NSIS Working Group
Internet-Draft
Intended status: Informational
Expires: September 6, 2007

L. Cordeiro
M. Curado
E. Monteiro
V. Bernardo
University of Coimbra
F. Racaru
M. Diaz
C. Chassot
LAAS
March 5, 2007

GIST Extension for Hybrid On-path Off-path Signaling (HyPath)
draft-cordeiro-nsis-hypath-03

Status of this Memo

Copyright Notice

Abstract

In a multi-domain Internet that offers QoS guaranties for applications, there is the need of signaling among the domain entities which are responsible for QoS management. Because different do the HyPath approach uses the NSIS protocol and interactions with

the local routing protocols to achieve an off path signaling in hybrid environments.

Table of Contents

1. Introduction

During past years, we assisted to 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 among others, multimedia, like VoIP, VoD, tele-engineering, telemedicine. These applications have new constraints and requirements concerning Quality of Service (QoS) 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. However, as presented before, some data flows need special handling in order to satisfy the application requirements, and thus it is necessary to address 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 signaling appears to be inevitable.

Our work aims at a context of a multi-domain Internet that offers QoS guarantees 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 [1]. At the present, a new requirement appears: signaling 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 signaling protocols have been proposed, especially in the IETF NSIS working group. NSIS [2] is a new approach that proposes to define a new signaling architecture. The goal of the NSIS protocol is to manipulate the network state related to data flows with the constraint that the messages will be processed on the nodes which also handle the data flows themselves ("path-coupled signaling"). This document discusses a NSIS multi-domain, multi- service, RM based Internet that allows off-path signaling. 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 [3].

This document uses a number of terms defined in [4]. The following

additional terms are used:

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

2. Off-path signaling state of the art

In the off-path approach, entities participating in the signaling 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...) MUST be signaled. These devices are not strictly on the data path; nevertheless the signaling protocol MUST arrive to interact with these devices. Off-path signaling has advantages, as presented in [4] and [5] :

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

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

2.1. Off-path signaling proposals

Several protocols have been proposed for the off-path signaling 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, as described in Figure 1, is a very simple protocol to be used between bandwidth brokers. It contains two principal PDUs:

o RAR (Resource Allocation Request)

o RAA (Resource Allocation Answer)

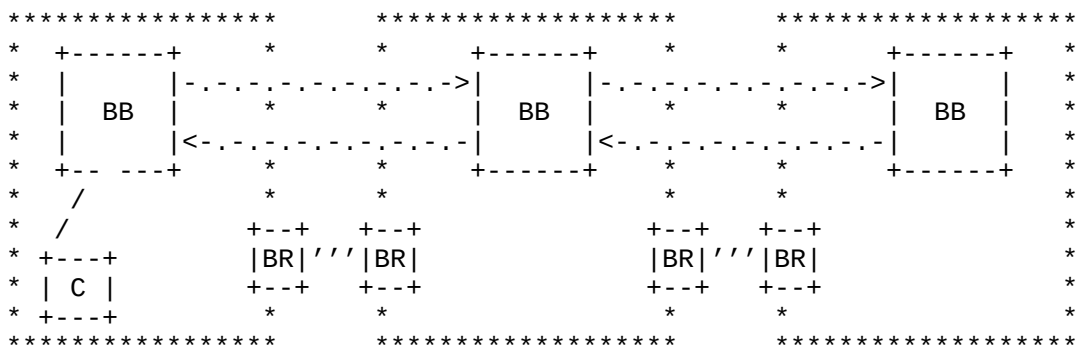

Figure 1: SIBBS protocol

The RAR message includes a globally well-known service ID, information related to the QoS request (class of service 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 the BB MUST have access to the border routers. SIBBS does not specify a particular protocol, but some examples are COPS, DIAMETER, SNMP

2.1.2. COPS-SLS

COPS-SLS [6][7] is an extension of the COPS (Common Open Policy Service) 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 in Figure 2

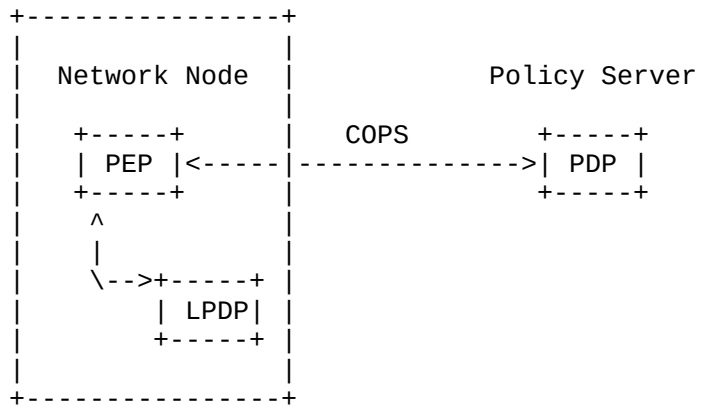

Figure 2: COPS protocol

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 signaling 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 behavior 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:

- PDP for the upstream domain, BB which sends the request, and
- PEP for the next the 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 the 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 or on the identification of border routers.

3. HyPath

The requirements for a hybrid on-path off-path approach for end-to-end (e2e) signaling 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 signaling 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 signaling when there are resource manager entities in the domains responsible for the domain QoS. In these situations the entities to be signaled are the RM entities and not only the network equipment (routers). An example of the normal NSIS signaling from a source RM to a destination RM is shown in Figure 3.

The normal way of work of the NSIS protocol [2], does not signal the RM servers in the data path, and it does not force the signaling 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 signaling 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.

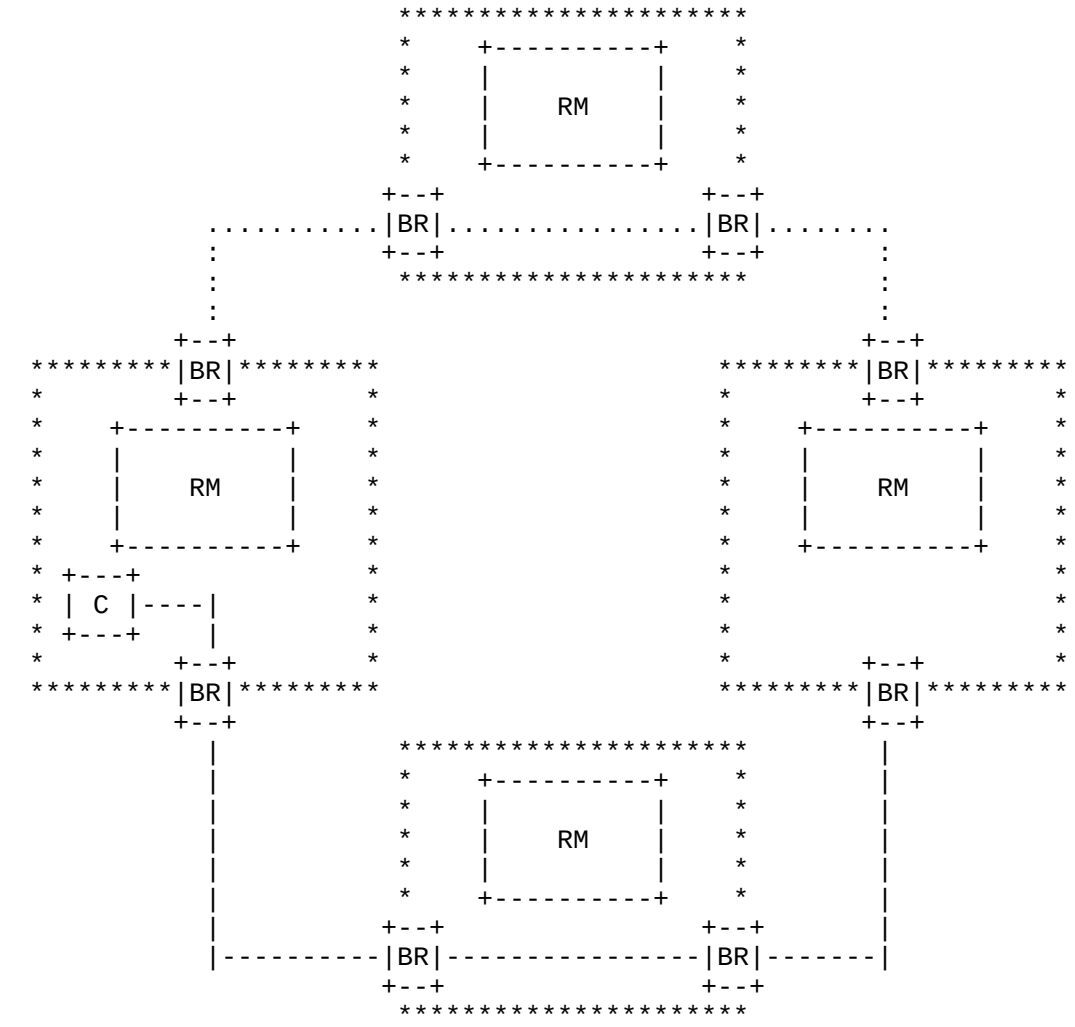

Figure 3: Normal NSIS signaling

The major requirements to achieve e2e network signaling are the following:

o signaling messages MUST follow the same path as the user data;

o All the RMs in the data path MUST be signaled.

The NSIS protocol as it is being defined in the IETF can not solve these two major requirements simultaneously. In order to fulfill the above requirements, a new Message Routing Method (MRM) for GIST needs to be conceived. This extension of the NTLP layer is named GIST Extension for Hybrid On-path Off-Path Signaling (HyPath).

With this approach the communication between layers, NSLP [2] and NTLP [9], works without changing their specifications. Figure 4 describes an example of the NSIS protocol architecture using the HyPath MRM.

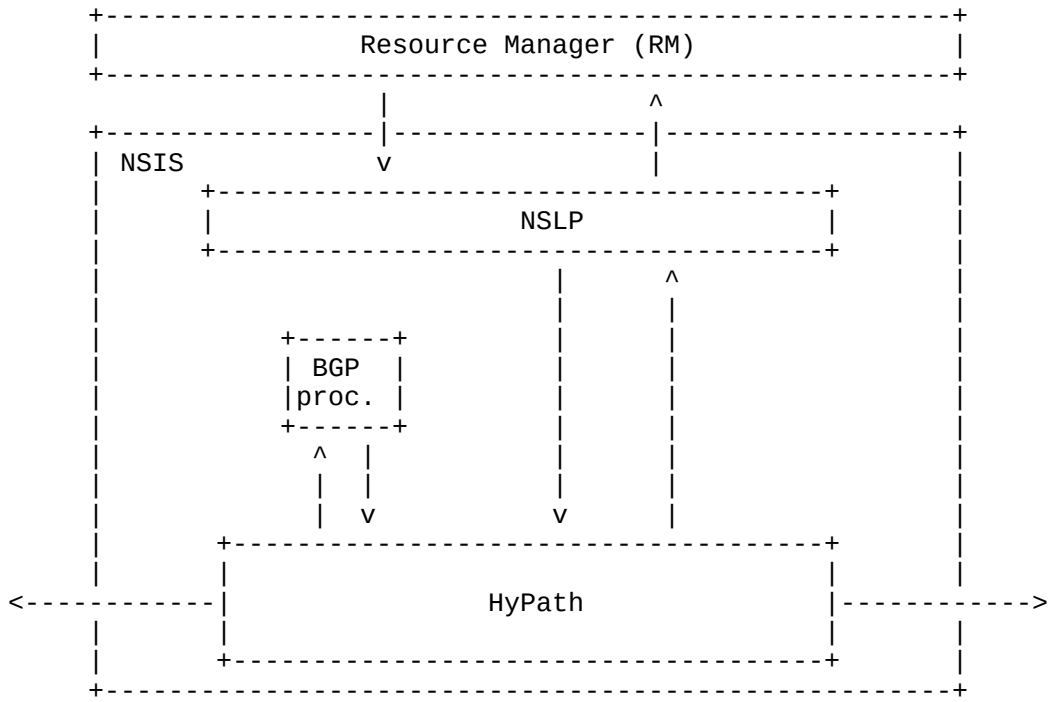

HyPath on the NSIS architecture
BGP proc = BGP processing + external functions

Figure 4: NSIS architecture with HyPath

In order to fulfill these e2e requirements the GIST processing MUST include the processing of this new MRM.

3.1. HyPath signaling

HyPath signaling provides a solution for the hybrid on-path off-path signaling. To better understand the new signaling Figure 5 illustrates the NSIS protocol with HyPath.

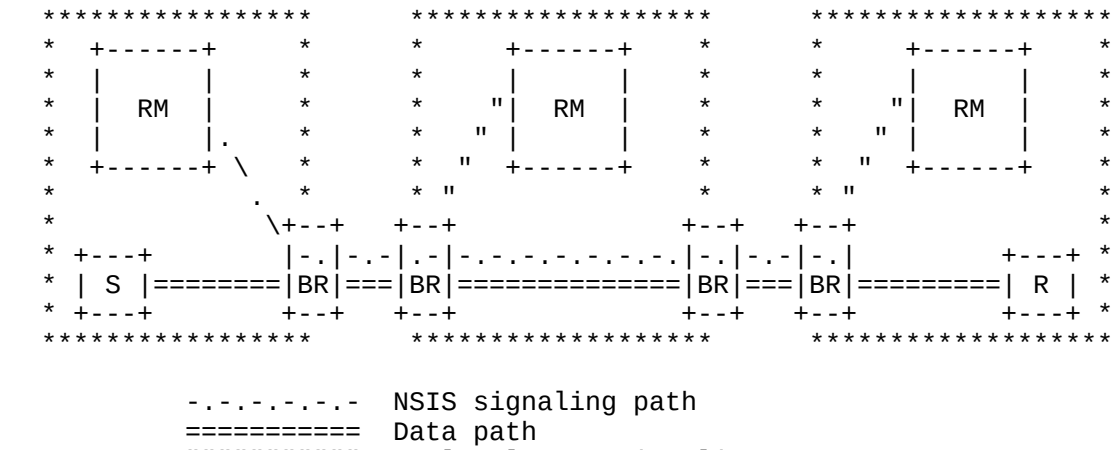

Figure 5: NSIS signaling with HyPath

When a user makes a QoS request to its local QoS system, NSIS signaling MUST occur in order to signal all RM's in the path . This signaling MUST follow the same path as the data. Therefore, in the first domain, in the local RM, HyPath MUST use an external function (as described in Section 3.3) to discover the local egress border router of the data.

Afterwards, HyPath sends a message to the egress border router. This message contains the normal NSLP payload and the HyPath MRM. This MRM includes additional information to fulfill the e2e signaling requirements (as described in Section 3.5).

Once in the egress border router, the NSIS signaling message, with

the additional information in the HyPath MRM, is forwarded to the end user.

In this scenario, all border routers intercept NSIS messages. As a result, in the following domain the NSIS signaling message is intercepted by the ingress border router. In this router the message is redirected to the local RM in order to make the local RM signaling.

After processing the received message, the RM server continues the signaling by sending a message back to the ingress border router. In the ingress border router the NSIS message is reinserted in the data path and continues.

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

With this architecture all the requirements to achieve e2e network signaling are met and no changes are needed in the definitions of the NSLP. The changes needed on the GIST layer are detailed later.

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 and HyPath complient. 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 (described in more detail in Section 3.4) 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.

The drawback of the approach described in Section 3.1 is that all border routers of all domains MUST be NSIS aware and HyPath compliant. Even though, in theory, this is a reasonable assumption, in practice we cannot guarantee that this happens. For this reason we define a heterogeneous solution (described in more detail in Section 3.4 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 signaling message, 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 the HyPath MRM (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 signaling 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 disadvantage 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. This information is obtained from the normal AS association procedure.

3.3. Usage of external routing protocols

The motivation for off-path signaling was described in the beginning of this document. If the signaling is decoupled from the data path (but still path-related signaling) 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 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 of the next domain MUST be identified in order to propagate the request.

The RM has access to the routing 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 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 routers or Router Reflectors [10].

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

- Discover the ingress border router;

- Discover the egress border router.

The discovery process of the ingress border router depends on the type of the actual and upstream domains, as described next.

If the actual and the upstream domains are NSIS domains, the ingress router is easy to retrieve. It is the border router that intercepts the NSIS message and redirects it to the RM.

If the actual domain is 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 3.5.1. The upstream RM interrogates the BGP table of its ingress border router and retrieves the address of the next domain.

Relying on the intra-domain routing (topology, traffic engineering) and BGP interactions 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 contain:

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 signaling 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 the new request must be addressed to the egress border router.

When this RM receives the message, it already has the IP address of the ingress border router. If this address is not the loopback address (a private IP address, for instance), it can obtain the

loopback address b from local BGP and topology configuration.

Dealing with a non-NSIS domain, is a similar procedure to the one presented when dealing with a NSIS domain.

All border routers communicate in the i-BGP session in order to discover the egress border router inside an AS. The egress border router is discovered using the BGP routing 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 neighbor. 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 the cartography of the domain (network topology) is stored.

The RM may obtain the IP address of a peer RM in an adjacent domain through the Service Level Agreement (SLA) between two adjacent domains which contains the loopback 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 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). 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 signaling 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 some disadvantages. On one hand, NSIS domains need to use HyPath in the border routers. On the other hand, in the non-NSIS domain it is needed an intensive

usage of external functions that extensively access the routing protocol.

A new solution is to integrate the two approaches 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 (between NSIS and non-NSIS domains). If the NSIS approach is used when the border router sends the message to the destination (through the data path) the signaling message is never intercepted by the next domain. Therefore, the next domain RM is never signaled.

To solve this problem, the NSIS domain MUST check the type of the next domain before sending any signaling 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 3.1, otherwise it is sent as described in Section 3.2.

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. NSIS architecture with HyPath

As described in the previous sections, with the usage of HyPath for the hybrid on-path off-path approach for e2e signaling across NSIS and non-NSIS domains requires that the RMs and the routers, especially the border routers, support HyPath.

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 signaling
  - Message reception and decoding
  - Sending messages

- In the egress border router
  - Start RM signaling 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 signaling

These functionalities are described in more detail in the next subsections.

3.5.1. HyPath on the Resource Manager

The MRM HyPath in the RMs is responsible for changing the destination address of the signaling message in order to assure that the right RM is signaled.

In the first domain (the domain where the network signaling starts) HyPath discovers the egress border router on 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 the NSIS message has already been received and the ingress border router information is included in the HyPath MRM. If the next domain and the current domain are NSIS domains, the message is sent to the ingress border router (IP address in the MRM) 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 use the normal GIST state to allow the usage of already established states and associations.

In the RM, HyPath requires a RM state table to store the previous RM IP address when the previous or the local domains are non-NSIS domains. In all other cases HyPath does not require any state.

3.5.2. HyPath in the Border Router

HyPath 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 signaling merges with the data path. From this point forward, if the message is always sent to the end user, the signaling path will follow the same path as the data path. In the other domains, the border router acts as an ingress border router where NSIS messages are intercepted.

In the border router, if the messages are received from the local RM (HyPath messages), 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 and the current network is neither the source nor the destination network, the message direction MUST be set with the original direction field of the HyPath MRM. The Border Router Address field of the HyPath MRM MUST be set with the address of the current machine and the border router flag set.

If the received message is to be sent to the final destination and if the border router flag is set, the Border Router Address field in the HyPath MRM 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 source address MUST be set to the BR address, the destination address MUST be set to the local RM address and the Border Router Address field MUST be updated with the local BR address.

3.5.3. The HyPath MRM

The current MRMs [9] do not solve the problems describe in the previous sections.

There are two possible approaches to adapt GIST with the previously described functionalities through the MRMs:

- A new MRM;
- A Path-Coupled MRM extension.

The next sub-sections describe these two approaches. The last sub-section presents the advantages and disadvantages of the two approaches.

3.5.3.1. A new MRM

The creation of a new MRM, is one of the approaches to have an hybrid on-path off-path approach for e2e signaling with NSIS. The creation of a new GIST MRM adapts the GIST signaling with the described HyPath functionalities in an independent way of the others MRMs.

The HyPath MRM takes the following format.

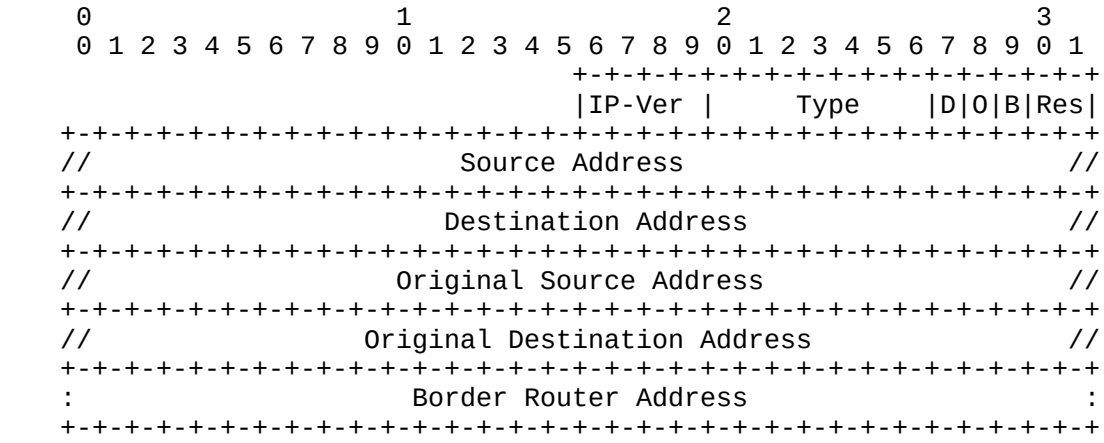

Type: type of HyPath message.

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 D: direction flag
Flag O: original direction flag
Flag B: means that the Border Router Address field is present

Source Address: data source address

Destination Address: data destination address

Original Source Address: original data source address

Original Destination Address: original data destination

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

This new MRM needs to be processed by GIST in a different way from the other MRMs. The Source and Destination Addresses are used by GIST to send the messages, but unlike most of the others MRMs these

two addresses are changed to perform the off-path signaling. The Original Source and Original Destination addresses are the data flow addresses that are kept unchanged. Taking into account the HyPath architecture, the Border Router Address is used to keep the border router address for the off-path signaling and to be used in the Source Address and Destination Address fields to perform the off-path signaling.

The HyPath MRM is also composed by three flags that can influence the GIST message processing. The D flag, like in other MRMs, indicates to GIST the direction of the message. The O flag, is used to store the D flag value whenever off-path is used in the signaling. If the B flag is set, HyPath message processing knows that the Border Router Address is present, otherwise this field is not checked.

3.5.3.2. Path-Coupled MRM Extension

This section describes another approach to extend the Path-Coupled MRM [9] with the new characteristics required to make an hybrid on-path off-path e2e signaling.

By adding additional fields to the current Path-Coupled MRM it is possible to maintain the usual path-coupled functionalities and, when required, have the additional HyPath functionalites. This can be achieved by adding a new flag, H Flag, in the reserved area of the MRM (after the already specified D Flag). NSIS entities that are configured to process HyPath messages (usually RM and border routers) can search for this H Flag in the Path-Coupled MRM, while other entities can simply ignore it and process the message normally (ignoring all extra HyPath fields in the MRM).

When the H Flag is present, new fields MUST be included on the MRM. These fields consist of a message type, followed by a flag stating if the border router address field exists, a flag to depict the message real direction and a reserved field for future use. Next, there are two fields containing the source and destination addresses of the message. Finally, there is an optional field containing the address of the border router of the domain.

The Path-Coupled MRM with the HyPath extension to support hybrid on-path off-path signaling takes the following format.

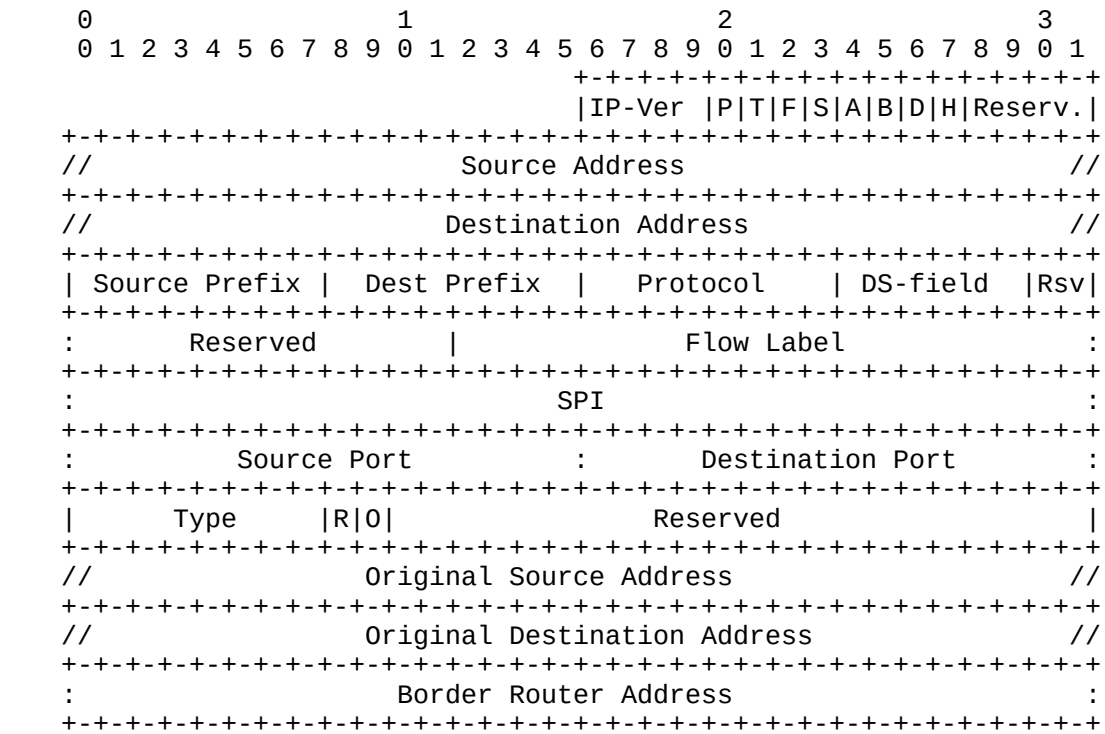

Type: type of HyPath message.

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 H: means that the HyPath information is present
Flag R: means that the Border Router Address field is present
Flag O: original direction flag

Original Source Address: original data source address

Original Destination Address: original data destination

Border Router Address: border router address to where the resource

manager MUST send the next message

3.5.3.3. Conclusion

The two solutions proposed previously (the new MRM and the Path-Coupled MRM extension) are able to provide a hybrid on-path off-path approach for e2e signaling. Both of these approaches allow the communication between the two NSIS layers, NSLP [2] and NTLP [9], without changing their specifications. The GIST changes to allow these solutions are few and SHOULD not influence the other functionalities.

With the new HyPath MRM approach there is no dependency of other MRMs and there is no need to change their processing in GIST. On the other hand this solution is only feasible if all NSIS Entities support HyPath MRM.

With the Path-Coupled MRM Extension approach, the NSIS entities that are not configured to process HyPath messages, can simply forward the message like the already defined Path-Coupled MRM (ignoring the additional fields in the MRM). Figure 8 shows an example of this scenario. In this scenario the Border routers "BR1" and "BR2" have HyPath enabled. The interior routers, "R1" and "R2", are NSIS routers with HyPath disabled (or even not supported). When the message is received by "R1", it is forwarded to "R2", using the normal Path-Coupled GIST processing.

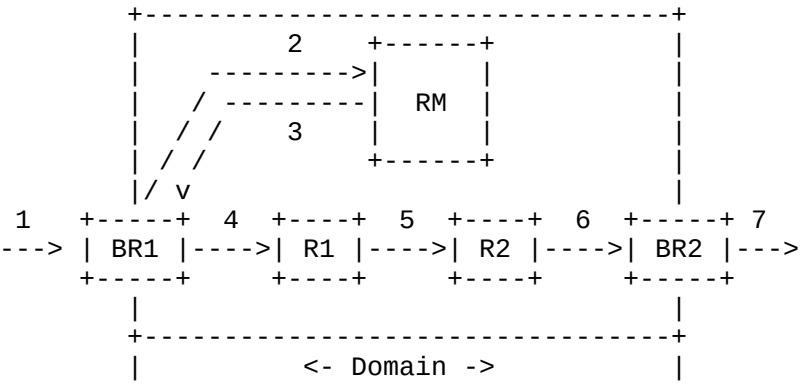

Figure 8: Non HyPath aware example

With the new HyPath MRM approach if on the signaling path exist, NSIS entities with no HyPath support the messages are discarded. The

Path-Coupled Extension approach solves this problem.

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

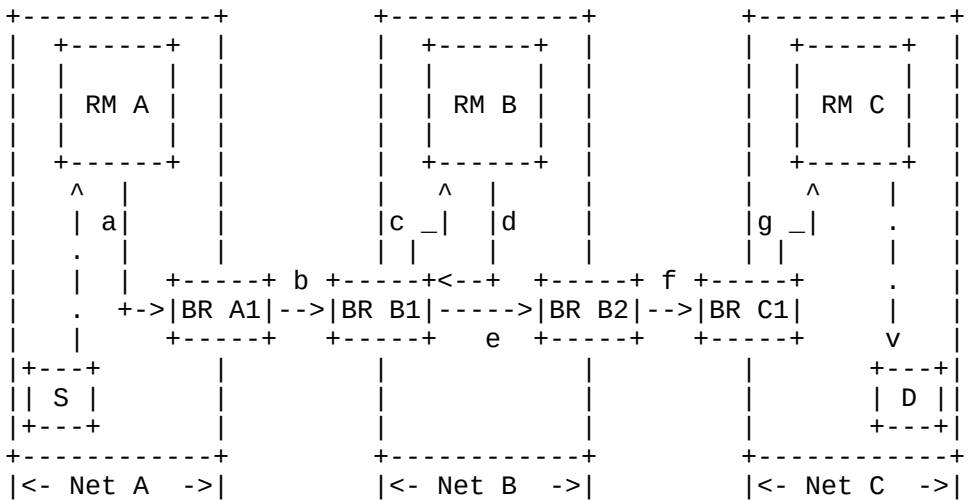

-.-> Protocol between users and the RM
---> NSIS messages with HyPath

Figure 9: HyPath architecture in the border routers

HyPath signaling starts in the RM A, where the NSLP requests HyPath to send a message from S to D. Before starting the signaling, 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.

When the BR A1 receives the message, knows it MUST be injected in the data path. BR A1 changes the received MRM with the original information available there (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) the message is then forwarded to the local RM, the RM B. The MRM information 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 it to the respective NSLP. If NSLP sends a response message, it is sent back to the respective BR, the BR B1. Before sending the message, the destination IP address MUST be changed with BR IP address information. One message that has no border router field set (in case it is the first message) 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 MUST 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 signaling stops and other type of signaling (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. Acknowledgements

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

5. Security Considerations

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

6. Conclusion

This draft presented an architecture in the context of a NSIS multi domain Internet that aims an off-path signaling 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, 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 proposes a solution called Hybrid Path, to involve the central entities on the signaling in the NSIS context.

Two solutions for hybrid on-path off-path signaling are presented in this draft. One of these solutions proposes the creation of a new MRM for GIST. On the other hand, an extension to the existing Path-Coupled MRM is also a way to solve this problem. Both SHOULD have a minimal impact on the NSIS architecture. A choice between these approaches MUST be made.

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

7. Open issues

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

8. Normative References

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

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

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

[4] Couturier, A. and O. Schelen, "On path and off path signaling for NSIS", draft-schelen-nsis-opopsig-00 (work in progress), June 2002.

[5] Hancock, R., "A Problem Statement for Partly-Decoupled Signalling in NSIS", draft-hancock-nsis-pds-problem-04 (work in progress), October 2006.

[6] Trang, N., Boukhatem, N., Doudane, Y., and G. Pujolle, "COPS-SLS - a service level negotiation protocol fort he internet", Editor IEEE Communication Magazine, vol. 40, no5, May 2002.

[7] Trang, N., Boukhatem, N., Doudane, Y., and G. Pujolle, "COPS-SLS usage for dynamic policy-based QoS management over heterogeneous IP networks", Editor IEEE Communication Magazine, vol. 17, no3, May 2003.

[8] Durham, D., Boyle, J., Cohen, R., Herzog, S., Rajan, R., and A. Sastry, "The COPS (Common Open Policy Service) Protocol", RFC 2748, January 2000.

[9] Schulzrinne, H. and R. Hancock, "GIST: General Internet

Signalling Transport", draft-ietf-nsis-ntlp-12 (work in progress), March 2007.

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

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

[12] Farrel, A., "A Path Computation Element (PCE) Based Architecture", draft-ietf-pce-architecture-05 (work in progress), April 2006.

[13] Manner, J., "NSLP for Quality-of-Service Signaling", draft-ietf-nsis-qos-nslp-12 (work in progress), October 2006.

[14] Rosenberg, J., Schulzrinne, H., Camarillo, G., Johnston, A., Peterson, J., Sparks, R., Handley, M., and E. Schooler, "SIP: Session Initiation Protocol", RFC 3261, June 2002.

Authors' Addresses

Luis Cordeiro
University of Coimbra
Polo II - Pinhal de Marrocos
Coimbra 3030-290
Portugal

Email: cordeiro@dei.uc.pt

Marilia Curado
University of Coimbra
Polo II - Pinhal de Marrocos
Coimbra 3030-290
Portugal

Email: marilia@dei.uc.pt

Edmundo Monteiro
University of Coimbra
Polo II - Pinhal de Marrocos
Coimbra 3030-290
Portugal

Email: edmundo@dei.uc.pt

Vitor Bernardo
University of Coimbra
Polo II - Pinhal de Marrocos
Coimbra 3030-290
Portugal

Email: vmbern@student.dei.uc.pt

Florin Racaru
LAAS
Avenue du Colonel Roche
Toulouse 31077
France

Email: fracaru@laas.fr

Michel Diaz
LAAS
Avenue du Colonel Roche
Toulouse 31077
France

Email: diaz@laas.fr

Christophe Chassot
LAAS
Avenue du Colonel Roche
Toulouse 31077
France

Email: chassot@laas.fr

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).