Next Steps in Signaling (nsis)
Internet-Draft
Expires: April 27, 2006

H. Cheng
J. Huang
T. Sanda
T. Ue
Panasonic

October 24, 2005

NSIS Flow ID and packet classification issues
draft-cheng-nsis-flowid-issues-02.txt

Status of this Memo

Copyright Notice

Abstract

In current NSIS signaling, there are two main functions depending on Flow ID, i.e. signaling message routing, data packet classification. Specifically, the same information carried by the MRI is also used to derive the packet classification at NSLP layer. This arrangement assumes that identical information is required by the two functions at two different layers, and thus has limitations. With the

introduction of NSIS applications in more complicated scenarios, such assumption can no longer hold. Therefore, keeping the dependency between the two functions hinders the development of NSIS. Efforts have been made in different NSLP layer applications to extend the relationship, e.g. QoS NSLP. This draft studied the possibility of disjoining the information for the two functions. Problems faced by the current system and different remedy options are discussed. With these details, it is intended to help the reader to evaluate the feasibility of redefining the packet classification information signaling in NSIS.

Table of Contents

1. Conventions used in this document

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 RFC2119 [1].

2. Introduction

With the NSIS framework design [2], signaling message routing and data flow classification are two important functions at two layers, i.e. NTLP and NSLP respectively. However, information to support the two functions is derived from the same element, Flow ID. This type of arrangement has limitation on the NSIS application, and creates problem when deployment setting is not as assumed. Therefore, this aspect of the design needs to be reviewed and improved.

The linkage between the two functions is based on the path-coupled signaling principle of the framework. In a path coupled MRM, the signaling message is suppose to go through the exact path of the data flow. To this end, NTLP layer protocol needs to use a MRI resembling the flow identification as much as possible. In the GIST draft [3], the MRI has been set as the Flow ID for the path-coupled MRM (or a simplified version of Flow ID for loose-end MRM). This essentially sets the information for the message routing and flow identification (packet classification) to be identical. The MRI is passed from the NSLP through a set of API to the NTLP layer when a NSLP message is to be sent, and passed from NTLP layer up to NSLP layer when a NSLP message is received. Due to the assumption that flow identification is always the same as MRI, NSLP applications, e.g. QoS NSLP [4], no longer has an element to carry the packet classification information. Thus, MRI is used by the NSLP layer to derive the packet classification information.

Obviously, such an arrangement brings problem when the routing information does not synchronize with the flow identification information. There are two possible cases, i.e. when the flow identification (classification) information has a wider scope than that used for routing, and when the flow identification information has a more specific scope than that use for routing. Section 4 discusses in detail of the scenarios where problems exist.

In an effort to bridge the difference between MRI and flow identification information, some NSLP applications introduced new objects. For example, QoS NSLP [4] included a PACKET_CLASSIFIER object. However, since the new object still heavily relies on the MRI to derive the packet classification information, it does not solve the problem.

Rather than trying objects to extend MRI, this draft investigate the possibility of using a disjoint object at NSLP layer to carry the packet classification information. With an extensible format, the new object can easily support multiple address pairs. This helps in a multihoming environment, and makes it possible to signal for flows

not specified by the MRI. Different aspects of introducing such an object are discussed in the draft, e.g. place of the object, NAT traversal issues, etc. With the recent discussion in the mailing list regarding packet classification, it is the intension of the draft to provide views on relevant issues to help the WG to reach consensus on the topic.

3. Terminology

The Terminology defined in [2], [3] and [4] applies to this draft. In addition, the following terms are used:

Filter List: A list of address attributes that can be used for classifying the packets associated with a particular flow.

Flow Identification and Data Packet Classification are used interchangeably in this draft.

4. Problem analysis for Flow ID usage

In the existing NSIS Working Group drafts, the Flow ID is defined as an identifier that "provide enough information such that the signaling flow receives the same treatment along the data path as the actual data itself" [2]. It is also identified that information to be used for the identifier includes Source IP address, Destination IP address, protocol identifier and higher layer (port) addressing, flow label, SPI filed of IPsec data, and DSCP/TOS filed. The reason for placing the Flow ID in the NTLP layer is to allow visibility and modification at the address boundaries [2]. There are multiple usages of Flow ID described in NSIS drafts.

In the GIST document [3], for a path-coupled MRM, the Flow ID is utilized as the Message-Routing-Information (MRI). The MRI is then required to be set at NSLP message sender. A formation of the Flow ID is provided as:

```
Flow-Identifier = Network-layer-version

                  Source-address prefix-length

                  Destination-address prefix-length

                  IP-protocol

                  Traffic-class

                  [flow-label]

                  [ipsec-SPI/L4-ports]
```

In the signaling applications, it is usually necessary to have information about the packet classification for the enforcement. Based on the description of the signaling applications drafts, this information has to be derived from the MRI. For example, in the QoS-NSLP draft [4], the packet classification information is derived from the MRI fields indicated by the PACKET_CLASSIFIER object. Therefore, it is assumed that the QNE would set the Packet Classifier in its Traffic Control module based on a subset of the MRI information.

In the QoS-NSLP draft [4], it is also mentioned that the Flow ID is used for the signaling state management. For example, it will help the QNE to detect a mobility event.

It is obvious from the above that the Flow ID is heavily depended on in the NSIS signaling, and therefore has multiple functions associated. This kind of overloading of the Flow ID works fine with

a simplified static network environment. However, with scales of the network growing and complexities increasing, adverse impacts are introduced to NSIS, with the Flow ID striving to meet the requirements of all the functions. Following sections provide detail analysis of three examples where Flow ID of current NSIS design may not function properly.

4.1. Multiple addresses involved session

According to the definition of the data flow, it is possible that multiple addresses are involved. The multiple addresses could be different IP addresses, different port numbers, or different higher layer protocol types, etc. With the development of the networking technology, it is becoming common to use multiple streams sessions. These streams may make use of different addresses, but could still goes through the same data path. Multihoming, multi-threading, and aggregation, etc., are the possible reasons that require different addresses to be signaled over one session. Below are some of the detail examples:

- Edge to edge signaling is a case that may involve multiple addresses in the flow signaled. For example, if several flows are aggregated at a domain entry point, the signaling sent for the aggregated flow would contain several sets of addressing information.

- The multiple addresses could also be the different higher protocol addresses, e.g. several port numbers used in a multiple thread session. This type of multi-thread method is widely used in some popular FTP clients. Usually, a download session would be limited by the server resource allocation and network condition. The throughput achieved is thus also limited. To boost the downloading speed, the FTP clients establish multiple connections with the server, and download the file simultaneously. This would achieve much higher throughput. In the process, the terminal device will use multiple ports in the communication for the multiple connections. Since all these concurrent connections belong to the same session, it is better for NSIS to signal all the port numbers in one single signaling message.

With the current definition of MRI, there is no way for the multiple addresses to be accurately represented. As quoted above, the MRI format only allows a simple prefix for the IP (source and destination) address fields. This may suit for the routing, e.g. to represent a group of address within a subnet. It is not sufficient for signaling the packet classification. For example, with two source addresses belong to the same subnet, the two data flows could be routed along the same path. However, if the source addresses are wildcarded and fed to the NSLP (eventually RMF) for enforcement,

every node in that subnet connecting to the same peer (destination address) would be classified into the same session. This is obviously not acceptable for the QoS enforcement. Other than the IP address prefix, current MRI format does not even provide any wildcarding or masking tools for other fields. Therefore, it is not possible to represent multiple address information of other fields, e.g. port number, traffic class, etc.

The latest QoS NSLP draft [4] introduced a PACKET_CLASSIFIER object to help deriving the packet classification information. However, due to its simple format, and heavy dependency on the existing MRI, it cannot solve the above problem. For the path-coupled routing MRM, the PACKET_CLASSIFIER data is just a two byte field with 8 bits utilized. It can only simply indicate which field of the MRI should be considered in deriving the packet classifier. Obviously, if the MRI cannot accurately represent the multiple address flow as described above, the PACKET_CLASSIFIER does not help either. The packet classifier constructed would be either too wide, i.e. when one or more fields are omitted, or too narrow, i.e. when every field is used.

To solve this problem, a method for signal accurate flow classification information needs to be provided. Different options are discussed in detail in section 5.1.

4.2. Predictive signaling scenario

Predictive routing support in NSIS is mentioned in [3], where the signaling is sent along the path that the data flow may or will follow in the future. The use of predictive routing is crucial for make-before-break reservation on predictive paths in mobility scenarios. Make-before-break is necessary for minimum QoS interruption at the time of handover [5], since performing NSIS signaling after the MN's actual handover causes service interruption due to the delay in signaling path establishment.

Example procedure of make-before-break is as following: when the MN intends to handover, it obtains the information of new subnetwork and a suitable proxy NE on the predictive path, such as new Access Router (nAR), by using proper mobility protocols e.g. FMIP [10] or CARD [11]; then the MN asks the nAR to perform NSIS signaling along predictive route. At this stage nAR may or may not obtain MN's nCoA which is necessary for generating packet classifier for the predictive path.

Even if the nAR doesn't have MN's nCoA, it is still desirable that the nAR can start the signaling, e.g. establish the NSIS signaling path for the MN, so that when the MN moves into the new network, the

QoS path is ready for use with least signaling. For the path coupled MRM, it is obvious that the actual data flow will go pass the nAR. Therefore, it is not a problem for the nAR to act as a proxy for the MN to pre-establish the signaling path, e.g. send out Query messages to find out the QNEs, and check out the resources available over the predictive path. However, with the current definition of NSIS, e.g. QoS NSLP [4], the Flow ID (MRI) in use will also dedicate the packet classifier for the flow, i.e. the PACKET_CLASSIFIER indicates which field of MRI to be used for packet classifier. Therefore, the nAR would not be able to generate a proper Flow ID for this signaling without knowing the MN's nCoA. If the nAR uses an arbitrary address to create the Flow ID, the whole signaling process needs to be repeated when the MN's nCoA is assigned. This means nAR cannot send any signaling message until obtaining MN's NCoA. This renders the usage of make-before-break limited.

Obviously, this problem calls for a possibility of separating the message routing information from the actual packet classification information. Different options to achieve that and considerations are discussed in section 5.2.

4.3. Packet classification information manipulations

In certain cases, the address information, e.g. port number, may change during the process of a flow. According to the current definition of Flow ID and MRI, it may result in the change of the Flow ID (which may trigger further NSIS signaling procedures).

An example of the address varying application process is the establishment of a H.323 session [12]. Protocol like H.323 establishes the session progressively with the help of different auxiliary protocols. Figure 1 shows a typical example of a H.323 session establishment process.

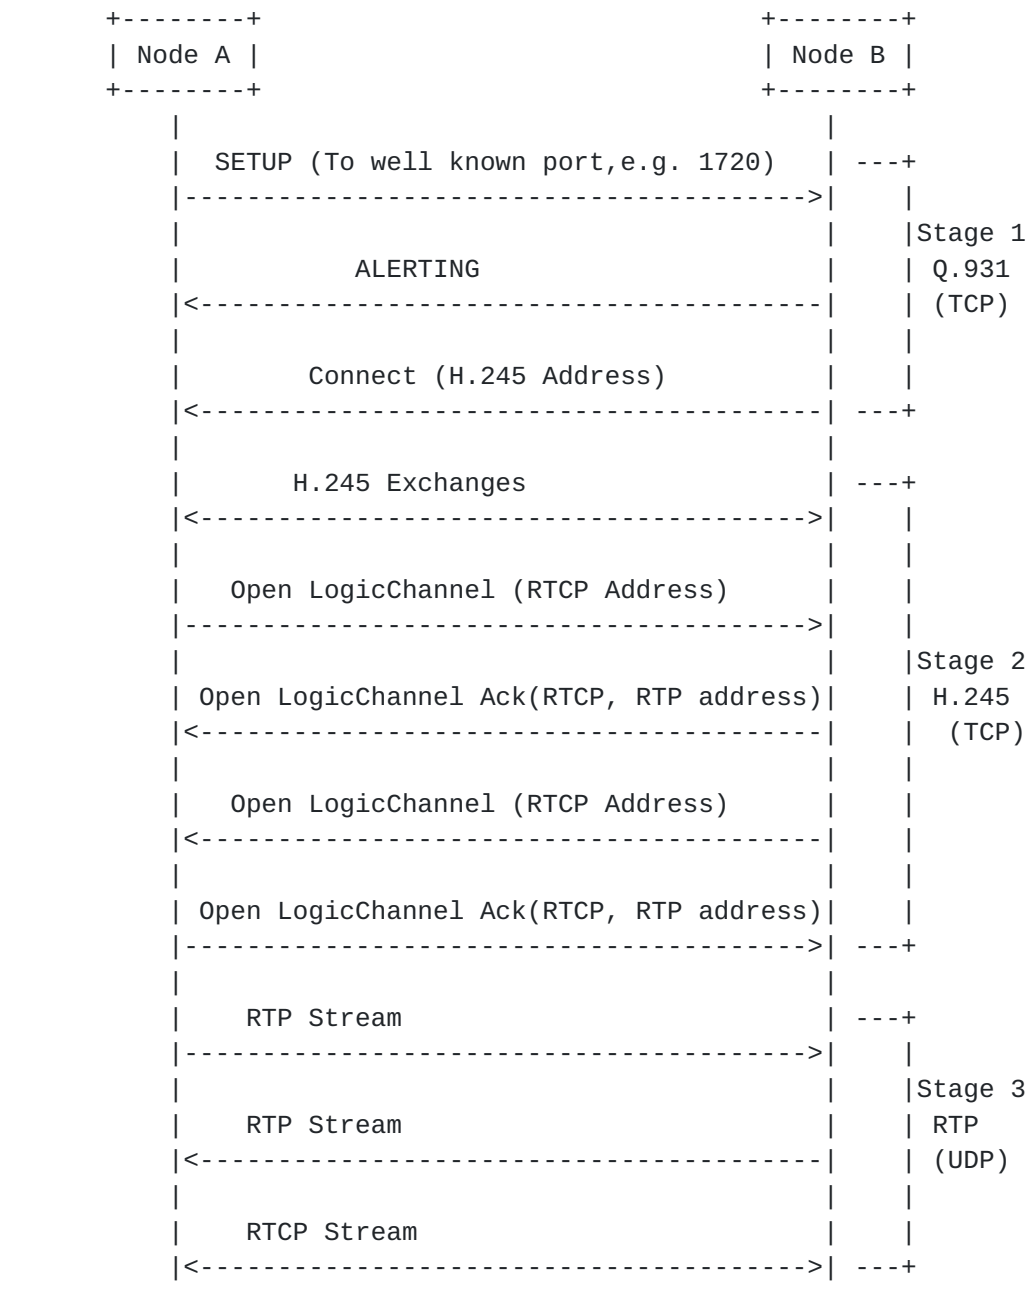

Figure 1. An example session establishment process using H.323

In the above example, the actual RTP session will only be initiated at stage 3, which means the port address may change several times before the first RTP data packet is sent out. However, the NSIS messaging needs to start in Stage 1, e.g. to open the firewall pinholes. The problem created from this is that when the address information changes, the Flow ID needs to be changed accordingly, so that NSIS aware nodes along the path can obtain the correct packet classification information. Otherwise, the local enforcer, e.g. the

RMF, could not perform its function correctly.

However, the change in the Flow ID means another round of NTLP layer path establishment according to GIST draft [3]. The reason is that the primary key of the message routing state is the combination of MRI, SID and NSLPID. Change of Flow ID means a primary key change and the message routing state has to be re-established. Besides this, in the process mentioned above, there could be multiple ports opened at the same time, and different transport protocol utilized. All these call for a simple way of changing the packet classification information along the signaling path without affecting the signaling state information.

Another type of scenarios that requires manipulation of the packet classification information of a session is the tunneling case, as described in [13]. In such cases, the flow binding requires the packet classification information to be merged, and modified during the lifetime of the session.

All these call for a simple way of changing the packet classification information along the signaling path without affecting the signaling state information. Section 5.3 provides a detail analysis of the options to deal with this requirement.

5. Separation of MRI and Packet Classification Information

In view of all the problems of the Flow ID in supporting scenarios described in previous section, there are three major issues to be resolved:

- NSIS signaling should allow accurate indication of multiple addresses information in the signaling of packet classification information.

- NSIS signaling should allow a separation of the message routing information and packet classification information when necessary. This will facilitates the NSIS signaling from non-end hosts.

- NSIS signaling should allow modification of the packet classification information without affecting the signaling and routing state.

There are different ways to achieve the above goals. Options and considerations are discussed in detail in the following subsections. With the discussion, the authors come to a conclusion that utilizing a separate object for carrying the packet classification information serves the above purpose and is a good trade off between different considerations.

The suggested object could take different form. Below is an example format for it:

Filter List ::= <List Length> <Action> <* Filter Spec>;

where the "Filter Spec" is basically a packet classifier that provides packet classification information. For example, the MF Classifier defined in RFC2475 [14] could be a good candidate. List Length indicates the number of Filter Spec included, and the Action indicates how to treat the previous filters of the same Flow ID. Example actions to be taken are: ADD, SUB, and REPLACE.

Following subsections provide detail discussion of the considerations for different scenarios.

5.1. Considerations with multiple addresses involved sessions

5.1.1. Different options for multiple addresses support

One obvious way to support the multiple address involved communication session is to use multiple signaling sessions for each set of address. However, this solution faces a couple of problems. One problem is the significant increase in signaling, processing, and

state management overheads. The increment of the overhead is in proportion to the number of address combinations. For example, when an application session makes use of five different ports on a MN, five separate NSIS sessions needs to be signaled. The other problem is that certain relationships have to be enforced between these signaling sessions. These sessions should share the same fate and probably share the resource reserved because they actually signal for the same application. This kind of relationship enforcement between signaling sessions may further complicate the signaling and state management logics.

Another possible way to carry the multiple address information is to signal multiple times, each with the same Session ID and a different Flow ID that accurately represents one of the address information in use. Obviously, this creates signaling overhead as well. Besides that, it may also cause difficulties in state management on NSIS nodes. For example, a QNE receiving several messages with different Flow IDs will not be able to decide the proper action. The REPLACE flag introduced in [4] provides a way to indicate the action to the QNEs. Nevertheless, it requires explicit signaling for the management of these different Flow IDs. For instance, if a mobility event happened, the state with the old Flow IDs should be replaced with a new set with new Flow IDs. Therefore, the signaling message needs to indicate state with which particular Flow IDs should be replaced, and which should be kept. This requires extra state management logic at the QNEs.

The root of the issues lies on the incapability of the MRI to carry multiple addresses information. Therefore, solution could be achieved by either modifying the MRI format to accommodate more complicated information or utilizing extra objects. For example, the MRI could add a masking/wildcarding tool for each of the fields.

Since the change of MRI format will affect numerous aspects of the NTLP layer implementation, this draft proposes to use a separate Filter List object in NSLP layer to carry data packet classification information. This way, the MRI/Flow ID is freed from the address dependency. It can then take any form or using a much relaxed criteria for masking, as long as it provides enough information to allow the NSIS nodes to route the message correctly. The Filter List is now carried in the NSLP layer, and therefore, should be set by the end nodes, which has the ultimate information about the application layer of the session. This means a faster and more accurate interaction between NSIS signaling and application becomes possible. Since the packet classifying is only meaningful to the NSLP application enforcer, e.g. the RMF in QoS NSLP case, having Filter List in the NSLP layer alleviates NTLP layer's burden on the processing of this information.

Following subsection describes deployment examples of using the new object.

5.1.2. Support of multiple addresses using Filer List

For the scenarios described in section 4.1, the use of Filter List solves the dilemma of using the MRI. As the Filter List allows accurate representation of the multiple addresses involved in the session, it does not require NTLP layer to care about the actual packet addresses. By doing this, the NTLP layer only needs to establish the correct signaling path. Simple wildcarding would then be useful to find out the common routing path. For example, for a session with different addresses, the NTLP layer can use a MRI with wildcarding without worrying the accuracy of the address presentation because the filter information is carried in the NSLP payload. The wildcarded MRI is only used for the route decision of the NSIS message, and the NSLP layer message with the Filter List can be configured by the corresponding entity, e.g. RMF, with correct settings. The introduction of the Filter List does not limit the NTLP layer's choice of Flow ID value. Instead, it gives NTLP more flexibility in selecting a proper Flow ID.

Usually the routing does not really require that all the fields of the MRI. For example, a normal router would only decide the route based on the destination address. Therefore most fields of the MRI could be set to a default value, e.g. 0x00. Since the state information stored at the NSIS node is indexed by both the Session ID and the MRI, the above means saving in the key length used. For example, in the current NSIS draft, the key for the MRI would need to at least include source address, destination address, etc, since it also is used for packet classifying. In the proposed scheme with Filter List, the key for the MRI only need to have the destination address if the routing is based on destination address only.

When a flow involves multiple addresses, it is possible to experience a divergence in the data path when the flow passes through a network section with load balancing supported. In this case, reservation on some of the split path may not be done correctly. There are several ways to address this issue. One possible method is to require the Load Balancing Initiator (LB-I) to duplicate the signaling message and send over every split path with the QSPEC adjusted accordingly. Another method is to force the LB-I to route all the packets identified by the Filter List to the same link. It is also possible for the QNI to be aware of the load balancing paths through the initial path discovery process, and construct the QSPEC and filter list accordingly.

5.2. Considerations with predictive signaling support

For the predictive signaling, one possible way to solve the problem is to pre-allocate an address for the MN for the predictive path. In this case, for example, the nAR can reserve an address for the MN when it proxy the signaling for it. When the MN actually moves to the new location, it must make use of the address pre-allocated by the nAR. For this option, there needs to have a mechanism that can guarantee the proper address for the MN can be pre-allocated. It infer a stateful address allocation scheme, since if stateless address formation is used, the nAR would not have the information to generate the new address. Another issue with this scheme is that it may cause waste in the address space. The MN's signaling over the predictive path could be simply a Query of the QoS level supported. MN may not eventually move to the new location. Therefore, requiring an allocation of an address for the MN for such a signaling decreases the utilization efficiency of the resources.

The use of a separate object in NSLP layer stands as another option to support the proxy predictive signaling. When the nAR does not know MN's nCoA, it generates a Flow ID with its own IP address, and sends signaling messages such as QUERY using it. Since the actual data flow has not started yet, actual packet classification information, i.e. Filter List, is not necessary. By sending the QUERY, path state between the nAR and CN is installed and QoS resource availability along the path is gathered as well. When the MN connects to the new subnetwork, the signaling path has already been established. The MN only has to update the state of the rest part of the new path and update the whole path with a RESERVE message containing Filter List. Since it is an update message, and all the signaling states are in place, it is much faster than a new RESERVE message with a different Flow ID.

Comparing the above two options, the use of separate packet classification object obviously provides a more flexible support for the predictive signaling.

5.3. Considerations for packet classification information manipulation

With the current NSIS Flow ID definition, the change of packet classification information means a change in the Flow ID, and thus requires another round of signaling. This usually results in a signaling state update on all the QNEs. Certain scenarios require a merge of the original and new flow states to achieve desired QoS enforcement results, e.g. as described in [13]. For this purpose, an intra-session association object, ASSOCIATE_FLOW_SESSION, is introduced in [13]. The new object helps to link different flows within a Session together, so that the packet classification

information could be manipulated on the QNEs.

Different from the approach taken by [13], the use of a separate Filter List object can achieve the similar purpose. Since the Filter List object introduced supports the Action indication, new Filter Spec could be progressively added when the new address information is available. For example, when a new port has been added to the communication session, the NSLP layer can issue a new message with the Filter List object carrying an Action field indicating "ADD". The receiving NSLP aware node could take this new port information and merge it with the existing packet classification information for the same Session ID and Flow ID.

Comparing the two options, the approach in [13] preserves the current MRI and Flow ID definition. It keeps the change in NSLP layer and slightly modifies the QNE behaviors. The disadvantage of this option is the signaling overhead introduced. Every signaling message needs to carry the new association object, which includes the additional Flow IDs. The other issue is the signaling using a new Flow ID when new packet classification information needs to be added. This may involves the path discovery procedure for the new Flow ID, which could be time consuming.

For the Filter List object option, the NSLP layer packet classification definition is modified. The classification information is no longer depending on the MRI. However, to signal for additional address information, the same Flow ID could be used. This means only an update of the signaling state on the QNE is necessary. Therefore, it can be faster.

5.4. Changes at the framework/NTLP layer

By introducing the Filter List object into the NSLP layer, the NTLP layer is relieved from signaling the actual data packet classification information. This looses its dependency on the data plane addressing information. Therefore, the design and operation of the NTLP layer has more flexibility. For example, the state pre-establishment and use of the proxy is much easier.

The current design of the Flow ID or the MRI can still be used in the NTLP layer for the message routing information signaling. However, it does not need to be bound by the actual addressing information of the data flow, e.g. the port number, etc. It does not require change in the current design of the GIST, but relaxed some of the requirements. Also, the state management can be simplified due to the above reasons. For example, the state information storage and retrieval can be speeded up as described in previous sections.

Since the NTLP does not need to modify the Flow ID or MRI every time the address changes, the state maintained on the NSIS nodes will be more stable, and thus requires less operation in updating.

5.5. Changes at the NSLP layer

Use of the Filter List in the NSLP layer requires some enhancement at the NSLP nodes. The new NSLP layer needs to be able to manage the packet classification information, e.g. the Filter List. It no longer depends on the NTLP layer to obtain the filter information, and therefore need to be able to process them. Most of the time, the Filter List should be passed directly to the local enforcer, e.g. the RMF in the QoS NSLP case together with the QSPEC. This means the NSLP layer does not need to include much extra functions.

For the operation of the scheme, the NSLP now need to have some interaction with the application or the addressing management. The NSLP application now needs to obtain all those addressing information from these entities to construct the Filter List at the end nodes. However, it does not mean the NSLP have to communicate with them directly. The current design could still be kept, so that the NSLP can communicate with them via NTLP using some extra APIs. It is for further study if it is better to have the NSLP interacts directly with the application layer.

6. Impact analysis

This section provides a detail analysis of the different aspects of introducing the Filter List object to the design of NSIS protocols. Some of the recent discussions in the mailing list on the topic are also covered.

6.1. Single object vs. separate objects

In the current NSIS framework, although the message routing and packet classification are two different functions, they uses the information derived from the same object, MRI. The reasoning for using a single object for the two functions is based on the address boundary traversing consideration. When a single object is used, the address boundary node, e.g. NAT, only need to modify this object. Information derived from the modified object could always be synchronized. Therefore, it is guaranteed that information used for the two functions are referring to the same flow.

A possibility of meeting the requirements listed in section 5 using a single object is to modify the MRI format. This way, the extra packet classification information introduced in the Filter List could be incorporated into MRI. However, it will affect almost every aspects of the current GIST operation, and therefore is not appropriate.

The use of a single object assumes that the signaled flow must be the same as that defined in MRI. However, in certain case, more flexibility is desirable. For example, the metering application [15], is expecting that some extra filter information could be carried in the NSLP layer other than the MRI.

When two disjoining objects are used for the two functions, the address boundary node needs to apply changes twice. In order to guarantee the synchronization of the information, exactly the same rules should be used on the two objects, e.g. same mapping used on the MRI should be used on Filter List. For normal address boundary nodes using a comparatively static rules, this should be easily achievable. No major issue could be envisaged, e.g. passing a NAT. If the address boundary node is using very dynamic rules for the address processing, problem may exist for the data flow, similar to the load balancing case.

6.2. Location of the packet classification object

When considering the location of the Filter List object, there are generally three options: in NTLP layer, in NSLP layer, or in QSPEC. These three options all have advantages and disadvantages. Following

brief discussions give an overview of the pros and cons. The considerations are focusing on the signaling overhead and NAT traversing.

6.2.1. Object placed in NTLP layer

Firstly, when the Filter List is used as NTLP layer object, it introduces unnecessary burden to the NTLP. Placing it in the NTLP layer means that the object needs to be included in every NSIS message. However, the packet classification information is only meaningful to the RMF, and is only necessary in certain NSLP messages, e.g. QoS NSLP RESERVE. Therefore, unnecessary signaling overhead is resulted.

When the Filer List is placed in the NTLP layer, it is the easiest to achieve NAT traversal support of the signaling. The NAT node only needs to be NTLP layer aware to process the information. Current NAT behavior specified in the GIST draft [4] will still be usable.

6.2.2. Object placed in NSLP layer

When the Filter List object is placed in the NSLP layer, the NSLP application knows when it should be included in the signaling. For example, a QoS NSLP NOTIFY message may not include the packet classification information object. This helps to reduce the overhead of the signaling compared with the case where it is placed in NTLP layer.

However, the packet classification object needs to be made stackable to support stackable QSPEC. The main reason for this is due to the direct relationship between the packet classification information and the specific QoS Model. Therefore, when the QSPECs are stacked, the corresponding packet classification information also needs to be preserved so that it could be referred correctly when the QSPECs are popped.

The NAT traversal aspect of this option is slightly complicated than the NTLP layer based choice. Basically, in this case, the NAT is required to be able to process the common NSLP layer object other than being NTLP aware. The Filter List object could be placed in a fixed position in NSLP layer when presented, e.g. with a flag indication. Then, the NAT is able to apply the same address mapping rule on the Filter List, similar to what it does on the MRI. It is also possible for the NAT to insert the NAT-Object that contains the mapping rules, so that the next hop NSLP aware node, e.g. QNE, could actually perform the modification to the Filter List.

6.2.3. Object placed in QSPEC

The last choice is to place the Filter List in the QSPEC. This seems logical, since the packet classification would be utilized directly by the RMF, which is decided by the specific QoS Model. In this case, the QoS Model could optimize the information to be carried in the Filter List.

However, this also means that the Filter List object is only accessible for a node with knowledge of the QoS Model of concern. Therefore, it will be extremely difficult for the signaling message to traverse an address boundary, e.g. a NAT. For this option to work in a NATed network, the NAT node must be the specific QoS Model aware, which is a hard requirement to meet.

6.3. General packet classifier vs. model specific classifier

As mentioned in section 6.2.3, the packet classification information is relevant to the QoS Model (eventually RMF) in use. Therefore, different QoS Model may have different requirements on packet classification format. For example, intra domain DiffServ signaling may only require DSCP information, whereas edge to edge signaling may require MF classifier. Therefore, there are options to have the packet classification defined according to the QoS Model or kept general.

The choice of the packet classifier depends on where the object is placed. In the case of section 6.2.1 and 6.2.2 (i.e. in NTLP layer or NSLP layer), the packet classifier needs to be general, because it is common for all the QoS Models. For the general packet classifier, the MF Classifier defined in the RFC2475 [14] is a good candidate.

Only in the case of section 6.2.3 (i.e. in QSPEC), the packet classifier could be made using fields that is required by the specific QoS model associated. In order to achieve this, the QSPEC needs to define a set of basic elements, so that, when needed, QoS Model specific packet classifier could be formed using the basic elements.

7. Possible optimization with the PACKET_CLASSIFIER object

As discussed in the previous section, the Filter List may not be necessary for all the scenarios. For those static and simple addressing cases, the conventional MRI way is sufficient. Therefore, to optimize the signaling efficiency, the Filter List object could be made optional with a flag set at the signaling header. Below is an example of using an indication derived from the PACKET_CLASSIFIER object in the QoS NSLP case.

For example, the modified method specific data format of the PACKET_CLASSIFIER object for a path-coupled routing MRM, as in section 5.1.3.5 of QoS NSLP [4], is specified as following:

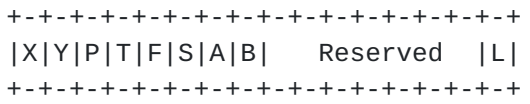

Wherein, the new flag L is to indicate if a separate Filter List is used in the NSLP for packet classification information. When the flag L is set, all the other fields should be set to zero.

Therefore, a QNE first checks if the L bit of the PACKET_CLASSIFIER object is set. If it is not set, the QNE will make use of the MRI and other fields of the PACKET_CLASSIFIER object to derive the packet classifier.

When the QNE finds out that the L flag is set, it will obtain the packet classification information directly from the Filter List object instead of the MRI.

8. Security Considerations

Major security issues for NSIS are addressed in [6], where the Section 4.4 mentions use of a Flow ID without source and destination IP addresses. If a Flow ID is used for traffic classification of data packets as well, then identity spoofing and injecting traffic is much easier since a packet only needs to be marked and an adversary can use a nearly arbitrary endpoint identifier to achieve the desired result.

The filter List method described in this draft allows the separation of the Flow ID and packet classification information. Different usage of the Flow can be employed, e.g. as described in [6]. Traffic classifier for data packets, however, may still use the conventional Flow ID information as a filter so that threat does not increase by using a Filter List.

9. Conclusion

This draft discussed issues faced by the NSIS design regarding its use of the Flow ID for carrying data packet classification information. A solution to the problems is proposed by introducing a Filter List object in the NSLP layer. This solution provides support for the scenarios that is not supported by the current NSIS framework. It also requires little change to the current NSIS design. Detail analysis of the impact to different components of the NSIS framework, NTLP and NSLP is provided, and it shows that the proposed solution is effective and easy to be incorporated into the current NSIS system.

10. Acknowledgements

This section contains the acknowledgements.

11. References

11.1. Normative Reference

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

[2] Hancock, R., "Next Steps in Signaling (NSIS): Framework", RFC 4080, Informational , June 2005.

[3] Schulzrinne, H. and R. Hancock, "GIST: General Internet Signaling Transport", Internet Draft draft-ietf-nsis-ntlp-08, Work in progress , September 2005.

[4] Manner, J., "NSLP for Quality-of-Service Signaling", Internet Draft draft-ietf-nsis-qos-nslp-08, Work in progress , October 2005.

[5] Chaskar, H., "Requirements of a Quality of Service (QoS) Solution for Mobile IP", RFC 3383, September 2003.

[6] Tschofenig, H., "Security Threats for NSIS", RFC 4081, Informational , June 2005.

[7] Ash, J., Bader, A., and C. Kappler, "QoS-NSLP QSpec Template", Internet Draft draft-ietf-nsis-qspec-06, Work in Progress , October 2005.

[8] Brunner, M., "Requirement for Signaling Protocols", RFC 3726, April 2004.

[9] Braden, R., "Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", RFC 2205, September 1997.

11.2. Informative References

[10] Koodli, R., "Fast Handovers for Mobile IPv6", RFC 4068, July 2005.

[11] Liebsch, M., "Candidate Access Router Discovery (CARD)", RFC 4066, July 2005.

[12] International Telecommunication Union (ITU) Telecommunication Standardization Sector (ITU-T), "Packet-based multimedia communications systems", ITU-T H. 323, July 2003.

[13] Shen, C., "NSIS Operation Over IP Tunnels", Internet Draft draft-shen-nsis-tunnel-00, Work in Progress , July 2005.

[14] Blake, S., "An Architecture for Differentiated Services", RFC 2475, Informational , December 1998.

[15] Dressler, F., "NSLP for Metering Configuration Signaling", Internet Draft draft-dressler-nsis-metering-nslp-02, Work in Progress , July 2005.

Authors' Addresses

Hong Cheng
Panasonic Singapore Laboratories
Block 1022, Tai Seng Industrial Estate
#06-3530, Tai Seng Avenue
Singapore 534415
Singapore

Phone: +65 6550 5477
Email: hong.cheng@sg.panasonic.com

Jack Huang
Panasonic Singapore Laboratories
Block 1022, Tai Seng Industrial Estate
#06-3530, Tai Seng Avenue
Singapore 534415
Singapore

Phone: +65 6550 5414
Email: jack.huangqj@sg.panasonic.com

Takako Sanda
Matsushita Electric Industrial Co., Ltd. (Panasonic)
5-3, Hikarino-oka, Yokosuka City
Kanagawa 239-0847
Japan

Phone: +81 46 840 5764
Email: sanda.takako@jp.panasonic.com

Toyoki Ue
Matsushita Electric Industrial Co., Ltd. (Panasonic)
5-3, Hikarino-oka, Yokosuka City
Kanagawa 239-0847
Japan

Phone: +81 46 840 5816
Email: ue.toyoki@jp.panasonic.com

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

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