechanisms or mathematic models. Instead an abstract view on authentication, authorization and charging is provided as far as relevant for NSIS and to QoS signaling in particular.

2 Terminology

Accounting terminology used in this document tries to be consistent with [2]. NSIS terminology is taken from [3]. The term Policy Decision Point (PDP) refers to the logical entity defined in [4].

Reverse charging denotes charging for the receiver of the data traffic in contrast to charging for the data traffic transmitter (i.e. sender).

3 The Relationship between Authorization and Accounting

RSVP is currently only deployed in fairly closed environment such as enterprise networks. In such an environment authorization usually means role based access control based on some group membership or special rights to use a service. Users are typically not charged directly for their generated QoS traffic nor for QoS reservations. If the signaling messages (and thereby the QoS reservation) travel beyond the administrative domain then the enterprise network is charged and not the individual end user directly.

With mobility and telecommunication networks today authorization can (or should) be seen in an abstract form as "Is one of the signaling participants able to pay for the reservation?". This abstraction is supported by the fact that reservations require some form of penalty in order not reserve too many resources.

It might therefore be correct to say that authorization is strongly interrelated to the availability of funds/credits and therefore with charging. Some service provider might use some additionally information based on the subscriber profile/information stored data to assist in the authorization process.

4 The two Accounting/Trust Models

4.1 New Jersey Turnpike Model

On the New Jersey Turnpike, motorists pick up a ticket at the toll booth when entering the highway. At the highway exit the ticket is presented and payment is made at the toll booth for the distance driven. An abstract form of this model is given in Figure 1 where security is provided in a peer-to-peer or network-to-network fashion since the

accounting and charging model is also accomplished in the same fashion.

```
 +--------------------+  +--------------------+  +--------------------+
 |             Network |  |             Network |  |             Network |
 |               X    |  |               Y    |  |               Z    |
 |                    |  |                    |  |                    |
 |                ----------->            ----------->                |
 |                    |  |                    |  |                    |
 |                    |  |                    |  |                    |
 +--------^-----------+  +--------------------+  +---------+----------+
          |                                                .
          |                                                .
          |                                                v
       +--+---+  Data                         Data      +--+---+
       | Node |  ====================================>  | Node |
       |  A   |  Sender                      Receiver  |  B   |
       +------+                                        +------+
```

Legend:

----> Peering relationship which allows neighboring networks/entities to charge each other for the QoS reservation and data traffic

====> Data flow

..... Communication to the end host

Figure 1: New Jersey Turnpike Model

The model shown in Figure 1 uses peer-to-peer relationships between different administrative domains as a basis for accounting and charging. Based on the peering relationship a chain-of-trust is established. There are several issues which come in mind when considering this type of model:

- Since accounting and charging requires some protocol interaction with the end host it is reasonable to assume that a QoS signaling protocol is not the first protocol executed between an end host and the attached network. Typically some network access protocols are executed which establish a relationship between the user and his home network (subscription-based scenario). A more detailed description of this environment is given in Section 6. Using this

  assumption a relationship between the visited/access network (the network where the end host is currently attached) and the user's home network for the purpose of charging is already established. Hence generating additional accounting records for QoS reservations and QoS data traffic does not require a major change.

- The price for a QoS reservation needs to be determined somehow and communicated to the charged entity and to the network where the charged entity is attached. The description of this model assumes that the data sender is charged. Section 6 addresses the issue of charging for either one of the two end points.

  Section A describes two mechanisms for price distribution: in-band (or probing) and out-off band price distribution protocols

- This architecture seems to be simple enough to allow a scalable solution (ignoring reverse charging, multicast issues and the problem of determining the cost for a reservation along the entire path).

- Depending on the signaling protocol and the price distribution protocol (especially in case of an in-band protocol) it might be possible that a malicious node is able to cause harm by modifying signaling messages in such a way that the end point is charged more than intended. (TBD: This issue needs to be elaborated in more details.)

Charging for the data sender applied to this model simplifies security handling by demanding only peer-to-peer security protection. Node A would perform authentication and key establishment. The established security association (together with the session key) would allow the user to protect QoS signaling messages. The identity used during the authentication and key establishment phase would be used by Network X (see Figure 1) to perform the so-called policy-based admission control procedure. In our context this user identifier would be used to establish the necessary infrastructure to provide authorization and charging. Signaling messages later exchanged between the different networks are then also subject to authentication and authorization. The authenticated entity thereby is, however, the neighboring network and not the end host.

The New Jersey Turnpike model is attracting because of its simplicity. In [5] S. Schenker et. al. discussed various accounting implications and introduced the edge pricing model. The edge pricing model is very similar to this model with the exception that mobility and the security implications itself are not addressed.

4.2 New Jersey Parkway Model

On the New Jersey Parkway highway, drivers have to deposit 20 or 25 cents every few miles, with toll booths in the middle of the road in addition to entrance or exit ramps. (With electronic toll tags, each such toll is deducted individually.)

```
+--------------------+  +--------------------+  +--------------------+
|            Network |  |            Network |  |            Network |
|               X    |  |               Y    |  |               Z    |
|                    |  |                    |  |                    |
|                    |  |                    |  |                    |
|                    |  |                    |  |                    |
|                    |  |                    |  |                    |
+-------^ -----------+  +----------^---------+  +-----^---+----------+
        |                          |                  |   .
        |+-------------------------+                  |   .
        ||+-------------------------------------------+   .
        |||                                               .
      +-+++--+  Data                        Data       +--+---+
```

```
      | Node |  ====================================>  | Node |
      |  A   |  Sender                       Receiver   |  B   |
      +------+                                         +------+
```

Legend:

```
 ----> Direct authorization and charging relationship

 ====> Data flow

 ..... Communication to the end host
```

Figure 2: New Jersey Parkway Model

In this model one of the NSIS end points (initiator or responder) is charged directly by the traversed domains along the path. In other words, each network charges the end point only for the incurred costs in its own network. Each network maintains only local pricing information. Figure 2 shows this model in case that the data sender is charged.

The following list gives some issues caused by the selection of this model:

- Since the end point probably does not have agreements with all traversed networks there is a need for a trusted third party for authentication, authorization and financial settlement. Such a trusted third party might be a clearing house.

- Authentication and authorization of reservation requests needs to be done on a per-reservation request basis. As a possible solution, the signaling messages might carry authorization tokens from the charged entity which can be verified at the intermediate domains. Each network might have to contact the clearing house. A route change might therefore require a new authentication and authorization of the end point for the purpose of charging.

- There are, however, some concerns related to scalability and

deployment. If the NSIS initiator is located in the end host (and the NSIS initiator is charged) then the number of end points may be too large to handle by a clearing house. Therefore, some kind of proxy in the access network which interacts with the clearing house on behalf of several end points may be required. Another approach is to use a distributed clearing house. If this model is deployed on an Internet-wide scale then there is a need for multiple clearinghouses that need to communicate with each other. This introduces additional complexity.

- A route change might require a new end-to-middle authentication/authorization for the purpose of charging. Hence a route change might not be handled locally anymore. This has an impact on the local repair mechanism. In the New Jersey Turnpike model a route change in the middle of the network does not require any interaction with nodes other than the involved ones. The New Jersey Parkway model is different since it might require an interaction with the end points and thereby destroying the local repair mechanism.

- To reduce state maintenance, processing and signaling message exchanges in the core network some sort of aggregation (see [6], [7], [8]) is used. Aggregation thereby causes per-flow end-to-end signaling messages to be hidden in the core network and a separate signaling message exchange to be used. Because the New Jersey Parkway model might require some interaction with an individual end host aggregation in the traditional sense might be much more difficult to deploy.

5 What should be charged?

In the above-description we assumed that data sender is simply charged for something. There are, however, some more fine-grained charging considerations which might affect the complexity of the interaction. In

Section 6 the question is consider which entity to charge. Closely related is therefore what a user can be charged for:

Signaling messages: Although it is possible to charge signaling message originators for generated messages it is currently rarely used. In some cases such a behavior denial of service attacks could be prevented or the misuse of end-to-end

signaling messages as a subliminal channel.

QoS reservations: Charging on the reservations implies charging on reserved resources regardless whether they are used or not.

Transmitted data traffic: Charging based on transmitted data traffic is based on the amount of bytes/packets that have been send by the data source. This type of charging will constrain the traffic volume of the data source but not the duration or amount of the reservation. Therefore, this type of charing can only be applied for QoS classes that allow for overbooking of resources (i.e. resources are not provisioned with regard to their specified peak rate).

Application cost: Finally there are costs associated to the usage of a particular service such as a video service. This cost might already include the cost of the above-mentioned costs for more end user transparency. Costs for applications are outside the scope of NSIS.

6 Who should be charged?

Which entity is charged is one of the most important set of components for a AAA framework. To provide the required functionality the following issues need to be addressed:

- Negotiation which entity is charged for which part of the costs

- Price Distribution

- Authorization of a reservation

- QoS Signaling

These individual steps might be combined and merged with the QoS signaling messages. As an example, in RSVP the signaling messages PATH or RESV might be used to piggyback information related to price distribution and charging. Whether this is possible depends on the flexibility of the signaling protocol, the number of round-trips executed by the signaling protocol and the semantic of the messages.

Subsequently the above-described issues are discussed in more detailed:

Negotiation which entity is charged:

First the end points need to negotiate/determine which end point will be charged for what part of the total cost. Once it has been decided the networks along the path (Parkway model) or the access networks have to be informed about this decision since they finally need to know where to get the money from. In existing telecommunication networks it is not only possible to provide this negotiation capability at the beginning of the session but also during an establish session or even afterwards. The term session in this context refers to a telephone call. Because of the stateless principle we assume that there is no such session concept and hence it is fair to say that the negotiation is done first (but with the option to be changed at any time).

In this context it is interesting to mention that ST-II [9] provides an object to indicate which entity to charge for the reservation. Such object is not included in the base RSVP RFCs. We believe that such information should belong to a QoS signaling protocol since it delivers the necessary information to the networks in order to setup the accounting and charging procedures.

In the literature (for example in [5], in [10] and in [11]) an additional degree of control has been introduced by allowing the sender and the receiver to share their costs according to a fractional part. Furthermore it is possible the the two parties share different types of costs (see Section 5). Hence it would be possible to charge the sender for the QoS reservation but to charge the receiver for application specific costs. It is needless to say that this adds additional complexity.

Price Distribution:

Aspects of price distribution are discussed in Section A but a summary of the most important issues is given in this section. Two problems arise when determining the price of the reservation. First, the price cannot be immediately referred from the IP address. Second, the asymmetry of routes at router and AS level (see [12]) and the possibility of asymmetric costs for a single link in the uplink or downlink direction requires that the direction is considered.

The process of price determination, price distribution and authorization/charging is likely to be periodic since the duration of the reservation is unknown. The soft-state principle used in NSIS requires periodic refresh messages to keep a reservation in place. Hence there is a question whether the price determination, price distribution and authorization/charging mechanism should be closely tight to this refresh interval. There is clearly a tradeoff between performance (computational and bandwidth requirements) and accuracy/efficiency. If price determination, price distribution and authorization/charging mechanism is bound to the refresh interval and the refresh messages are transmitted at a very high rate then substance overhead might be caused.

From a user perspective it seems to be important that cost transparency is provided and that the end host has the ability to determine the cost of a reservation and has the ability to perform cost control.

Authorization of a reservation:

Whenever authorization is discussed in this context then the ability to provide assurance for charging is meant. This is, however, only of interest where an end host is participating in the signaling message exchange and depending on the chosen model which part of the signaling path is considered. For intra-domain traffic (traffic within an administrative domain) authorization is much simpler: An incoming signaling message hitting a router within the domain is authenticated and verification is required to ensure that the message is transmitted from a known router within the same domain. This assumes that the borders are properly protected and discard unprotected signaling message from other domains.

Authorization for a QoS reservation from an end host to the attached network requires some guarantee that a particular entity can be charged for the cost of the reservation. The charged entity is likely to be one of the end hosts. The establishment of the necessary infrastructure is either based on a per-session communication (e.g. micro-macro payment protocols, authorization tokens) or more traditionally as part of the network access procedure (e.g. AAA communication). Depending on the model (NJ Turnpike or NJ Parkway model) and

on the choice for charging of the data sender or the data receiver per-session established authorization setup might be required. From a performance perspective the per-session based approach is less favorable.

QoS Signaling:

Finally there is the question how the above-described steps should be most efficiently combined with the behaviour of a QoS signaling protocol.

Principally either the data sender or the data receiver can be charged for a QoS reservation. Since signaling protocols are typically characterised as either sender- or receiver-initiated, an answer has to be provided which approach allows a better integration with various charging strategies. Unfortunately, it is not possible to consider only charging for the data sender since charging for the data receiver is often used in todays telecommunication networks (e.g. 800# numbers, collect calls).

In this version of the document we mainly focus on the simpler NJ Turnpike model. Future versions will extend the descriptions to the NJ Parkway model.

To simplify representation the AAA infrastructure is not shown in Figure 3, 4, 5 and in 6. Hence to get a complete picture the reader has to take the AAA infrastructure into account. This might involve interaction with local AAA servers, interaction with a Credit Control Server for the purpose of real-time cost and credit control as described in [13] or home AAA servers in case of mobility as depicted in Section 7.

6.1 Sender initiated reservations with charging for the data sender

This model is the simplest in relationship with the NJ Turnpike model since the data sender S which triggers the reservation is also charged. The necessary charging infrastructure is likely to be established as part of network access authentication and the interaction with a AAA infrastructure. When AS 1 receives a QoS reservation request which asks for the establishment of a QoS reservation then an authorization check can immediately be executed. Authorization might not only be based on credit availablility but also on some information stored with the

subscriber's contract such as contract type or some form of policy which might also be distributed as part of the initial network access procedures or on-demand accessible via the AAA infrastructure. The subscriber's contract is a business relationship between the user and his home provider.

To provide cost control for the data sender it is likely that additional communication between AS 1 and the initiator S is necessary to distribute the necessary information. The initiator S might want to know the price for the QoS reservation in advance before issuing a QoS reservation message (RESV message in Figure 3 based on the RSVP terminology). Hence for in-band price distribution a separate roundtrip

is required. For out-of-band price determination such a roundtrip can be omitted but some tariff or price information has to be communicated between the sender S and the access network (AS1 in our case) - if not already known for some other reason.

For in-band price distribution each network (or even each router) along the path accumulates cost and AS 1 charges S for the total amount. Based on the existing peering relationship between neighboring networks each provider charges its neighboring provider. This procedure might be comparable with the postal service where a customer gives a letter to a post post office and delegates responsibility to perform the required shipping. The post office might itself delegate the responsibility to other companies to transport the letter to its final destination. The sender pays for the total amount for the shipment at the post office which knows the total cost for the entire delivery. Each participating party receives the monetary amount negotiated with its "peer" based on the previously agreed price. A similar description is provided in [14].

```
+-----+ RESV +-----+ RESV +-----+ RESV +-----+ RESV +-----+
|AS 1|----->|AS 2|----->|AS 3|----->|AS 4|----->|AS 5|
+-----+      +-----+      +-----+      +-----+      +-----+
  ^                                                   |RESV
  |RESV                                               v
+-----+                                             +-----+
| S  |                                              | R  |
+-----+                                             +-----+
                     Data Traffic
```

```
        ==================================================>
          Charging (S -> AS1 -> AS2 -> AS3 -> AS4 -> AS5)
        ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>

  Legend:

  ----> Signaling message
  ====> Data flow
  ~~~~> Charging direction
```

Figure 3: Sender-initiated reservation with charging for the data sender

6.2 Sender initiated reservation with charging for the data receiver

Charging for the data receiver is more complex in comparison to charging for the data sender. The reason is not due to the QoS signaling machinery - such as sender- or receiver-initiated reservations but caused by the complicated charging relationships. The following description tries to describe the problem in more detail which is depicted in Figure 4:

When AS 1 receives the RESV signaling message from S which indicates that R is charged for the price of the QoS reservation then AS 1 needs some assurance that the entity R is willing to pay for the indicated reservation. Hence a plain identifier containing the identity of R is insufficient to provide enough assurance.

Hence the sender needs to possess some from of authorization token which allows AS 1 to establish the necessary association to a party which is able to provide the financial settlement. Following the idea of such an authorization token the subsequently described interaction is necessary.

An authorization token previously send from R to S and then transmitted to AS 1 might allow AS 1 to establish the necessary infrastructure (possibly to a trusted third party or to R's home network) to execute a real-time credit check and to be able to charge R via this

infrastructure by AS 1 for a given QoS reservation. Then the QoS reservation is done in the same way as a sender-initiated reservation with charging for a data sender. The total cost for the session cannot be fully determined during the reservation setup because the duration of a call and other factors are unknown at the beginning. Hence periodic communication is necessary between AS 1 and a trusted third party or R's home network. R needs to be given a mechanism to allow the QoS reservation and therefore the costs to be restricted without always transmitting authorization tokens to the data sender for periodic re-authentication and re-authorization procedures.

Note that the sender S communicate the name of the data senders access network (in this case AS 1) to the receiver R. This allows the data receiver R to request an authorization token for a specific network with the indicated QoS parameters to include some additional restrictions in the token.

It is not very likely that the data receiver R provides direct payment to S before triggering a QoS reservation. Such an infrastructure is not likely to be available.

6.3 Receiver initiated reservation with charging for the data receiver

The properties of the sender initiated reservation with charging for the data receiver described-above are similar to those of a receiver


```
+-----+ RESV +-----+ RESV +-----+ RESV +-----+ RESV +-----+
|AS 1|----->|AS 2|----->|AS 3|----->|AS 4|----->|AS 5|
+-----+      +-----+      +-----+      +-----+      +-----+
   ^                                                  |RESV
   |RESV                                              v
+-----+                                            +-----+
| S   |                                            | R   |
+-----+                                            +-----+
                       Data Traffic
   =================================================>
                  Payment (R -> AS1 or R -> S)
   <~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
     Charging ([S ->] AS1 -> AS2 -> AS3 -> AS4 -> AS5)
```

```
        ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>

  Legend:

  ----> Signaling message
  ====> Data flow
  ~~~~> Charging direction
```

Figure 4: Sender-initiated reservation with charging for the data receiver

initiated reservation.

When AS 1 receives a PATH signaling message then S has to indicate that R is willing to pay for the QoS reservation. Unfortunately the PATH message (with the semantics defined within RSVP) cannot be used to determine the price of a reservation since the receiver is allowed to change the QoS parameters. Hence the computed price might only serve to compute the upper-bound and would therefore only serve R as a hint. AS 5 cannot use an out-of-band price distribution mechanism because of asymmetric routes. Hence price distribution can only be probing based (in-band).

Finally after a successful reservation the receiver R (or some party associated with R) has to provide a financial settlement with AS 1 to transfer the desired QoS costs.

A major question is therefore whether it is possible for R to provide financial settlement with AS5 although the reservation price is

determined from S to R (data flow direction).

AS 1 therefore has to determine the price for the reservation and communicate the accumulated price along the path to AS 5 and to R.

TBD: Is it possible for R establish a financial settlement with AS5 to provide peer-to-peer charging in the reverse direction(R -> AS 5 -> AS 4 -> AS 3 -> AS 2 -> AS 1) although authorization for the RESV message

would be required at AS 1?

```
  +-----+ PATH +-----+ PATH +-----+ PATH +-----+ PATH +-----+
  |AS1  |.....>|AS2  |.....>|AS3  |.....>|AS4  |.....>|AS5  |
  |     |<-----|     |<-----|     |<-----|     |<-----|     |
  +-----+ RESV +-----+ RESV +-----+ RESV +-----+ RESV +-----+
    ^ |                                            .  ^RESV
PATH. v RESV                                  PATH v  |
  +-----+                                         +-----+
  |  S  |                                         |  R  |
  +-----+                                         +-----+
                       Data Traffic
     =================================================>
                  Charging (R -> AS1 or R -> S)
     <~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~
       Charging ([S ->] AS1 -> AS2 -> AS3 -> AS4 -> AS5)
     ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>

  Legend:

  ----> Signaling message
  ====> Data flow
  ~~~~> Charging direction
```

Figure 5: Receiver-initiated reservation with charging for the data receiver

6.4 Receiver initiated reservation with charging for the data sender

When the sender S transmits a PATH message neither S nor AS 1 are able to determine the cost for the reservation solely based on the semantic of the PATH message. The PATH message is forwarded towards the data receiver. R then finally decides about the reservation and its

parameters but S is charged for the reservation.

It seems to be difficult for the sender S to restrict the QoS parameters selected by the receiver R when transmitting the RESV message. It would therefore be better if either a double roundtrip is used or if the semantics of the PATH message is changed.

```
  +-----+ PATH +-----+ PATH +-----+ PATH +-----+ PATH +-----+
  |AS1 |.....>|AS2 |.....>|AS3 |.....>|AS4 |.....>|AS5 |
  |    |<-----|    |<-----|    |<-----|    |<-----|    |
  +-----+ RESV +-----+ RESV +-----+ RESV +-----+ RESV +-----+
    ^ |                                              .  ^RESV
PATH. v RESV                                   PATH v  |
  +-----+                                           +-----+
  | S  |                                            | R  |
  +-----+                                           +-----+
                       Data Traffic
     =================================================>
       Charging (S -> AS1 -> AS2 -> AS3 -> AS4 -> AS5)
     ~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~>

  Legend:

  ----> Signaling message
  ====> Data flow
  ~~~~> Charging direction
```

Figure 6: Receiver-initiated reservation with charging for the data sender

6.5 NJ Parkway Model Example

The following example shows the implications for a sender initiated reservations with charging for the data sender based on the NJ Parkway model.

The sender needs some mechanisms to provide information to all intermediate domains which request independent charging from the data sender (i.e. from S). This mechanism can be provided by the following procedures:

- Information carried within the NSIS protocol (e.g. OSP tokens) which immediately allow the intermediate domain to contact some trusted third party (such as a clearing house).

- The possibility for an intermediate network to request authentication / authorization from the data sender S via NSIS. Such a mechanism might therefore be similar to SIP.

- An out-of-band mechanism which is triggered by intermediate networks to request authentication and authorization from intermediate networks.

In-band price distribution (or probing) is difficult to use since the data sender is not aware of the QoS reservation costs along the entire path (without a previous query). Out-of-band price distribution might provide this functionality but a separate interaction with each domain to the end host is required. When transmitting some sort of authorization tokens it might be useful for the data sender S to have information about the QoS reservation costs of all individual intermediate domains along the path.

7 Implication of Mobility

This section addresses some of the implications of mobility. Starting with a simple model at the beginning which allows limited mobility in the same administrative domain some basic observations are made. Extending the basic model to support to mobility to support mobility where both users are registered at the same home network but roam to different access networks (different from the home network). Finally even this restriction is abolished.

7.1 Simple model without mobility

In Figure 7 two nodes are attached to a single administrative domain either in a non-mobile environment (traditional enterprise network) or with limited mobility in this network. No roaming agreements are necessary and even authentication during network access might be simplified due to a larger degree of freedom for selecting the proper security infrastructure (for example Kerberos everywhere). To provide authorization of a QoS reservation request role based access control might be used since momentary authorization might not be applicable in an enterprise network. Instead users or groups with specific rights might be allowed to trigger QoS reservations. In this case it might not even be necessary to communicate information who is charged for which information to the network elements. Inter-domain communication for QoS signaling messages and for AAA communication is not required.

```
   +-------------------------------------------------------------------+
   |                                                                   |
   |                            +-------+                Network       |
   |                           -+ AAA/  +--                 X          |
   |                        --- | PDP   |  ----                        |
   |                     ---    +-------+      -----                   |
   |                  ---                           -----              |
   |               ---                                   ----          |
   |          -----                                           ---      |
   |      +---+----+                                       +---+----+  |
   |      | Router |                                       | Router |  |
   +------|   1    |---------------------------------------|   n    |--+
          +---+----+                                       +---+----+
              |                                                |
           +--+---+                                        +--+---+
           | Node |                                        | Node |
           |  A   |                                        |  B   |
           +------+                                        +------+
```

Figure 7: Simple model without mobility

7.2 Split between access and home network(s)

With Figure 8 the basic environment described in Figure 7 is extended by allowing end hosts to roam to networks (denoted as access network) beyond their home networks. As part of the network access authentication the end host is authenticated to its home network involving entities (such as the local AAA server in the access network). AAA inter-domain communication is required. The QoS signaling messages stay within the same access network which is different than the home network. Additionally the user might be registered at different home networks. These networks primarily serve the purpose of providing a guarantee that the indicated user requesting resources (and network access) is able to pay. This functionality can be provided by a traditional telecommunication network, by a credit card company or by something

similar.

In comparison to the previous model it is likely that role-based access control is not sufficient for the purpose of QoS reservation request authorization. Hence it might be necessary for the end hosts to decide which entity (user A at node A or user B at node B) has to be charged for which resource (QoS reservation, QoS data traffic, etc.). The access network then collects accounting records and transmits bills to the indicated home network of the authenticated user. Since the QoS

signaling messages travel only within a single administrative domain it is not necessary to address issues raised in Section 4.

```
+----------------------+                      +----------------------+
| +-------+            |                      | +-------+            |
| | AAA   |    Home    |                      | | AAA   |    Home    |
| |       |    Network |                      | |       |    Network |
| +--+----+    (User A)|                      | +--+----+    (User B)|
|    |                 |                      |    |                 |
|    |                 |                      |    |                 |
+----+-----------------+                      +----+-----------------+
     |                                             |
     +---------------------------+-----------------+
                                 |
+--------------------------------+-----------------------------------+
|                             +--+----+               Access Network |
|                             | AAA/  |                              |
|                             | PDP   |                              |
|                             +--+----+                              |
|          +---------------------+-----------------------+           |
|          |                                             |           |
|      +---+-----+                                   +---+-----+     |
|      | Router  |                                   | Router  |     |
+------|   x     |-----------------------------------|   y     |-----+
       +---+-----+                                   +---+-----+
           |                                             |
        +--+----+                                     +--+----+
        | Node  |                                     | Node  |
        |  A    |                                     |  B    |
        +-------+                                     +-------+
```

Figure 8: Split between access and home network(s)

7.3 Global mobility

As an extension of the previous model global mobility is considered where users are subscribed at different home networks and they roam in different networks. The networks between the two access networks (X and Y), which are traversed by the QoS signaling message, are omitted. This scenario addresses issues discussed in Section 4 and 6. Note that the AAA Broker is not necessarily required if the two home networks (of user

A and B) share a business and trust relationship (and consequently a security association).

```
                             +----------+
                             |   AAA    |
     +-----------------------+  Broker  +----------+
     |                       +----------+          |
+----+-----------------+                      +----+-----------------+
| +--+---+             |                      | +--+---+             |
| | AAA  |     Home    |                      | | AAA  |     Home    |
| |      |     Network |                      | |      |     Network |
| +--+---+     (User A)|                      | +--+---+     (User B)|
|    |                 |                      |    |                 |
|    +----+            |                      |    +----+            |
|         |            |                      |         |            |
+---------+------------+                      +---------+------------+
          |                                             |
+---------+------------+                      +---------+------------+
|         |            |                      |         |            |
|      +--+----+       |                      |      +--+----+       |
|      | AAA/  |       |                      |      | AAA/  |       |
|      | PDP   |       |                      |      | PDP   |       |
|      +---+---+       |                      |      +---+---+       |
|          |           |                      |          |           |
```

```
|   Access Network X    |                       |   Access Network Y    |
|           |           |                       |           |           |
|       +---+-----+     |                       |       +---+-----+     |
|       | Router |      |                       |       | Router |      |
+-------|   x    |------+                       +-------|   y    |------+
        +---+-----+                                     +---+-----+
            |                                               |
         +--+----+                                       +--+----+
         | Node  |                                       | Node  |
         |  A    |                                       |  B    |
         +-------+                                       +-------+
```

Figure 9: Global mobility

8 Security Considerations

This document describes the implications of two accounting and charging models (i.e. the New Jersey Turnpike and the New Jersey Parkway model) for NSIS QoS signaling. As excepted, there are implications for the

security architecture. The New Jersey Turnpike model is based on the peer-to-peer security and the chain-of-trust. This model, although often criticised, serves as the basis for RSVP and some of its mechanisms such as local repair and the aggregation mechanism. The second model, the New Jersey Parkway model, relaxes the assumption of the first model. The introduced end-to-middle authentication adds additional complexity.

This document does not discuss concrete security mechanisms for both models, instead the implications are presented at an abstract level. Hence it is not useful to give detailed security requirements and threats.

Based on the topics discussed in this draft the NSIS working group should decide on which model QoS signaling should be based. Additionally it is necessary to discuss sender- and receiver-initiated signaling and finally the impacts of price distribution need to be addressed.

9 Open Issues

* Non-repudiation is a security property where one party is later unable to deny the execution of a specific action. For QoS signaling this might be a desirable property. When added to a signaling protocol this property, unfortunately, is not for free. Hence it is an open question whether real-world applications and architectures demand this property.

10 Acknowledgements

Place you name here.

A Price Distribution

How much an entity is charged for individual parts of a QoS reservation (see Section 5) is mainly a matter of business/marketing decisions and will not be discussed in this document and is outside the scope of NSIS. The task of determining the price is called pricing. Unfortunately the price of a QoS reservation cannot easily determined based on the structure of the IP address unlike with E.164 phone numbers. Depending on the chosen price distribution mechanism implications for an NSIS signaling protocol exist.

Principally, two ways of price distributions can be identified:

Out-of-band price distribution: Using this approach the inter-domain prices for certain destinations a distributed by mechanisms executed separate from a NSIS in-path signaling protocol. In case of out-of-band price distribution it is required that the price is determined based on destination AS

and the ASes along the path to this network. If the price for one or more networks along this path then some additional signaling is required. The main assumption of this scheme is that the information obtained by the BGP-based sink tree mechanism provides a good approximation to the path subsequently taken by the later data packets.

In-band probing: The signaling messages enable some functions to query the costs along the path to determine the costs between the source and the destination. To discover the networks along the path is fairly simple if a signaling protocol used used

(in-band probing). As a disadvantage a signaling protocol needs to carry new objects and additional processing is required at each network along the path. Hence it is required that each network understands and processes these objects.

In the past a number of price distribution protocols have been proposed which have a strong relationship to the signaling machinery, since they share common properties:

- The determined price depends on the route between source network and destination network. Protocols which allow their objects to be embedded into the signaling protocol (in-band probing) are able to more accurately determine the path and therefore the associated costs.

- Some flexibility and extensibility is required to allow autonomous systems to determine the price for a QoS reservation independently of other domains.

- Signaling price information between various networks suffers from the same signaling protocol requires (such as scalability, "in-path" discovery, etc.) as QoS signaling protocols. To tackle scalability similar mechanisms for aggregation are therefore considered such as those used in [7].

- Unfortunately none of the proposals cares about the issues described in Section 4 introduced by the two different models.

In [11] an in-band probing approach is presented which allows price information to be communicated. The pricing object is updated along the path to reflect costs. The idea of the Simple RSVP protocol [15] also seems to follow a similar strategy.

RNAP [10] is a proposal which allows both in-band probing and out-of-band price distribution. The out-of-band price distribution mechanism is modeled according to the same principles as BGRP's aggregation and protocol mechanism [7]. The aggregation mechanism of BGRP is inspired by

BGP [16]. A very similar idea was chosen by the Border Pricing Protocol (BPP) [17], which uses the same aggregation mechanism but only allows out-of-band price distribution.

The Tariff Distribution Protocol (TDP) [18] is an attempt to define message formats (using XML, HTML, plain text or even in a binary format e.g. JAVA classes) and attributes for exchanging tariff information either in an in-band (for example with RSVP) or out-of-band fashion. Instead of exchanging price information in [18] tariff are communicated. The term tariff is thereby defined as: "A tariff is a set of rules for calculating the charge advices for session of one service" (see Section 2 of[18]). The difference between charge and charge advice is also described in Section 2 of [18]. Unlike in other proposals aggregation is not considered.

In the Billing Information Protocol (BIP) [19] only attributes used to deliver price information are described (in BNF notation). The current specification mainly addresses SIP as a transport mechanism but can be used for other protocols as well.

Related to the work described above is the Open Settlement Protocol [20] which is mainly focused on charging and not on price distribution. Hence it should be seen as complementary to the above schemes which could be used to support the New Jersey Parkway Model described in Section 4.

The work in the Internet Open Trading Protocol (IOTP) working group (see [21] for the IOTP Version 1.0 specification) aims to map real world transactions to the internet and is as such a superset of the functionality described in this document.

B Authors' Addresses

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

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

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

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

C Bibliography

[1] 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.

[2] B. Aboba, P. Calhoun, S. Glass, T. Hiller, P. McCann, H. Shiino, G. Zorn, G. Dommety, C. Perkins, B. Patil, D. Mitton, S. Manning, M. Beadles, P. Walsh, X. Chen, S. Sivalingham, A. Hameed, M. Munson, S. Jacobs, B. Lim, B. Hirschman, R. Hsu, Y. Xu, E. Campbell, S. Baba, and E. Jaques, "Criteria for evaluating AAA protocols for network access," RFC 2989, Internet Engineering Task Force, Nov. 2000.

[3] R. Hancock, I. Freytsis, G. Karagiannis, J. Loughney, and S. V. den Bosch, "Next steps in signaling: Framework," Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[4] R. Yavatkar, D. Pendarakis, and R. Guerin, "A framework for policy-based admission control," RFC 2753, Internet Engineering Task Force, Jan. 2000.

[5] S. Shenker, D. Clark, D. Estrin, and S. Herzog, "Pricing in computer networks: Reshaping the research agenda," in Proc. of TPRC 1995 , 1995.

[6] F. Baker, C. Iturralde, F. L. Faucheur, and B. Davie, "Aggregation of RSVP for IPv4 and IPv6 reservations," RFC 3175, Internet Engineering Task Force, Sept. 2001.

[7] P. Pan, E. Hahne, and H. Schulzrinne, "Bgrp: Sink-tree-based aggregation for inter-domain reservations," in Journal of Communications and Networks, Vol. 2, No. 2, pp. 157-167, http://www.cs.columbia.edu/

pingpan/papers/bgrp.pdf , 2000.

[8] Y. Bernet, P. Ford, R. Yavatkar, F. Baker, L. Zhang, M. Speer, R. Braden, B. Davie, J. Wroclawski, and E. Felstaine, "A framework for integrated services operation over diffserv networks," RFC 2998, Internet Engineering Task Force, Nov. 2000.

[9] C. Topolcic, "Experimental internet stream protocol: Version 2 (ST-II)," RFC 1190, Internet Engineering Task Force, Oct. 1990.

[10] X. Wang and H. Schulzrinne, "Rnap: A resource negotiation and pricing protocol," in International Workshop on Network and Operating Systems Support for Digital Audio and Video (NOSSDAV'99), pages 77--93, Basking Ridge, New Jersey

[11] M. Karsten, J. Schmitt, and R. Steinmetz, "An embedded charging approach for rsvp," in International Workshop on Quality of Service '98, Napa, California, USA , May 1998.

[12] L. Amini and H. Schulzrinne, "Observations from router-level internet traces," in DIMACS Workshop on Internet and WWW Measurement, Mapping and Modeling, (Piscataway, New Jersey) , Feb. 2002.

[13] H. Hakala et al. , "Diameter credit control application," Internet Draft, Internet Engineering Task Force, June 2002. Work in progress.

[14] J. Gerke, H. Ritter, J. Schiller, and K. Wehrle, "Elements of an open framework for pricing in the future internet," in Proceedings of the Conference on Quality of future Internet Services (QofIS 2000), pages 300--311, Berlin

[15] K. Fujikawa and Y. Goto, "Simple resource ReSerVation protocol (SRSVP)," Internet Draft, Internet Engineering Task Force, Feb. 2000. Work in progress.

[16] Y. Rekhter and T. Li, "A border gateway protocol 4 (BGP-4)," RFC 1771, Internet Engineering Task Force, Mar. 1995.

[17] V. Oberle, H. Ritter, and K. Wehrle, "Bpp: A protocol for exchanging pricing information between autonomous systems," in Proceedings of HPSR 2001 (IEEE Workshop on High-Performance Switching and Routing), Dallas (USA) , May 2001.

[18] O. Heckmann et al. , "Tariff distribution protocol (TDP)," Internet Draft, Internet Engineering Task Force, Mar. 2002. Work in progress.

[19] R. Prasanna, "Bip: Billing information protocol," Internet Draft, Internet Engineering Task Force, 2002. Work in progress.

[20] 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."

[21] D. Burdett, "Internet open trading protocol - IOTP version 1.0," RFC 2801, Internet Engineering Task Force, Apr. 2000.