Hybrid on-path off-path approach for end-to-end signalling across NSIS and non-NSIS domains (HyPath)

Luís Cordeiro[1], Marilia Curado[1], Edmundo Monteiro[1],
Florin Racaru[2], Michel Diaz[2], Christophe Chassot[2]

[1]Dept. of Informatics Engineering, Univ. of Coimbra,
Polo II - Pinhal de Marrocos, 3030-290 Coimbra, Portugal
{cordeiro, marilia, edmundo}@dei.uc.pt
[2]Laboratory for Analysis and Architecture of Systems,
Avenue du Colonel Roche, 31077 Toulouse, France
{fracaru, diaz, chassot}@laas.fr

Abstract. In a multi-domain Internet that offers QoS guaranties for applications, there is the need of signalling among the domain entities that are responsible for the management of QoS. Because different domains have different network protocols and topologies, there is no solution that is able to signal these entities using an off path approach, and in particular that is able to interwork with the on path. The HyPath approach uses the NSIS protocol and interactions with the local routing protocols to have an off path signalling in these hybrid environments.

Keywords. HyPath, off path, signalling, NSIS, QoS, hybrid.

1. Introduction

In the past years, we assisted at a common rise of new technologies in the telecommunication and computer science fields. This evolution led to the emergence of new types of applications involving multimedia, like VoIP, VoD, tele-engineering and telemedicine. These applications have new constraints and requirements concerning parameters such as delay and jitter. Therefore, new services are required besides those given by the actual Internet.

Nowadays, all packets in the Internet receive the same treatment. As presented before, some data flows need special processing in order to satisfy the application requirements, and thus it is necessary to address Quality of Service (QoS) issues. The internet is an interconnection of networks, comprising different domains, called Autonomous Systems (AS), managed independently, especially in what concerns QoS strategies. In order to support QoS for communications over several domains, intra and inter-domain QoS signalling appears to be inevitable.

Our work aims at defining a global architecture able to provide a QoS signalling within a of a multi-domain Internet context that offers QoS guaranties for applications. Inside the domains, the QoS is managed through central entities, that are in charge of installing and handling QoS based on internal rules. This concept was introduced in the Differentiated Services (DiffServ) domains, and is associated with Bandwidth Brokers [1]. At the present, a new requirement appears: signalling must take place, not only among devices strictly on the data path, but also among new domain-optimised central entities, that we call hereafter Resource Manager (RM).

Several signalling protocols have been proposed, especially in the IETF NSIS working group [2]. The goal of the NSIS protocol is to manipulate the network state related to data flows with the constraint that the signalling protocol will be processed on the nodes which also handle the data flows themselves ("path-coupled signalling"). This paper discusses a NSIS multi-domain, multi-service, RM based Internet that allows also includes an off-path signalling. The main issue addressed is the inter-operability between NSIS and non-NSIS domains.

This paper is structured as follows: Section 2 presents the state of the art of off path signalling protocols; Section 3 describes the NSIS protocol and the off path signalling problem, followed by the HyPath approach to solve the off path problem; Section 4 presents and describes some HyPath use cases; Finally, Section 5 comprises the main conclusions of the paper.

2. Off path signalling state of the art

This section addresses off path signalling issues focusing in the description of the state of the art and the main current proposals.

The expression “off path” is used to convey the situation where entities participating in the signalling process are not bound to the path followed by the data flows. The most common example is when particular entities inside a domain, which have special responsibilities such as QoS, policy control and servers must be signalled. These devices are not strictly on the data path; nevertheless the signalling protocol must arrive to interact with these devices. Off-path signalling has advantages, as presented in [3] and [4], on the other hand, off-path signalling must answer new challenges such as discovering the next hop and synchronisation with Interior Gateway Protocol (IGP) and Exterior Gateway Protocol (EGP) routing protocols. Several protocols have been proposed for the off-path signalling in a bandwidth broker-based multi-domain DiffServ model, as described in the remaining of this section.

2.1 SIBBS (Simple Inter-domain Bandwidth Broker Protocol)

The SIBBS protocol [5] has been defined by the QBone Signalling Workgroup and it aims to be used on DiffServ bandwidth broker-based domains. In the QBone testbed, each network is a DiffServ domain supporting one or more globally well known forwarding services built from fundamental DiffServ blocks. SIBBS is a very simple protocol to be used between bandwidth brokers. It contains two principal PDU: RAR (Resource Allocation Request); RAA (Resource Allocation Answer).

The RAR message includes a globally well-known service ID, information related to the QoS request (class of services and bandwidth) and a destination IP address, a source IP address, an authentication field, and the other parameters of the service. The sender can be the client host, a BB or a proxy. The RAA message contains the answer to a RAR PDU. Receiving a RAR message, a BB:

- Authenticates that the request is indeed from a peer bandwidth broker;
- Determines the egress router (interface) from its (inter-domain) routing tables;
- Checks that the requested resources fall within the SLS;
- Ensures that there are sufficient resources within the domain to support the flow from the ingress border router;
- Determines whether the flow may be accepted according to the policies of the domain.

If the required resources are available, the request is propagated successively through the inter-domain path to the last BB. Afterwards, this last BB returns a RAA message to its immediately upstream BB and the process is continued until the originating BB. This process is concluded with an admission of the QoS request. Resources are confirmed by means of refresh messages, sent periodically.

2.2 COPS-SLS

COPS-SLS [6, 7] is an extension of the Common Open Policy Service (COPS) protocol [8] for SLS management in a multi-domain environment. COPS is a client/server protocol designed for the management of policy based networks. The basic model of COPS is presented below:

- The PDP (Policy Decision Point) is the central entity in charge of making the decisions (for itself or for other elements of the network). The PEP (Policy Enforcement Point) is the point where the policies are applied, such as a router. The optional Local Policy Decision Point (LPDP) can be used by the device to make local policy decisions in the absence of a PDP;
- COPS is a request/response protocol that allows a PEP (router) to interrogate its PDP about the action to perform once an event has occurred (for instance, if a signalling message arrived);
- COPS-PR is an extension of COPS with the goal to force the application of a policy in the PEP without any prior request.

COPS-SLS has the same behaviour as SIBBS: a request is propagated from one BB to the other in each domain of the data path. Each BB has a double role, namely, as a PDP for the upstream domain, BB which sends the request, and as a PEP for the next BB domain.

Compared to SIBBS, COPS-SLS adds some features to the protocol, as the renegotiation of classes of service in case of failure of admission control. The communication between BB and border routers is assured by the COPS-PR protocol. However, COPS-SLS does not provide any specification on the discovery of the next BB or on the identification of border routers.

3. HyPath

The HyPath approach to the end-to-end signalling across NSIS and non-NSIS domains is described in this section. We start by describing the signalling protocol

NSIS followed by the description of the off path signalling problem and the need for a solution. Afterwards, the HyPath approach is described.

3.1 NSIS

Next Steps in Signalling (NSIS) [2] is a new protocol being developed in the Next Steps in Signalling Working Group of The Internet Engineering Task Force (IETF). This group is responsible for standardizing an IP signalling protocol following a two-levels signalling paradigm with QoS signalling as the first use case.

In the development of this protocol there is the intention to re-use the Resource ReSerVation Protocol (RSVP) [9] mechanisms, whenever this is possible since these mechanisms have already been widely tested.

Figure 1 shows the simplest signalling configuration using NSIS. A single data flow is running from an application in the sender to the receiver via routers R1, R2, and R3. Each host and two of the routers contain NSIS Entities (NE) that exchange signalling messages about the flow.

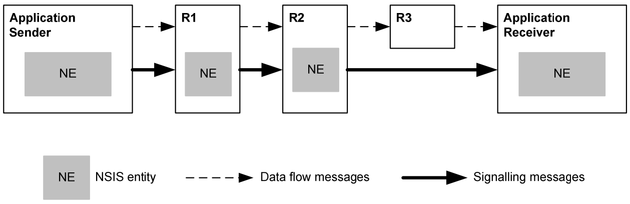

Figure 1 – Simple NSIS configuration

As stated before, NSIS is being developed as a two-layer modular solution, comprising the signalling transport layer and the signalling application layer. With this approach the transport of the signalling messages and the signalling application are separate, which allows the protocol to be used for more general purposes.

The signalling transport layer, GIST, is responsible for moving signalling messages among network entities. This process should be independent of the signalling applications. The signalling application layer contains the specific functionalities of the signalling applications. This two-layer protocol model allows the support of various signalling applications such as QoS and NAT & Firewall, as described in Figure 2.

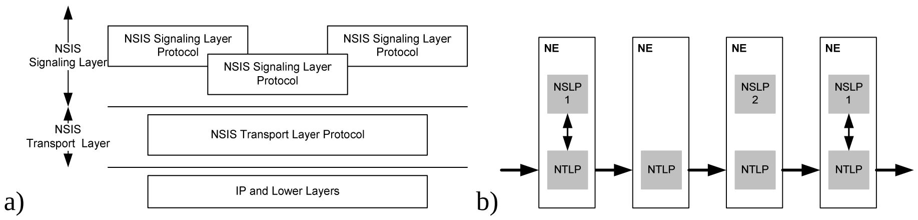
NSIS Signaling Layer
NSIS Signaling Layer Protocol
NSIS Signaling Layer Protocol
NSIS Signaling Layer Protocol
NSIS Transport Layer
NSIS Transport Layer Protocol
IP and Lower Layers
a)
b)
NE
NE
NE
NE
NSLP 1
NSLP 2
NSLP 1
NTLP
NTLP
NTLP
NTLP

Figure 2 – Two-layer protocol model

3.2 The off path signalling solution, HyPath

The requirements for a hybrid on-path/off-path approach for end-to-end signalling across NSIS and non-NSIS domains are not fully solved by the NSIS protocol as it is being defined currently in the IETF NSIS working group. There is the need to have network signalling between specific entities in domains (not only the routers in the data path like the normal on-path solution). This is the case of QoS network signalling when there are resource manager entities in the domains responsible for the domain QoS. In these situations the entities to be signalled are the RM entities and not only the network equipment (routers).

The normal way of work of the NSIS protocol, not only does not signal the RM servers in the data path, but also does not force the signalling to follow the same path as the user data (because the source and destination are different and the domains can have different routing policies based on local source IP addresses). Therefore the resource reservation will not be properly done on the data path.

The major requirements to achieve end-to-end network signalling are the following:

- Signalling messages must follow the same path as the user data;
- All the RMs in all ASs of the data path must be signalled.

The NSIS protocol as it is being defined in the IETF can not solve these two major requirements simultaneously. In order to fulfil the above requirements, a middle layer between the two NSIS layers was conceived. This layer is named HyPath (Hybrid Path).

To be able to connect the HyPath with the NTLP layer [10] and the NSLP layer [11] without altering their specifications, the HyPath needs to be a middle layer between the NTLP layer and the NSLP layer (the already defined interfaces must not be changed). Therefore, the HyPath interface with the NTLP layer must be the same as the NSLP layer interface already defined. Likewise the interface with the NSLP layer must be the same as the presently defined NTLP layer interface.

The operation of NSIS with the additional HyPath in the border routers and RMs in all domains is illustrated in Figure 3.

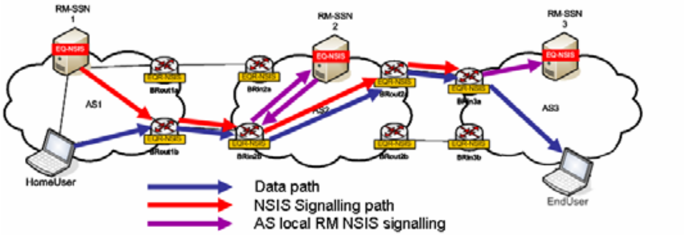

Figure 3 - NSIS signalling with HyPath

When a user makes a QoS request to the local QoS system, NSIS signalling must occur in order to signal all on path RM. This signalling must follow the same path as the data. Therefore, in the first domain, the HyPath in the local RM uses an external function (described latter in this section) to discover the local egress border router of the data.

Afterwards, the HyPath asks the NSIS transport layer to send a NSIS message to the egress border router. This message contains the NSLP payload and some additional HyPath information (described latter in this section). Once in the egress border router, the NSIS signalling message is sent to the end user.

In this scenario, all border routers intercept NSIS messages, and are HyPath aware. Therefore, in the following domain the NSIS signalling message is intercepted by the ingress border router. In this router the message is redirected to the local RM server to make the local RM signalling.

After processing the received message, the RM server continues the signalling sending a message back to the ingress border router. The signalling is restarted in the ingress border router and the NSIS message continues to the next domain.

These procedures continue in all domains until the last domain is reached and the signalling stops in the RM server.

With this architecture all the requirements to achieve end-to-end network signalling are met and no changes are needed in the definitions of the NTLP and NSLP layers. Specifically, in the first domain egress border router the data path and the signalling path meet. From that point on, if the NSIS signalling message is always sent to the end user, the message will follow the data path (the routing roles will be the same).

Non-NSIS domains. The drawback of the approach described is that all border routers of all domains must be NSIS aware. Even though in theory this is a reasonable assumption, in practice we can not guarantee that this happens. For this reason we define a heterogeneous solution that works when border routers are not NSIS aware (non-NSIS domains) and the only information available is provided by the routing protocol of the domain. Not being able to rely on NSIS interception in the border router, the solution is to rely on the routing protocol.

In non-NSIS domains, when the RM intends to send a signalling message, the HyPath uses an external function (described latter in this section) to discover the local egress border router of the data path and the next RM IP address. With this information, a NSIS message with the NSLP payload and some additional HyPath information (described latter in this section) is sent directly to the RM of the next domain.

Using again the external function to discover the local egress border router of the data path and the next RM IP address, the NSIS signalling message is sent to the RM of the next domain. The procedure described is repeated until the last domain is reached.

In this approach, the signalling messages do not follow the data path, but they follow all the RMs in the data path.

The inconvenience of this approach is the extensive usage of the external functions. Since these functions are used in all non-NSIS domains, this approach would have an impact on the processing time and on the amount of resources used.

Usage of external routing protocols. If the signalling is decoupled from the data path (but still path-related signalling) two general problems need to be solved:

- the RM must discover the ingress and egress points through which the data path will pass in its domain; this information is needed in order to continue the NSIS signalling and to perform an admission control between the ingress and the egress border routers and on the inter-domain link;

- in non-NSIS domains, the RM in the next domain must be identified in order to propagate the request.

The RM is supposed to have an access to the BGP tables of the border routers of its domain, and to be able to interrogate the BGP tables. This interrogation may be implemented as a request/response protocol via telnet or ssh. The main information in the BGP routing table after rejecting unacceptable routes is:

- the list of the accessible destination IP networks;
- for each of these IP networks:
    - the next router address (next-hop) in the adjacent domain; this information is carried up in the messages inside the AS (i-BGP session);
    - the list of the Autonomous Systems successively traversed (AS path), from the adjacent domain to the AS destination domain (i.e. which contains the destination IP network);
- For each border router: the neighbor routers addresses with whom it has established BGP session (neighbor); they can be which are either border router or Router Reflectors [12].

The way to discover ingress and egress border routers is the following:

a) Discover the ingress border router

   i) If we deal with a NSIS domain and the upstream domain is also a NSIS domain, the ingress router is easy to retrieve. As described in the HyPath signalling, it is the border router that intercepts the NSIS message and redirects it to the RM.

   ii) If we deal with a NSIS domain, and the upstream domain is a non-NSIS domain, the ingress router is retrieved from the message received by the RM. In this case, the upstream RM sends the message directly to the local RM as explained in Section 3.2. This upstream RM interrogates the BGP table of its ingress border router and retrieves the next-hop IP address.

   The goal for the RM of the AS2 is to retrieve the ingress border router in the next domain (AS3). The RM of the AS2 interrogates this table and retrieves the next-hop address of R31. This operation is done only with adjacent non NSIS domains. Then, it passes this address in the signalling message to the RM of AS3.

   A particular case is when the IP address of the next-hop is not distributed through the internal routing protocol (a private IP for instance), and thus, the new request must be addressed to the egress border router. The procedure to discover this router is presented in the next section.

   iii) If we deal with a non-NSIS domain, the procedure is similar to the one presented in point ii)

b) Discover the egress border router:

   i) Inside an AS, all border routers communicate via i-BGP sessions. The egress border router is discovered using the BGP table of the ingress border router. If we deal with a full mesh iBGP (all border routers are connected on iBGP), then the egress border router is a neighbour.

ii) As an alternative, if the domain uses Route Reflectors, either the attribute ORIGINATOR_ID or the domain topology can be used to find the egress border router for the data path.

For non-transit traffic (i.e. traffic originating inside the domain) the RMcan use a database (similar to TED for the PCE Element [13] [14]) where it is stored cartography of the domain (network topology).

We now explain how a Resource Manager can obtain the IP address of a peer Resource Manager in an adjacent domain.

The RM is configured by the administrator of the domain. One solution is to consider that the Server Level Agreement (SLA) between two adjacent domains contains also the IP address of the RM present in the domain. As the administrator is aware of the SLAs, it can configure the RM with all peer RM addresses.

Another solution follows the SIBBS proposal [5] that suggests to retrieve the bandwidth broker address via a DNS mechanism (the BB for each domain is to be named bb.<domain_name> and put it in a CNAME record in the DNS). We propose to have a similar approach. However, instead of using the domain name, we propose to associate the AS number to an RM IP address. When a RM needs to obtain the next RM IP address, it checks the BGP table to find the AS path to the destination. In the AS path, it finds the next AS number and, based on one of the mechanisms presented before, it does the correspondence AS number <-> RM IP address.

3.3 HyPath architecture

The NSIS and the non-NSIS solutions presented are able to work but, as discussed in the previous sections, have disadvantages. On one hand, in the NSIS solution it is mandatory to use NSIS and HyPath aware routers. On the other hand, in the non-NSIS solution it is needed an intensive usage of external functions that extensively access the routing protocol.

A new solution is to integrate the two solutions presented. Particularly, with this approach, in NSIS domains it is used the NSIS solution and in non-NSIS domains it is used the non-NSIS solution.

The difficulty of this hybrid approach is the interaction between domains that have different solutions implemented. When sending signalling messages from one NSIS domain to a non-NSIS domain, the information issued by the ingress border router is not sent within the signalling message and therefore cannot be retrieved locally.

To solve this problem, the NSIS domain must check the type of the next domain before sending any signalling message. This information is obtained from the normal AS association procedure. If the next domain is a non-NSIS domain, the message must be sent as described in the non-NSIS domain case, otherwise it is sent as described in the NSIS domain case.

This approach implies that NSIS domains connected with non-NSIS domains need to determine the type of the next domain, increasing not only the response time but also the complexity of the solution. If a NSIS domain is only connected to other NSIS domains the solution is very simple and light weight.

As described in the previous sections, the additional functionalities for the Hybrid on-path/off-path approach for end-to-end signalling across NSIS and non-NSIS domains implies the usage of HyPath in the RMs and in the border routers.

The main HyPath functionalities are the following:

- In the Resource Manager: discovery of the egress border router of the first domain; discovery of the ingress border router after a non-NSIS domain; RM signalling; message reception and decoding; sending messages;
- In the egress border router: start RM signalling in the first domain;
- In the ingress border router: message interception and sending them to the local RM; reception of the local RM response message and continuation of RM signalling.

These functionalities are described in more detail in remaining of this section.

HyPath in the RM. The HyPath in the RMs is responsible for changing the destination of the signalling message so the right RM is signalled.

In the first domain (the domain where the network signalling starts) the HyPath starts by discovering the egress border router of the data path using an external function. If the next domain (discovered using the external function) is a NSIS domain, the message is sent to the egress border router. Otherwise, the ingress border router and the IP address of the RM of the next domain in the data path must be discovered using again an external function. Afterwards, the message is sent directly to the IP address of the next domain RM.

If a domain is not the first domain, it means that a HyPath message has already been received and there is state stored in the HyPath database. If the next domain and the current domain are NSIS domains, the message is sent to the ingress border router (IP address in the database) to be forwarded through the same path as the data. If the next domain is a non-NSIS domain, then again, an external function must be used to discover the ingress border router and the IP address of the RM of the next domain in the data path. Afterwards, the message is sent directly to the IP address of the next domain RM.

The messages to be sent upstream first need to query the HyPath database for the upstream RM IP address (state stored when a downstream message is received) and then are sent directly to the RM.

When a message is received from GIST, the HyPath information in the message must be recorded in the database and the NSLP layer payload must be sent to the NSLP.

HyPath in the border router. In the border routers the NSLP layer is not needed, so the HyPath acts as a normal NSLP.

The HyPath module in the border router has two different functionalities depending if it is an egress or ingress border router. In the first domain, the border router acts as the egress router where the signalling merges with the data path. From this point forward, if the message is always sent to the end user, the signalling path will follow the same path as the data path. In the other domains, the border router acts as an ingress border router where HyPath NSIS messages are intercepted.

In the border router, if the messages are received from the local RM, they are forwarded to the end user with the IP Router Alert Option flag. If the messages are intercepted, they are forwarded to the local RM.

4. Use cases of HyPath

In this section the HyPath solution is described using a group of diagrams that intend to show the flow of the signalling messages through the several domains (NSIS and non-NSIS domains) between the sender and the receiver.

Figure 4 shows a diagram with the sequence of messages in the HyPath signalling in NSIS domains. To make the signalling of all RM servers between the domains AS1 and AS4 when all domains have border routers with NSIS (and the HyPath), the RM in the AS1 domain starts by using an external function to discover the local egress border router of the data path and then sends a NSIS message to that router (1). In the border router, the message is then sent to the destination (2), through the same path as the data. In the ingress border router of the next domain, AS2, the NSIS message is intercepted and sent to the local RM (3). After the appropriate processing in the RM, if the signalling is to continue to another domain, the local RM sends a message back to the ingress border router (4). The border router will then continue the signalling through the same path as the data (5). This processing continues until the signalling message reaches the last domain (AS4) where the signalling can stop on the RM or in the destination IP address of the signalling message.

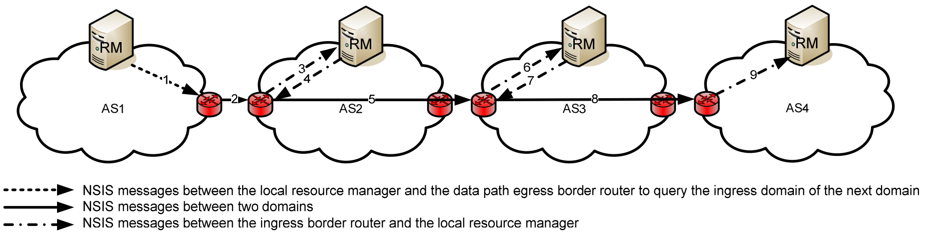
RM
RM
RM
RM
1
2
3
4
5
6
7
8
9
AS1
AS2
AS3
AS4
NSIS messages between the local resource manager and the data path egress border router to query the ingress domain of the next domain
NSIS messages between two domains
NSIS messages between the ingress border router and the local resource manager

Figure 4 - HyPath messages in signalling between NSIS domains

Figure 5 shows a diagram with the sequence of messages in the HyPath signalling in non-NSIS domains. When none of the domains have NSIS in the border routers, the previous approach cannot be used because the interception of NSIS messages is not possible. Therefore, in order to make a RM signalling between the AS1 and AS2 domains, the first RM uses an external function to discover the local egress border router of the data path, and makes a request to this border router to inquire the next domain AS number and its ingress border router (1).

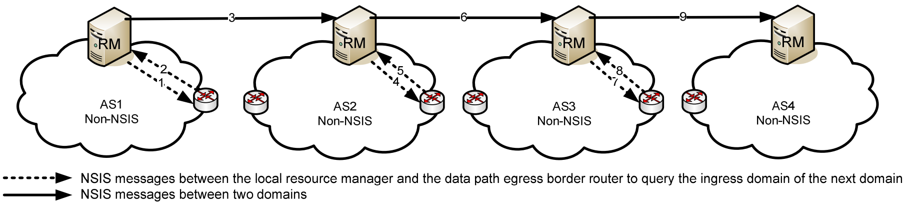
3
6
9
RM
RM
RM
RM
2
1
5
4
8
7
AS1
Non-NSIS
AS2
Non-NSIS
AS3
Non-NSIS
AS4
Non-NSIS
NSIS messages between the local resource manager and the data path egress border router to query the ingress domain of the next domain
NSIS messages between two domains

Figure 5 - HyPath messages in signalling between non-NSIS domains

The border router replies with the required information (2) obtained from the inter-domain routing protocol. With this information and the RM IP addresses of the adjacent domains contained in the SLAs, the NSIS signalling message is sent to the next RM (3) through a path that can be different from the data path. This processing continues until the signalling message reaches the last domain (AS4), where the signalling stops.

In a transition phase, when not all domains are able to have NSIS border routers, a hybrid solution must be applied in order to allow the signalling between the RMs of two domains. Figure 6 shows a diagram with the sequence of messages in the HyPath signalling between NSIS and non-NSIS domains. This sequence of messages is similar to the two previous use cases, depending only on the type of the next domain. The type of the next domain (NSIS or non-NSIS domain) is available from the SLAs between the adjacent domains. If a NSIS domain does not have adjacent non-NSIS domain, it can use the first use case method, otherwise the SLAs must be checked to see if the next domain is a non-NSIS. If the next domain is a NSIS domain, then the first use case method can be used again, otherwise the second use case must be used. A non-NSIS domain always uses the second use case method.

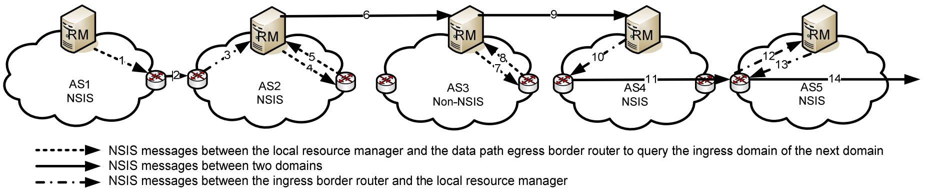
RM
RM
RM
RM
RM
6
9
1
2
3
5
4
8
7
10
11
12
13
14
AS1
NSIS
AS2
NSIS
AS3
Non-NSIS
AS4
NSIS
AS5
NSIS
NSIS messages between the local resource manager and the data path egress border router to query the ingress domain of the next domain
NSIS messages between two domains
NSIS messages between the ingress border router and the local resource manager

Figure 6 - HyPath messages in signalling between NSIS and non-NSIS domains

5. Conclusions

In this paper was presented an architecture in the context of a NSIS multi domain Internet that aims at off path signalling when a hybrid solution is required (for instance NSIS is not implemented in all domains).

Currently, an increasing number of applications claim special treatment for their packets in order to satisfy new requirements in terms of delay, loss and jitter. Inside an AS, the QoS management is often delegated to a central entity which has a global view of network topology. This entity is also aware of QoS availability inside and on the inter domain links of the domain. In order to signal these entities, which are not on the data-path, this paper proposed a solution called Hybrid Path to involve the central entities on the signalling in the NSIS context.

The proposed Hybrid on-path off-path approach for end-to-end signalling across NSIS and non-NSIS domains is being used in the EuQoS project (http://www.euqos.org). Nevertheless, it can also be used for all signalling that needs to signal specific entities in all the domains in the data path.

Acknowledgments

This work was partially funded by the European Union 6th Framework Programme under contract IST FP6 IP 004503 EuQoS Integrated Project.

References

[1] Blake, S;, Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W. “An Architecture for Differentiated Services”, RFC 2475, December, 1998

[2] Next Steps in Signaling (NSIS) Charter, http://www.ietf.org/html.charters/nsis-charter.html

[3] Hancock, R., Kappler, C., Quittek, J., Stiemerling, J., “A Problem Statement for Partly-Decoupled Signalling in NSIS”, draft-hancock-nsis-pds-problem, October 2005

[4] Schelen, O., Couturier, A., Bless, R., Geib, R., Dugeon, O., “Path-coupled and Path-decoupled Signalling for NSIS“, draft-schelen-nsis-opopsig-01.txt, November 2002

[5] Simple Inter-domain Bandwidth Broker Signaling, QBone Signaling Design Team, http://qbone.internet2.edu/bb/

[6] Trang Nguyen TM, Boukhatem N., Doudane, Y.G., Pujolle, G. “COPS-SLS – a service level negotiation protocol fort he internet“, IEEE Communication Magazine, vol. 40, n°5, May 2002

[7] Trang Nguyen TM, Boukhatem N., Doudane, Y.G., Pujolle, G, “COPS-SLS usage for dynamic policy-based QoS management over heterogeneous IP networks“, IEEE Communication Magazine, vol. 17, n°3, May-June 2003

[8] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R., Sastry, A. “The COPS (Common Open Policy Service) Protocol” RFC 2748, January 2002

[9] Braden, R., Zhang, L., Berson, S., Herzog, S., Jamin, S., “Resource ReSerVation Protocol (RSVP)”, RFC 2205

[10] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signalling Transport", draft-ietf-nsis-ntlp-09, September 2005

[11] Bosch, S., "NSLP for Quality-of-Service signalling", draft-ietf-nsis-qos-nslp-09, October 2005.

[12] Bates, T., Chandra, R., Chen, E., "BGP Route Reflection - An Alternative to Full Mesh IBGP", RFC 2796, April 2000

[13] Path Computation Element (PCE) Charter, http://www.ietf.org/html.charters/pce-charter.html, 2005

[14] Farrel, A., Vasseur, J-F., Ash, J., “A Path Computation Element (PCE) Based Architecture“, draft-ietf-pce-architecture-04.txt, January 2006