NSIS Working Group
Internet Draft
Expires: February 2007

L. Cordeiro
M. Curado
E. Monteiro
V. Bernardo
University of Coimbra
F. Racaru
M. Diaz
C. Chassot
LAAS
October 2006

Hybrid on-path off-path approach for end-to end signalling accross NSIS domains (HyPath) draft-cordeiro-nsis-hypath-02

Status of this Memo

Copyright Notice

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, the HyPath approach uses the NSIS protocol and interactions with the local routing protocols to have an off path signalling in hybrid environments.

Table of Contents

1. Introduction

In the past years, we assisted at a common rise of new technologies in the telecommunication and computer science field. This evolution led to the emergence of new types of applications involving multimedia, like VoIP, VoD, tele-engineering, telemedicine, etc.. 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 QoS (Quality of Service) 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 a context of a multi-domain Internet that offers QoS guaranties for applications. Inside a domain, 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 DiffServ domains, and is associated with Bandwidth Brokers [6]. At the present, a new requirement appears: signalling MUST take place, not only among devices strictly on the data path, but also among these central entities, that we call hereafter Resource Manager (RM).

Several signalling protocols have been proposed, especially in the IETF NSIS working group. 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 document discusses a NSIS multi-domain, multi-service, RM based Internet that allows an off-path signalling. The main issue addressed in this document is the inter-operability between NSIS and non-NSIS domains.

1.1 Terminology and Abbreviations

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [10].

This document uses a number of terms defined in [5]. The following additional terms are used:

- o QoS: Quality of Service
- o RM: Resource Manager, central entity of a domain in charge of the QoS management.
- o NSIS domain: an administrative domain where the RM and at least all border routers are NSIS aware.
- o Non-NSIS domain: an administrative domain where only the RM is NSIS aware.
- o E2E: end-to-end

2. Off-path signalling state of the art

In the off-path approach, 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 (QoS, policy control, servers, etc\205) 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 [4] and [5]:

- o independence between the signalling plane and the forwarding plane
- o introduction of flexibility allowing entities such as proxies to be signalled even if they are not on the data path;
- o functioning with new routing protocols or traffic engineering mechanisms (QoS routing, q-BGP, etc\205);
- o better adapted for mobility.

On the other hand, off-path signalling MUST answer new challenges such as discovering the next hop and synchronisation with IGP (Interior Gateway Protocol) and EGP (Exterior Gateway Protocol) routing protocols.

2.1 Off-path signalling proposals

Several protocols have been proposed for the off-path signalling in a bandwidth broker-based multi-domain DiffServ model.

2.1.1 SIBBS (Simple Inter-domain Bandwidth Broker Protocol)

The SIBBS protocol has been defined by the QBone Signaling Workgroup and it aims to be used on DiffServ bandwidth broker-based domains. In the QBone testbed, each network is a differentiated service (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 borkers. It contains two principal PDU:

- o RAR (Resource Allocation Request)
- o RAA (Resource Allocation Answer)

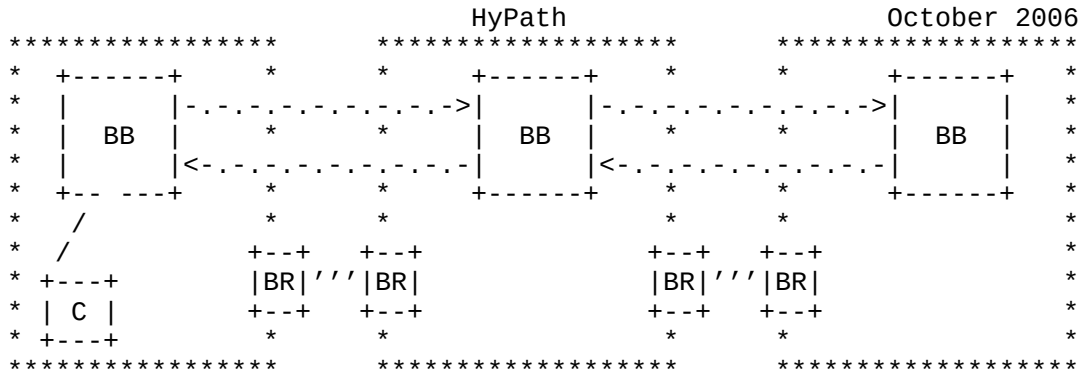

<-.-.-.-.-> = signalling message betwwen BB
---------- = message between client and BB
C = client
BB = Bandwidth Broker BR = Border Router
Figure 1 \226 SIBBS protocol

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. The communication between BB is supposed to be reliable, i.e. using TCP.

Receiving a RAR message, a BB:

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

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

In order to perform the configuration of the border routers, the BB MUST have access to this one. SIBBS does not specify a particular protocol, but some examples are COPS, DIAMETER, SNMP, etc..

2.1.2 COPS-SLS

COPS-SLS [7, 8] is an extension of the COPS (Common Open Policy Service) protocol [9] 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:

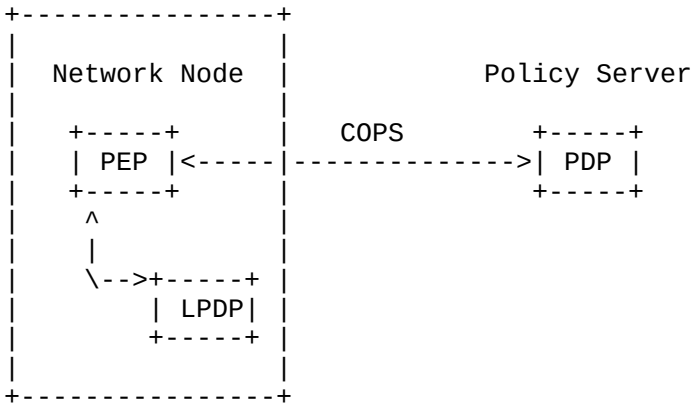

Figure 2 \226 SIBBS protocol

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:

- o PDP for the upstream domain, BB which sends the request, and
- o PEP for the next BB domain.

Compared to SIBBS, COPS-SLS adds some features to the protocol, as 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. COPS-SLS does not provide any

specification on the discovery of the next BB discovering or on the identification of border routers.

3. HyPath

The requirements for a hybrid on-path/off-path approach for e2e 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). An example of the normal NSIS signalling from a source RM to a destination RM is shown in Figure 3.

The normal way of work of the NSIS protocol [1], 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). As presented in Figure 3, the signalling message could not follow the same inter-domain path from the sender domain to the receiver domain. Therefore the resource reservation will not be properly done on the data path.

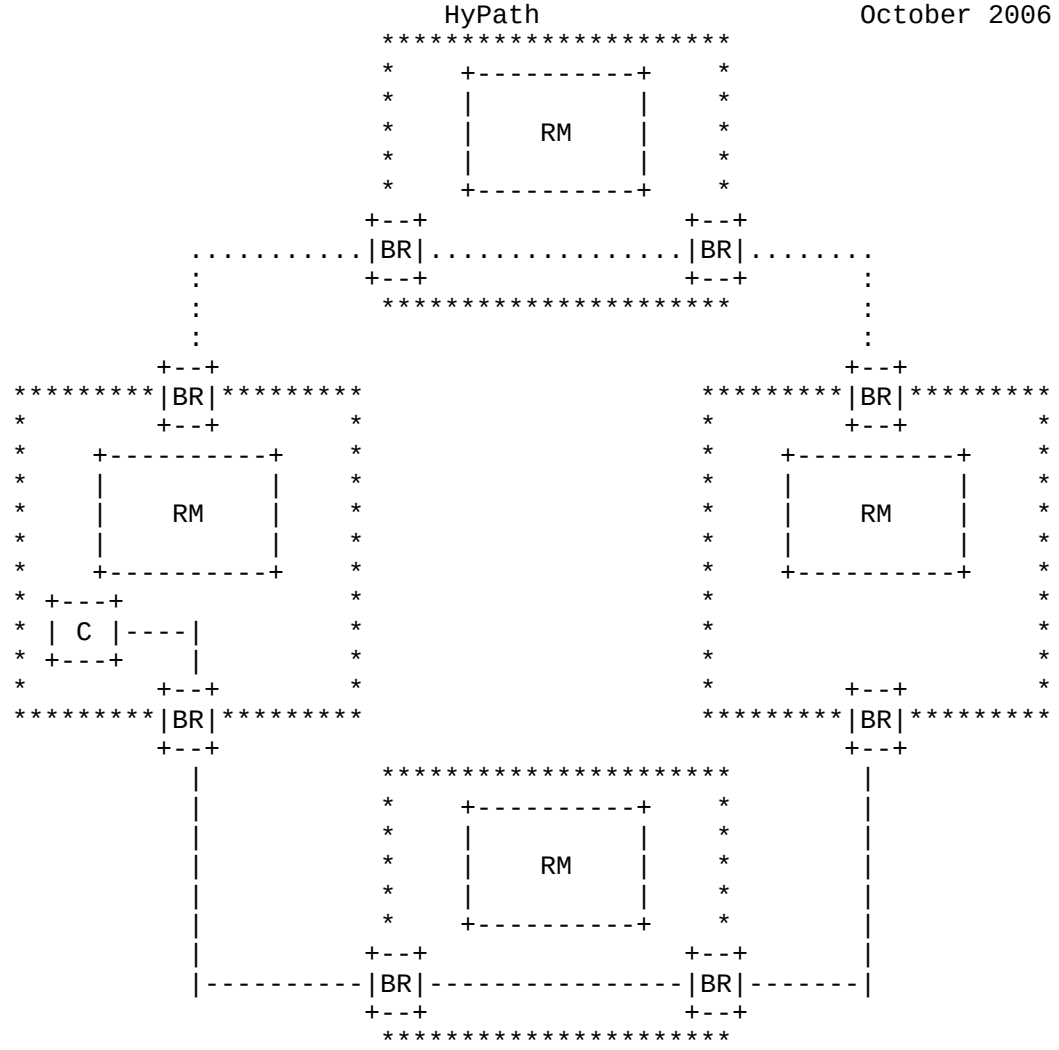

Figure 3 \226 Normal NSIS signalling

The major requirements to achieve e2e (end-to-end) network signalling are the following:

- o Signalling messages MUST follow the same path as the user data;
- o All the RMs in 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 [2] and the NSLP layer [3] 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). Figure 4 describes an example of the NSIS protocol architecture including the HyPath.

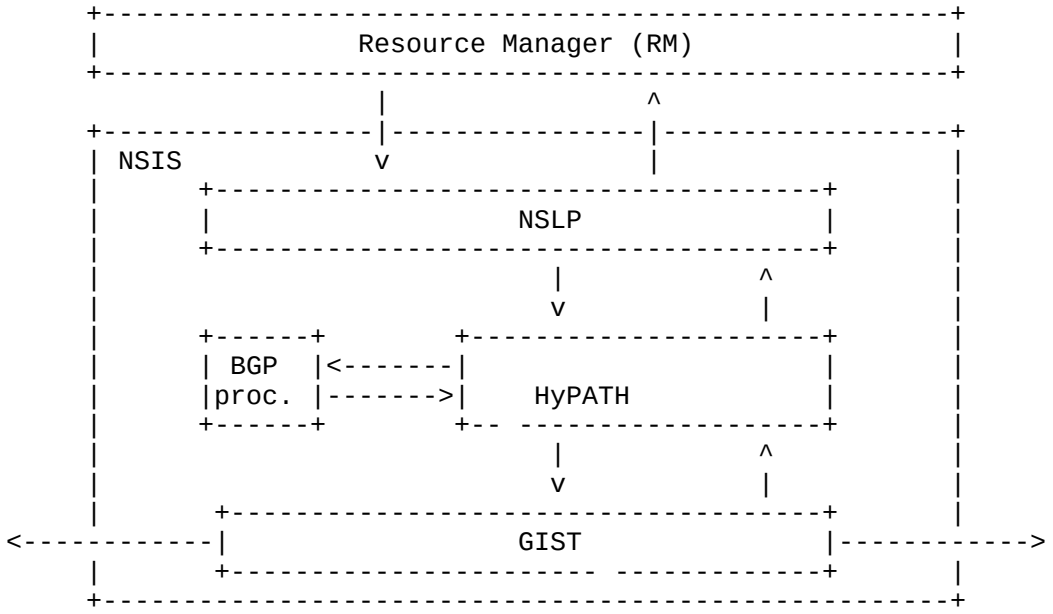

HyPath on the NSIS architecture
BGP proc = BGP processing + external functions

Figure 4 \226 HyPath in the NSIS architecture

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 defined NTLP layer interface.

3.1 The new signaling

The operation of NSIS with the additional HyPath in the border routers and RMs in all domains is illustrated in Figure 5.

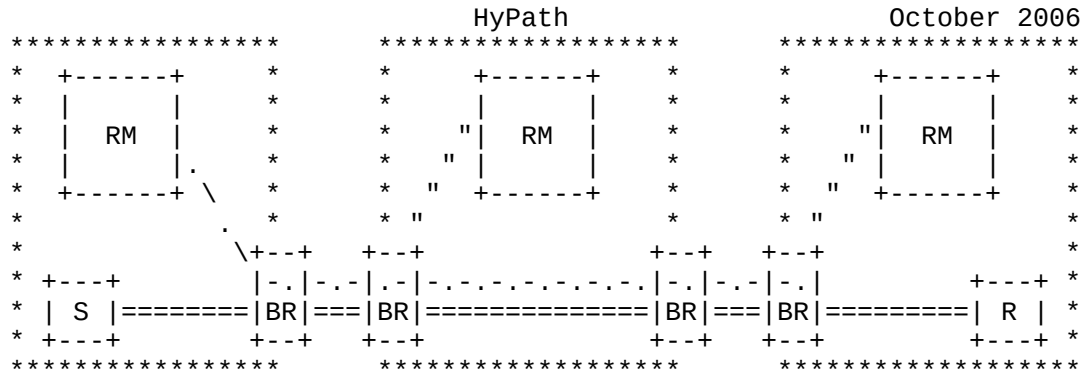

Figure 5 - NSIS signalling with HyPath

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 in Section 3.3) 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 in Section 3.5).

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 e2e 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).

3.2 Non-NSIS domains

The drawback of the approach described in Section 3.1 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 in Section 3.3) 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 in Section 3.5) 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 signaling messages do not follow the data path, but they follow all the RMs in the data path.

The inconvenience with 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.

3.3 Usage of external routing protocols

In the beginning of this document the motivation for off-path signaling was described. If the signaling is decoupled from the data path (but still path-related signaling) two general problems need to be solved:

- The Resource Manager 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 signaling 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 Resource Manager in the next domain MUST be identified in order to propagate the request.

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

- Accessible destination network list (IP prefixes);
- For each prefix:
  - next router address (next-hop) in the adjacent domain; this information is carried up in the messages inside the AS (i-BGP session);
  - List of Autonomous Systems successively traversed (AS path), from adjacent domains to the AS destination domain for the destination network;
- For each border router: address of neighbor routers with whom it has established BGP session (neighbor) which are either border router or Router Reflectors \226 RFC 2796.

The approach 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 signaling (Section 3.1), 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 sent 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. In the upstream AS2 domain the BGP R21 tables contains:

Network	Next-Hop	Path
*>ip form AS3	R31	i 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.

When this RM receives the message, it has already the IP address of the ingress border router. If this address is not a public address (a private IP address for instance), it can obtain this public address from local BGP and topology configuration.

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 in the i-BGP session. 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 Resource Manager can use a database (similar to TED for the PCE Element [11] [12]) where it is stored cartography of the domain (network topology).

In the next paragraph we 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 SLA (Server Level Agreement) between two adjacent domains contains also the PI 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 [3] 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.

In this section we described some mechanisms that allow for the solution of major issues of the off-path signalling path, namely to discover the next hop to signal and how to interact with external routing protocols such as BGP.

3.4 Heterogeneous solution

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 Section 3.2, otherwise it is sent as described in Section 3.1.

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.

3.5 HyPath architecture

As described in the previous sections, the additional functionalities for the Hybrid on-path/off-path approach for e2e 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 the next subsections.

3.5.1 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. If the RM is not the first or the last RM of the signalling, the MRI direction flag (dFlag) MUST be changed. This change allows the usage of the previous associations.

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.

3.5.2 HyPath in the border router

In the border routers the NSLP layer is not needed, so the HyPath acts as a normal NSLP. Figure 6 shows the NSIS architecture in a border router with HyPath.

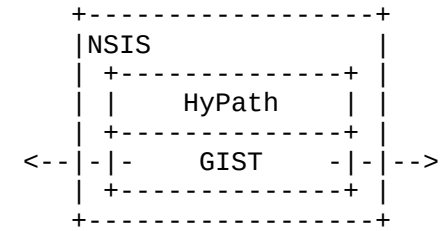

Figure 6 \226 HyPath architecture in the border routers

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. If the messages are intercepted, they are forwarded to the local RM.

If the received message is from the local RM (message type 3) and the current network is neither the source nor the destination network, the message direction MUST be set with the D Flag field of the payload. The Border Router Address field of the payload MUST be set with the address of the current machine and the B Flag set.

If the received message is of type 1 (to be sent to the final destination) and if the B Flag is set, the Border Router Address field in the payload MUST be checked. If the Border Router Address belongs to the local domain, the message is forwarded, only updating the Border Router Address field. If the Border Router Address is not from the local domain the message MUST be sent to the local RM. In this case the MRI source address MUST be set to the BR address, the MRI destination address MUST be set to the local RM address and the HyPath payload Border Router Address field MUST be updated with the local BR address. When sending the message to the RM, if the message direction is upstream, the source and destination SHOULD be swapped.

3.5.3 HyPath payload

HyPath needs to manage state in the RM and border routers to be able to handle all types of messages. Therefore it needs to send specific information together with the network signalling message.

A Hypath message consists of a message type, followed by a flag stating if there is the border router address field, a reserved field for future use and a length field. Next, there are two fields containing the source and destination addresses of the message. Finally, there is a field containing the address of the border router of the domain.

The field of the source address is filed with the address of the source of the data sender. The destination field is filed with the address of the data destination. The border router field is filed with the address of the egress border router of the domain discovered by Hypath through the external function, described above, or by the ingress border router that intercepted the HyPath message.

The binary payload description is presented next.

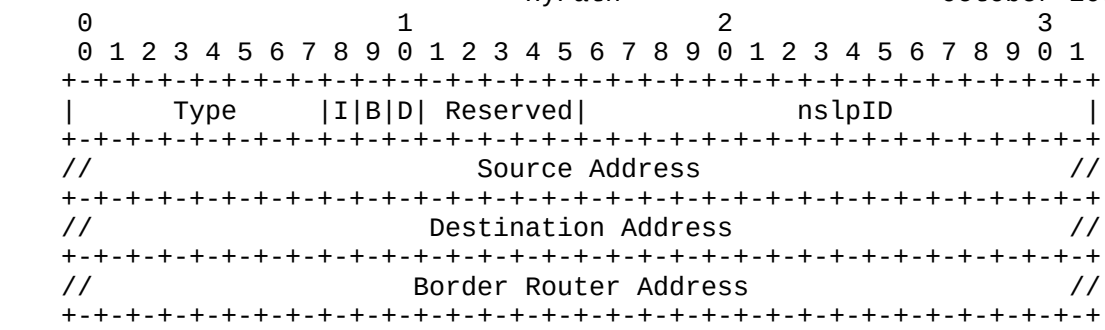

Type: type of HyPath message.

1 - Standard message: a message that follows the data path and it is intercepted in the border routers

2 - Border router message: a message that it is sent from the border router to the RM

3 - Resource Manager message: a message that is sent either to the border router of the actual domain or to the RM of the next domain, depending on the type of the next domain (NSIS or non-NSIS)

Flag I: IP version [0 = IPv4 | 1 = IPv6]
Flag B: means that the border router address field is present
Flag D: message direction flag

Source Address: data source address

Destination Address: data destination address

Border Router Address: border router address to where the resource manager MUST send the next message

3.6 HyPath multiple domain example

To better understand how HyPath works in a multiple domain situation, a three domains example is described next. This example is composed by three networks (A, B and C), three RMs (RM A, RM B and RM C), four border routers (BR A1, BR B1, BR B2 and BR C1) and two users (S, source user, and D, destination user).

Network A is composed by the source user, S, RM A and BR A1. Network B is composed by RM B and two BRs, BR B1 and BR B2. Network C is composed by the destination user, D, RM C and BR C1. Network A connects to network B through BR A1 and BR B1 and network B connects to network C through BR B2 and BR C1. The connection

between the network A and network C MUST be through the network B. The data path between the source user and the destination user is the following sequence: S -> BR A1 -> BR B1 -> BR B2 -> BR C1 -> D.

Figure 7 illustrates the three domains network described with the signalling messages required to make a signalling between the source and the destination users. In this case the source user, S, requests the local RM, using one signalling protocol (independent from HyPath) that initiates HyPath signalling to the destination user.

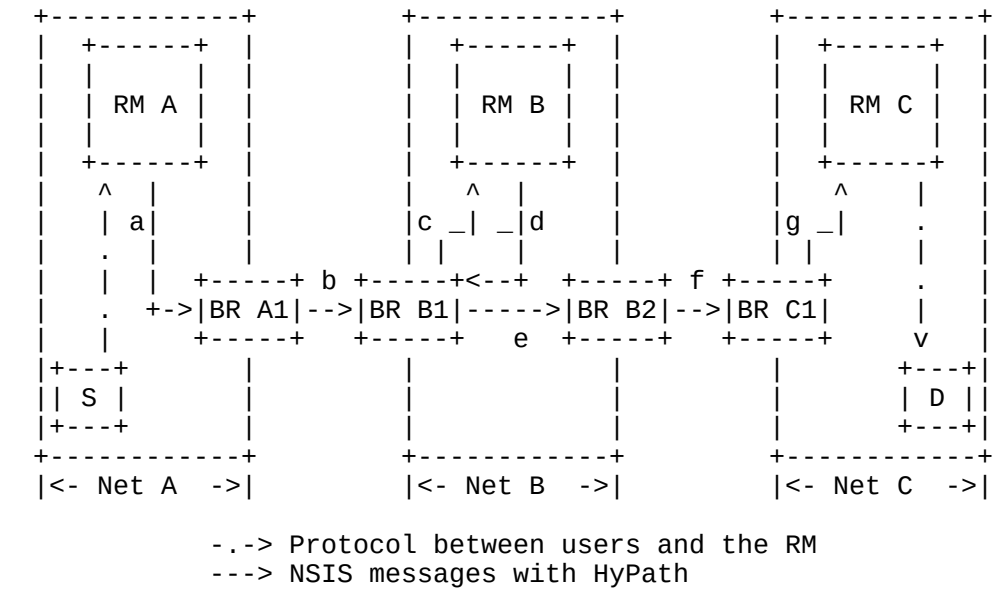

Figure 7 \226 HyPath architecture in the border routers

The HyPath signalling starts in the RM A, where the NSLP requests HyPath to send a message from S to D. Before starting the signalling, the local data path egress border router needs to be discovered (using an external function). With this information, the RM A sends the message to the respective egress border router, BR A1. This message is type 1.

When the BR A1 receives the message, knows it MUST be injected in the data path because it is a type 1 message. BR A1 changes the received MRI with the information available in the HyPath payload (source IP address, destination IP address and direction) and sends the message towards the destination. This is the first message that follows the data path, like as if it had been generated by the source user.

In the network B (the next domain) the message is intercepted by the ingress border router (BR B1). This border router checks if the message needs to be sent to the local RM or not. Since the message was received from another domain (by checking the BR Address field in the HyPath payload) the message is then forwarded to the local RM, the RM B. The MRI is changed in order to send the message to the local RM. Moreover, the source is changed to the BR IP address and the destination is changed to the local RM IP address. The direction of the message is also set to downstream.

In the RM B, when Hypath receives a message, sends the NSLP payload to the respective NSLP (if it exists). If NSLP sends a response message, it is sent back to the respective BR, the BR B1. Before sending the message, the HyPath state table needs to be searched for the specific session to find the respective BR IP address. One message that has no state (in case it is the first message or the state was removed due to a time out) is treated as a new signal. With the BR B1 IP address, the message is sent upstream (using the already established states) to BR B1.

When a message is received in BR B1 from the RM B (because of the message type) the message is injected in the network like in BR A1. The message is then sent to the destination.

In network B the message is intercepted by the egress border router, the BR B2. In this BR, the message does not need to be sent to the local RM because the previous node is a local node (BR B1). This means that no processing is needed and the message is forwarded to the destination. Only the BR IP address in the HyPath payload SHOULD be updated.

In network C the message is again intercepted by the ingress border router, BR C1. The ingress border routing procedure is the same as described above, and the message is sent to the local RM (RM C).

In the RM C the node is treated as the last node because the destination user is in the RM network. Here, the HyPath signalling stops and other type of signalling (protocol independent from the HyPath like the Session Initiation Protocol [13]) can interact with the destination user.

When the NSLP requests a response message to a HyPath message in the RMs, the procedure is similar to the one described previously. The message SHOULD be sent upstream and, if new associations are required, the GIST layer SHOULD support their creation upstream.

4. Hypath implementation example

This section describes an example of Hypath implementation architecture. Figure 8 shows the Hypath processing of messages received from the NSLP layer and from the GIST layer.

The architecture was developed in order to allow for the connection of multiple signalling applications to Hypath. Therefore, for each NSLP there is an NSLPThread to handle the messages exchanged between the NSLP and Hypath layers. The Hypath layer is connected to the NTLP layer, GIST, through the NTLPThread. This thread handles all the communication between Hypath and NTLP.

The message flow from NSLP to GIST is as follows: when a specific NSLP sends a message, the corresponding NSLPThread receives the request and queues it on the SendQueue for further processing. Afterwards, the SendProcessing thread consumes the request from the SendQueue and, it processes each request accordingly. Finally, the result it is sent to the NTLPThread in order to be transported by GIST.

The message flow from GIST to NSLP is as follows: when the NTLPThread receives a message from GIST, the message is queued in the RcvQueue. Then, the RecvProcessing module consumes the message from the RcvQueue and processes it. If the current Hypath is on a Border Router, the message is sent back to the NTLPThread to be sent to the local Resource Manager. Otherwise the message is sent to the corresponding NSLPThread. If no NSLPThread is registered for the received message, the message is sent back to the NTLPThread to continue the signaling

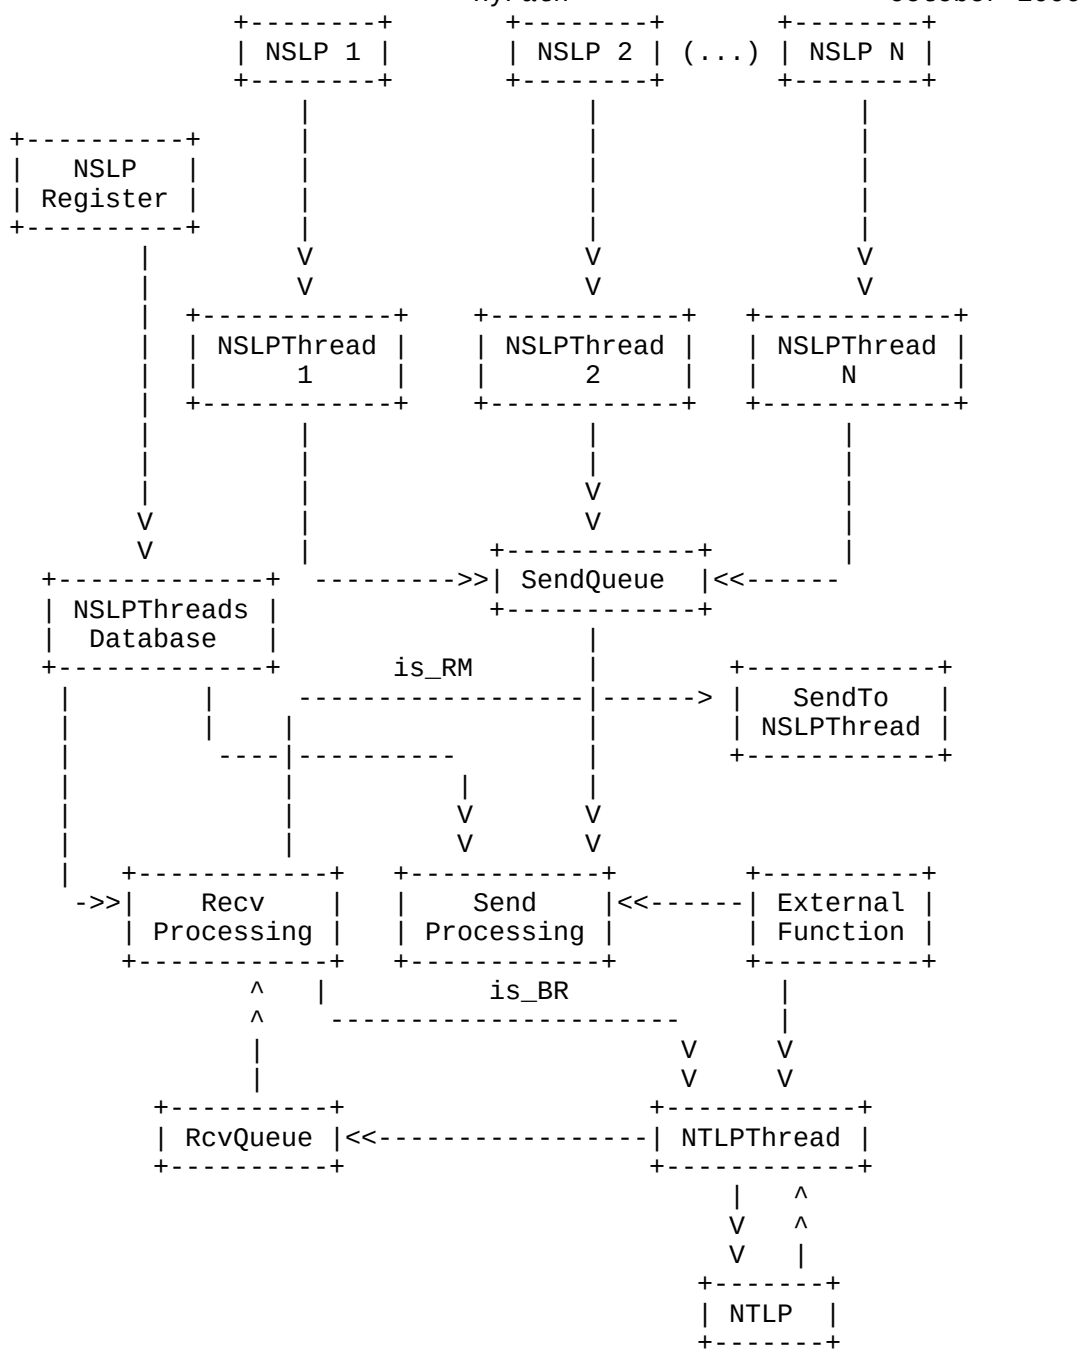

Figure 8 \226 HyPath architecture

5. Acknowledgements

The authors would like to thank Luís Conceição, Vitor Bernardo and David Palma for their contribution on the implementation and testing of GIST, Hypath and QoS-NSLP.

6. Security Considerations

This section describes the security considerations related to the HyPath and this will be discussed and clarified later.

7. Conclusion

This draft presented architecture in the context of a NSIS multi domain Internet that aims an off-path signalling when a hybrid solution is required (fir instance NSIS is not implemented in all domains). Currently, increasing number applications claim special treatment for their packets in order to satisfy new requirements in term of delay, loss, jitter, etc. Inside an AS, the QoS management is often delegated to e 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 draft proposed a solution to involve the central entities on the signalling in the NSIS context, solution called Hybrid Path.

This work, a Hybrid on-path off-path approach for e2e signalling across NSIS and non-NSIS domains, aims specifically at the EuQoS project (http://www.euqos.org), but also for all network signalling that needs to signal specific entities in all domains in the data path.

6. Open issues

This section describes the open issues related to the HyPath and this will be discussed and clarified later.

7. References

[1] Hancock, R., Karagiannis, G., Loughney, J., and S. Van de Bosch, "Next Steps in Signaling (NSIS): Framework", RFC 4080, June 2005.

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

[3] Bosch, S., "NSLP for Quality-of-Service signalling",

draft-ietf-nsis-qos-nslp-09, October 2005.

[4] Hancock, R., Kappler, C., Quittek, J., Stiemerling, J., \223A Problem Statement for Partly-Decoupled Signalling in NSIS\224, draft-hancock-nsis-pds-problem-02.txt, October 2005.

[5] Schelen, O., Couturier, A., Bless, R., Geib, R., Dugeon, O., \223Path-coupled and Path-decoupled Signaling for NSIS \223, draft-schelen-nsis-opopsig-01.txt, November 2002

[6] Blake, S;, Black, D., Carlson, M., Davies, E., Wang, Z., Weiss, W. \223An Architecture for Differentiated Services\224, RFC 2475, December, 1998

[7] Trang Nguyen TM, Boukhatem N., Doudane, Y.G., Pujolle, G. \223COPS-SLS \226 a service level negotiation protocol fort he internet\223, IEEE Communication Magazine, vol. 40, n°5, May 2002

[8] Trang Nguyen TM, Boukhatem N., Doudane, Y.G., Pujolle, G, \223COPS-SLS usage for dynamic policy-based QoS management over heterogeneous IP networks\223, IEEE Communication Magazine, vol. 17, n°3, May-June 2003

[9] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R., Sastry, A. \223The COPS (Common Open Policy Service) Protocol\224 RFC 2748, January 2002

[10] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[11] \223Path Computation Element (PCE) Charter, http://www.ietf.org/html.charters/pce-charter.html\223, 2005.

[12] Farrel, A., Vasseur, J-F., Ash, J., \223A Path Computation Element (PCE) Based Architecture\223, draft-ietf-pce-architecture-04.txt, January 2006

[13] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A. and all, \223SIP: Session Initiation Protocol", RFC 3261, June 2002.

8. Author’s Addresses

Luís Cordeiro
Dept. of Informatics Engineering
Univ. of Coimbra
Polo II - Pinhal de Marrocos
3030-290 Coimbra, Portugal
Email: cordeiro@dei.uc.pt

Marilia Curado
Dept. of Informatics Engineering
Univ. of Coimbra
Polo II - Pinhal de Marrocos
3030-290 Coimbra, Portugal
Email: marilia@dei.uc.pt

Edmundo Monteiro
Dept. of Informatics Engineering
Univ. of Coimbra
Polo II - Pinhal de Marrocos
3030-290 Coimbra, Portugal
Email: edmundo@dei.uc.pt

Vitor Bernardo
Dept. of Informatics Engineering
Univ. of Coimbra
Polo II - Pinhal de Marrocos
3030-290 Coimbra, Portugal
Email: vmbern@student.dei.uc.pt

Florin Racaru
Laboratory for Analysis and
Architecture of Systems
Avenue du Colonel Roche
31077 Toulouse
France
Email: fracaru@laas.fr

Michel Diaz
Laboratory for Analysis and
Architecture of Systems
Avenue du Colonel Roche
31077 Toulouse
France
Email: diaz@laas.fr

Christophe Chassot
Laboratory for Analysis and
Architecture of Systems
Avenue du Colonel Roche
31077 Toulouse
France
Email: chassot@laas.fr

9. Intellectual Property Statement

The IETF takes no position regarding the validity or scope of any

Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat 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 implementers or users of this specification can be obtained from the IETF on-line IPR repository at http://www.ietf.org/ipr.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

10. Disclaimer of Validity

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

11. Copyright Statement