ION Working Group
INTERNET DRAFT
<draft-ietf-ion-nhrp-flowext-00.txt>

Lou Berger
FORE Systems
Rob Enns
Berkeley Networks
Expires: December 1998

NHRP Flow Extension

June 22, 1998

Status of this Memo

Abstract

This document presents an extension to NHRP [RFC2332] that enables resolution of NBMA next hop addresses based on destination flow information. This extension also enables NHSs to relay simple forwarding policy to source stations (NHCs).

1. Introduction

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [RFC2119].

The NBMA Next Hop Resolution Protocol provides a mechanism for a station to resolve a destination internetworking layer address into the corresponding NBMA subnetwork addresses of the "NBMA next hop." As defined in [RFC2332], NBMA next hop subnetwork addresses are

resolved based on an internetworking layer address and, possibly, an address prefix.

Resolution based on destination addresses may not be adequate for all cases. This document is motivated by one such case. The case addressed by this document is when the NBMA next hop is dependent on the contents of a data packet, i.e. its type of traffic. The extension presented in this document provides the ability to resolve NBMA next hops based on destination and additional data packet information. We refer to the presented extension as the "Flow Extension."

There are multiple possible reasons why an NBMA next hop may depend on the contents of a data packet. One example is configured administrative policy at serving or transit NHSs. The choice of selecting next hop based on just destination address, or selecting next hop based on additional data packet information is considered to be a local policy decision. The presented extension will interoperate with NHSs and NHCs that don't support the Flow Extension as well as those that implement disparate selection policies.

2. Discussion

The Flow Extension enables the communication of a limited amount of policy information to the source NHC from NHSs along the routed path. NHRP currently supports a very limited amount of policy communication via a resolution NAK. With the resolution NAK, a source NHS can be informed that a resolution request is administratively prohibited. The Flow Extension allows the communication of additional policy information. This additional policy information indicates which types of data traffic, or flows, can be directed to which NBMA next hops.

The Flow Extension supports the communication of one final piece of policy information. With the Flow Extension, the source NHC can be informed that matching data packets should be silently discarded. When a source NHC drops such traffic, network resource usage is reduced since such traffic would be discarded further along the routed path. Source NHCs are not required to discard any data traffic and may forward any data packets along the routed path. Currently, source NHCs that cannot resolve an NBMA next hop are expected to forward data packets along the routed path.

The rest of this section covers general issues related to Flow Extension creation and processing. Specific format requirements and Resolution Request and Reply message processing impact are covered in sections 3 and 4 respectively.

2.1 Policy Communication

The objective of the Flow Extension is twofold. The first is to enable the resolution of a single destination address to more than one NBMA next hop address based on flow information. The second is to allow selected traffic to a particular destination address to be resolved to one or more NBMA next hop addresses while other traffic to the same destination address is not resolved. When traffic is unresolved, the extension allows the source NHC to be informed that such traffic should be dropped. NHCs are permitted to continue handling unresolved traffic in the same fashion as described [RFC2332].

To achieve the two objectives of the Flow Extension, the extension is included in NHRP Resolution Request and Reply messages. In Request messages, the Flow Extension contains a description of the specific flow associated with the request, the flow matching capabilities of the requester, and the policy information from the NHSs that have already processed the Request message. Reply messages are handled in the traditional manner per [RFC2332] and also include the policy information from the serving NHS.

When a source NHC creates a Resolution Request with a Flow Extension, it describes its capabilities and the flow associated with the request. The requesting NHC also initializes the policy information fields that will be updated by downstream NHSs to indicate an unrestrictive policy. NHSs add information based on their local administrative policy. This information can indicate that the flow is administratively prohibited and should be dropped at the source NHC, that the specific flow is acceptable or, lastly, that the flow falls within a set of acceptable flows. NHSs processing Request messages may indicate that their policy is more restrictive than the policy described in the received message and are not allowed to indicate a less restrictive policy. Serving NHSs add their policy information when generating a Resolution Reply message.

Once the Reply message reaches the requester, the requester is expected to follow the policy information contained in the Reply. This includes only sending data packets matching the described flow to the listed NBMA next hop addresses, and following the drop indication in NAK messages. The requester is permitted to handle NAK Reply messages and associated data packets in the traditional, pre Flow Extension, fashion. The requester is also permitted to forward data packets matching the described flow via the routed path rather than following the information in a non-NAK Reply message.

2.2 Compatibility

There are no compatibility issues with implementations that do not support the Flow Extension. [RFC2332] requires that the extension be transported even when a processing station does not support the extension. Flow Extension aware stations may generate and handle multiple resolution requests for the same destination each with a different flow descriptions. NHSs that do not support the extension will see such flow specific requests as revalidation requests all for the same destination.

From the policy standpoint, NHSs that do not support the extension are able to verify that a request is administratively acceptable, but do so just on a destination rather than flow basis. When a particular Resolution Request message containing a Flow Extension is processed by a combination of extension supporting and non-supporting NHSs, the corresponding Reply message will reflect flow based policy from extension supporting NHSs and the coarser grained destination based policy from non-supporting NHSs.

2.3 NHS Handling of Non-Served NHC Resolution Requests

The NHRP specification [RFC2332] requires that NHCs only send Resolution Request messages to their serving NHS, but it places no corresponding requirement on NHSs. In order to ensure that a Resolution Request message follows the routed path, NHSs that support the Flow Extension MUST respond to Request messages which contain the Flow Extension from non-served NHCs with an Error Indication.

2.4 Security Considerations

The Flow Extension is used to communicate a limited amount of policy information. Consequently there are a number of security related issues. The issues have to do with both the communication of the policy information and with the use of the policy information.

2.4.1 Trust of NHRP Stations

The Flow Extension is used to relay policy information between NHCs and NHSs. For the extension to be useful NHSs and NHCs must rely upon each other to relay policy information correctly, to properly update the information and to act appropriately based on the information. Potential threats include downstream and reverse path NHSs overriding communicated policy information, and NHCs that send NBMA Next Hops traffic not included in the communicated policy information.

In environments where reliance (and trust) between NHSs and NHCs

posses an unacceptable risk, the NHRP Authentication Extension SHOULD be used to identify trusted NHRP speakers and the Flow Extension SHOULD only be supported (processed from and sent to) between authenticated speakers. In environments where this approach is unacceptable or unavailable support for the Flow Extension SHOULD be disabled. With Extension processing disabled, the threat becomes the same as that posed by standard NHRP.

2.4.2 Dissemination of Policy Information

With the Flow Extension, an NHS' forwarding policy is communicated to other NHSs and the requesting NHC. In some cases the communication of this policy information will be acceptable, in others the dissemination of such information may be undesirable. In the later case, the communicated policy can be limited to a minimum. If it is unacceptable to communicate any policy information, again, support for the Flow Extension SHOULD be disabled.

It should be noted that the Resolution Reply messages containing Flow Extensions may be carried along a different routed path from the serving NHS to the requesting NHC. Forcing symmetric routing was considered in order to limit the dissemination of policy information, but was found to result in incompatibility with existing NHRP implementations.

2.4.3 Policy Information Updates

An NHS' local policy information may be updated at anytime. Such an update may relate to previously requested resolution information. The new policy information will eventually be communicated once the source NHC refreshes the resolution information. Between the time the information is changed and the resolution is refreshed, the source NHS may handle data packets based on the out of date policy information.

The handling of data packets based on old policy information may be an issue for some environments. In cases where this is an issue, NHSs can force the removal of the related resolution information. To do this, an NHS SHOULD cache information related to Resolution Requests that contain a Flow Extension, and then issue Purge Request messages for destinations that are affected by local policy updates.

3. NHRP Flow Extension Format

Compulsory = 0
Type = To Be Assigned
Length = variable, see format definitions

The NHRP Flow Extension is carried in NHRP Resolution Request and Reply messages to convey flow related information. A source NHC includes the extension to indicate that the extension is supported. Within the extension, the source NHC identifies the flow associated with the resolution request, indicates its flow matching capabilities, and initializes the policy information related fields. Transit NHSs update the policy related fields based on their local policies. The NHS responding to the Request message, the "responder", copies a received Flow Extension to the corresponding Resolution Reply message and updates the policy related fields based on its local policies.

The format of the extension varies based on the traffic being described. In all formats, the extension has request and policy information related fields. The request related fields describe the flow from the source NHC's perspective. The policy information related fields are used to relay the policy found along the routed path to the destination, and may be set by transit and serving NHSs. The first byte of the extension indicates the type of flow information associated with the request. The source NHS selects the type based on its capabilities. The traffic type values defined in this document are:

Value	Label	Description
0x00	Reserved	Illegal Value
0x01	IPv4	Generic IPv4 Header
0x02	IPv6	Generic IPv6 Header
0x03	IPv4-TCP/UDP	IPv4 Header with TCP/UDP like ports
0x04	IPv6-TCP/UDP	IPv6 Header with TCP/UDP like ports
0x05	IPv4-IPSEC	IPv4 Header with IPSEC SPI
0x06	IPv6-IPSEC	IPv6 Header with IPSEC SPI
0x07-0x7F	Reserved	To be assigned by IANA
0x80-0xFF	Reserved	Allocated to the ATM forum

A Flow Extension containing an unrecognized traffic type values SHALL be treated as an unrecognized extension.

3.1 NHRP Flow Extension Format: IPv4

This format is used to describe any type of IPv4 data packet. It is expected to be used when a more specific description is not defined or supported.

Extension Length = 16

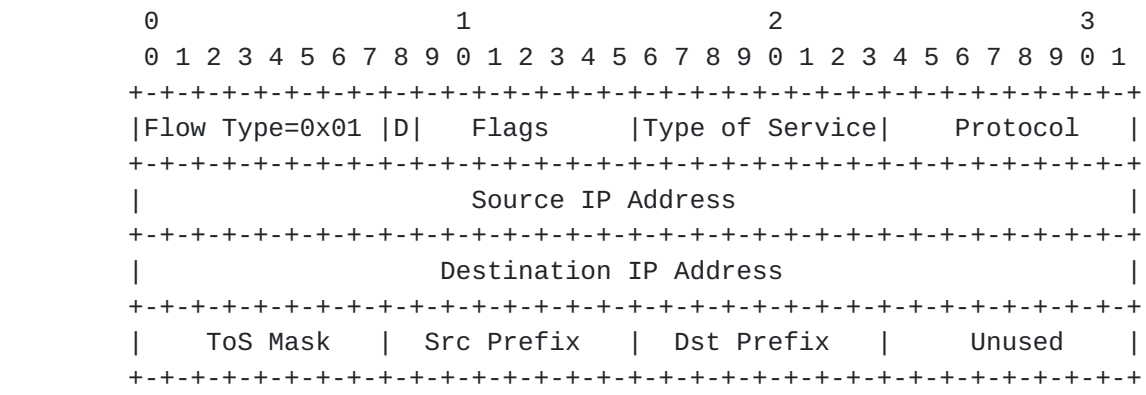

Flow Type

Indicates type of flow information carried in the extension.

D bit

When set, the "D" bit indicates that data packets matching the flow information SHOULD be silently discarded by the source NHC. If the source NHC forwards matching data packets, such packets MUST be forwarded via the routed path.

This bit is set to zero (0) in Resolution Request messages and may be set by NHSs. If a Flow Extension is received with a set value (1), an NHS MUST forward the extension with the "D" bit set.

Flags

The following flags are defined IPv4 flows:

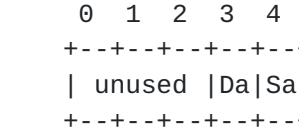

Da

When set, this bit indicates that flows may be identified using prefix matching on the Destination IP Address field of the IP Header. Prefix matching uses a specified number of bits from the address field rather than the whole field. This bit is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Sa

When set, this bit indicates that flows may be identified using prefix matching on the Source IP Address field of the IP Header. This bit is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

P
When set, this bit indicates that flows MUST be identified using the value in the Protocol field. This bit is only meaningfull if the Protocol field is non-zero. This bit MUST be cleared by the source NHC and MAY be set transit or serving NHSs. When a Flow Extension is received with this bit set, an NHS MUST NOT clear this bit.

T
When set, this bit indicates that flows may be identified based on the Type of Service field of the IP Header. This bit is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Unused bits and fields MUST be set to zero (0) by the source NHC and not modified by transit or serving NHSs.

Type of Service
Value of the IP Type of Service field for the associated flow. This value is set by the source NHC and MAY NOT be modified by transit or serving NHSs. This field is only meaningful when the "T" bit is set.

Protocol
Value of the next level protocol field for the associated flow. A value of zero (0) indicates that the source NHC will not include this field in identification of the associated flow. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Source IP Address
The source IP address for the associated flow. A value of zero (0) indicates that the source NHC will not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Destination IP Address
The destination IP address of the associated flow. A value of zero (0) indicates that the source NHC will not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

ToS Mask
This field indicates the bits that must be set in a data packet's Type of Service field in order to be associated with the flow. A value of zero (0) indicates that the field SHALL NOT be included in flow identification. This field is only meaningful when the "T" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs add their policy information by setting cleared bits in the field. NHSs MUST NOT clear bits in the field.

Src Prefix
This field carries the number of bits of the Source IP Address field to be used in identifying data packets associated with the flow. A value of zero (0) indicates that the field SHALL NOT be included in flow identification. A value of 32 indicates a host address, i.e. all bits should be used. Values greater than 32 are illegal. This field is only meaningful when the Source IP Address field is non-zero and when the "Sa" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST NOT decrease the value.

Dst Prefix
This field carries the number of bits of the Destination IP Address field to be used in identifying data packets associated with the flow. A value of 0 or 32 indicates a host address, i.e. all bits should be used. Values greater than 32 are illegal. This field is only meaningful when the Destination IP Address field is non-zero and when the the "Da" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST set the value to 32 to indicate a host address policy. NHSs MUST NOT decrease the value.

3.2 NHRP Flow Extension Format: IPv6

This format is used to describe any type of IPv6 data packet. It is expected to be used when a more specific description is not defined or supported.

Extension Length = 48

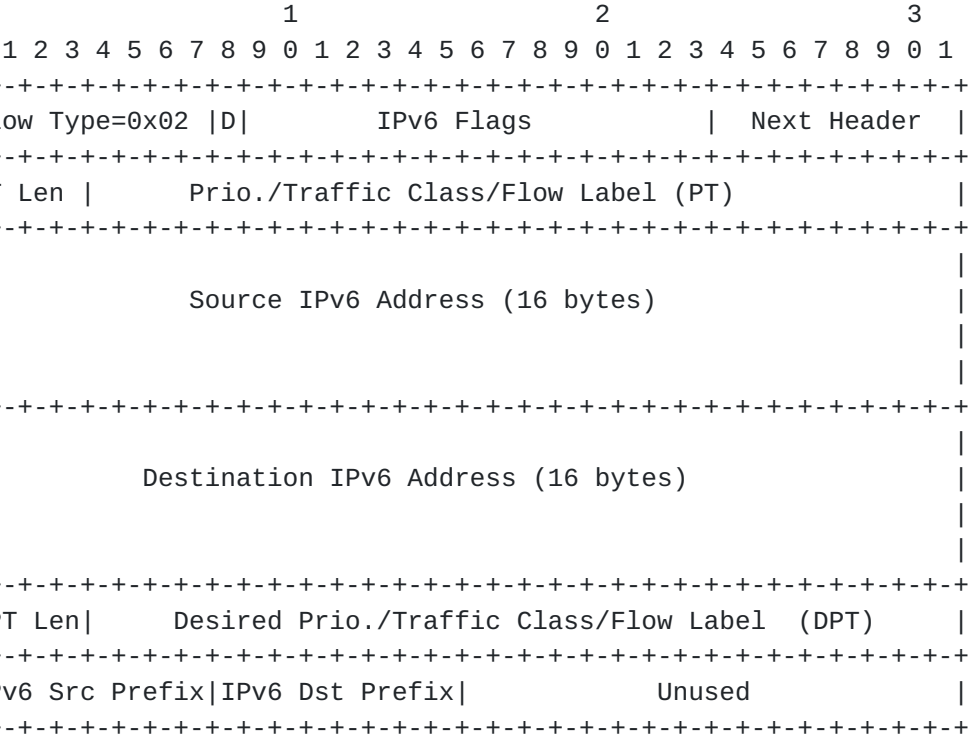

IPv6 Flags

The following flags are defined IPv6 flows:

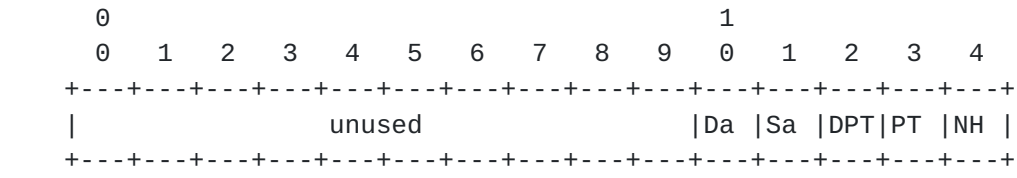

DPT

When set, this bit indicates that flows MUST be identified based on the Prio./Traffic Class and Flow Label fields of the IP Header. This bit MUST be zero when the PT bit is zero. This bit MUST be cleared by the source NHC and MAY be set transit or serving NHSs. When a Flow Extension is received with this bit set, an NHS MUST NOT clear this bit.

PT

When set, this bit indicates that flows may be identified based on the Prio./Traffic Class and Flow Label fields of the IP Header. This bit is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

NH
When set, this bit indicates that flows MUST be identified using the value in the Next Header field. This bit is only meaning full if the Next Header field is non-zero. This bit MUST be cleared by the source NHC and MAY be set transit or serving NHSs. When a Flow Extension is received with this bit set, an NHS MUST NOT clear this bit.

Next Header
Value of the IPv6 Next Header field for the associated flow. A value of zero (0) indicates that the source NHC will not include this field in identification of the associated flow. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

PT Len
Indicates the number of bits in the Prio./Traffic Class (first) portion of the PT field. This field is set by the source NHC and is not modified by transit or serving NHSs. The source NHS MAY set this value to zero (0). This field is only meaningful when the "PT" bit is set.

Prio./Traffic Class and Flow Label (PT)
The 28-bit PT field for the associated flow. The PT field is composed of two parts. The first part carries either the IPv6 priority value or the the proposed IPv6 Traffic Class field. The second part is the IPv6 Flow Label. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs. This field is only meaningful when the "PT" bit is set.

Source IPv6 Address
The source IPv6 address for the associated flow. A value of zero (0) indicates that the source NHC will not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Destination IPv6 Address
The destination IPv6 address of the associated flow. A value of zero (0) indicates that the source NHC will not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

DPT Len
This field indicates the number of bits in the Prio./Traffic Class (first) portion of the DPT field. This field is only meaningful when the "DPT" bit is set. Legal values are between 0 and 28. This field may contain a different value than the PT Len field.

This field MUST be initialized by the source NHC to a value of zero (0). Transit and serving NHSs MAY set this field only when the received Flow Extension has the DPT bit cleared and the PT bit set. If an NHS sets this field, it MUST also set the DPT bit.

Desired Prio./Traffic Class and Flow Label (DPT)

This field carries the Prio./Traffic Class and Flow Label fields associated with the flow. The Prio./Traffic Class portion of the DPT field indicates the bits that must be set in a data packet's Prio./Traffic Class field in order to be associated with the flow. The Flow Label portion of the DPT field contains the value that must be in a data packet's Flow Label field in order to be associated with the flow.

The division of of the Prio./Traffic Class field is indicated by the DPT Len field. The Prio./Traffic Class field is contained in the first (leftmost) DPT Len bits of the DPT field. The Flow Label portion is in the last (rightmost) 28 minus DPT Len bits of the DPT field. This field is only meaningful when the "DPT" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs add their Prio./Traffic Class related policy information by setting cleared bits in the Prio./Traffic Class portion of the DPT field. NHSs MUST NOT clear bits in the Prio./Traffic Class field. NHSs MAY set the Flow Label portion of the field when the received Flow Extension has the DPT bit cleared and the PT bit set. If an NHS sets the Flow Label portion of the field, it MUST also set the DPT bit.

IPv6 Src Prefix

This field carries the number of bits of the Source IPv6 Address field to be used in identifying data packets associated with the flow. A value of zero (0) indicates that the field SHALL NOT be included in flow identification. A value of 128 indicates a host address, i.e. all bits should be used. Values greater than 128 are illegal. This field is only meaningful when the Source IPv6 Address field is non-zero and when the "Sa" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST NOT decrease the value.

IPv6 Dst Prefix

This field carries the number of bits of the Destination IPv6 Address field to be used in identifying data packets associated with the flow. A value of 0 or 128 indicates a host address, i.e. all bits should be used. Values greater than 128 are illegal. This field is only meaningful when the Destination IPv6 Address

field is non-zero and when the the "Da" bit is set.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST set the value to 128 to indicate a host address policy. NHSs MUST NOT decrease the value.

3.3 NHRP Flow Extension Format: IPv4-TCP/UDP

This format is used to describe IPv4 data packets carrying a protocol that supports TCP/UDP-like ports. Specifically protocols that carry a 16-bit source port field at the start of the transport header and 16-bit destination port field starting at bit 16 of the transport header.

Fields and flags not listed have the same meaning as defined in previous sections.

Extension Length = 28

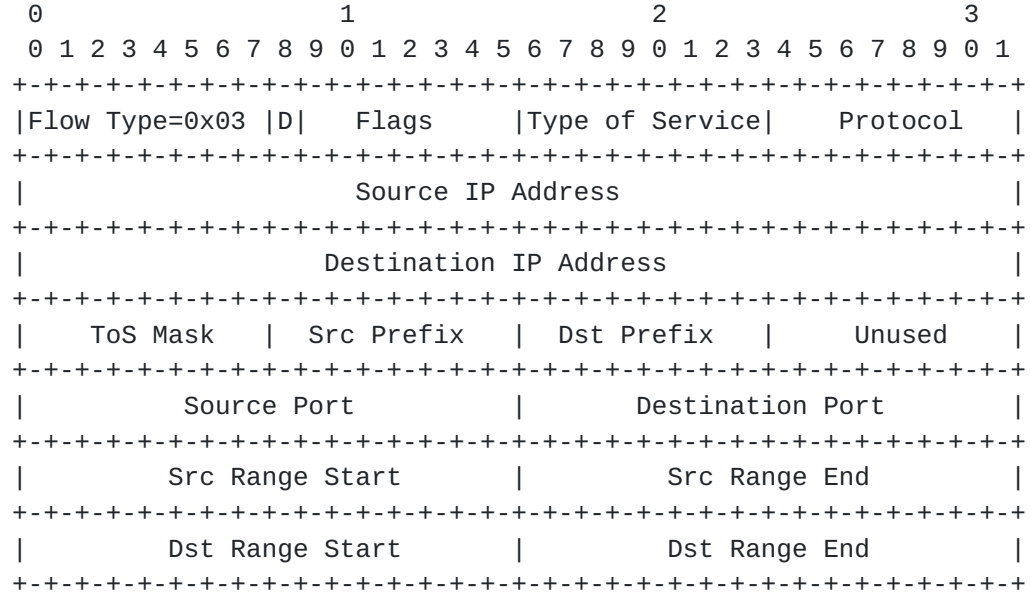

Flags

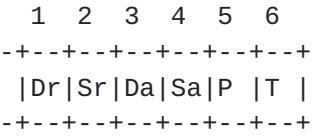

Dr

When set, this bit indicates that flows may be identified using a range of acceptable values in the Destination Port field of the transport header. This bit is set by the source NHC and MAY NOT

be modified by transit or serving NHSs.

Sr
When set, this bit indicates that flows may be identified using a range of acceptable values in the Source Port field of the transport header. This bit is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Source Port
Value of the Source Port field of the transport header for the associated flow. A value of zero (0) indicates that the source NHC has not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Destination Port
Value of the Destination Port field of the transport header for the associated flow. A value of zero (0) indicates that the source NHC has not include this field in flow identification. This field is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

Src Range Start
This field carries the lower bound on Source Port field values that will be considered to be associated with the flow. A zero (0) value indicates that the Source Port field of the transport header is not included in flow identification. This field is only meaningful when the Source Port field is non-zero. When the "Sr" bit is cleared (0), an NHS setting this field MUST set the Src Range Start and Src Range End fields to the same value. This field MUST NOT be set to a value greater than the Source Port field.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST NOT decrease the value.

Src Range End
This field carries the upper bound on Source Port field values that will be considered to be associated with the flow. This field is only meaningful when the Source Port and Src Range Start fields are non-zero. When the "Sr" bit is cleared (0), an NHS setting this field MUST set the Src Range Start and Src Range End fields to the same value. This field MUST NOT be set to a value less than the Source Port field.

This field MUST be initialized by the source NHC to a value of all ones (0xffff). NHSs MAY decrease the value based on their policy information. NHSs MUST NOT increase the value.

Dst Range Start
This field carries the lower bound on Destination Port field values that will be considered to be associated with the flow. A zero (0) value indicates that the Destination Port field of the transport header is not included in flow identification. This field is only meaningful when the Destination Port field is non-zero. When the "Dr" bit is cleared (0), an NHS setting this field MUST set the Dst Range Start and Dst Range End fields to the same value. This field MUST NOT be set to a value greater than the Destination Port field.

This field MUST be initialized by the source NHC to a value of zero (0). NHSs MAY increase the value based on their policy information. NHSs MUST NOT decrease the value.

Dst Range End
This field carries the upper bound on Destination Port field values that will be considered to be associated with the flow. This field is only meaningful when the Destination Port and Dst Range Start fields are non-zero. When the "Dr" bit is cleared (0), an NHS setting this field MUST set the Dst Range Start and Dst Range End fields to the same value. This field MUST NOT be set to a value less than the Destination Port field.

This field MUST be initialized by the source NHC to a value of all ones (0xffff). NHSs MAY decrease the value based on their policy information. NHSs MUST NOT increase the value.

3.4 NHRP Flow Extension Format: IPv6-TCP/UDP

This format is used to describe IPv6 data packets carrying a protocol that supports TCP/UDP-like ports.

Fields and flags not listed have the same meaning as defined in previous sections.

Extension Length = 60

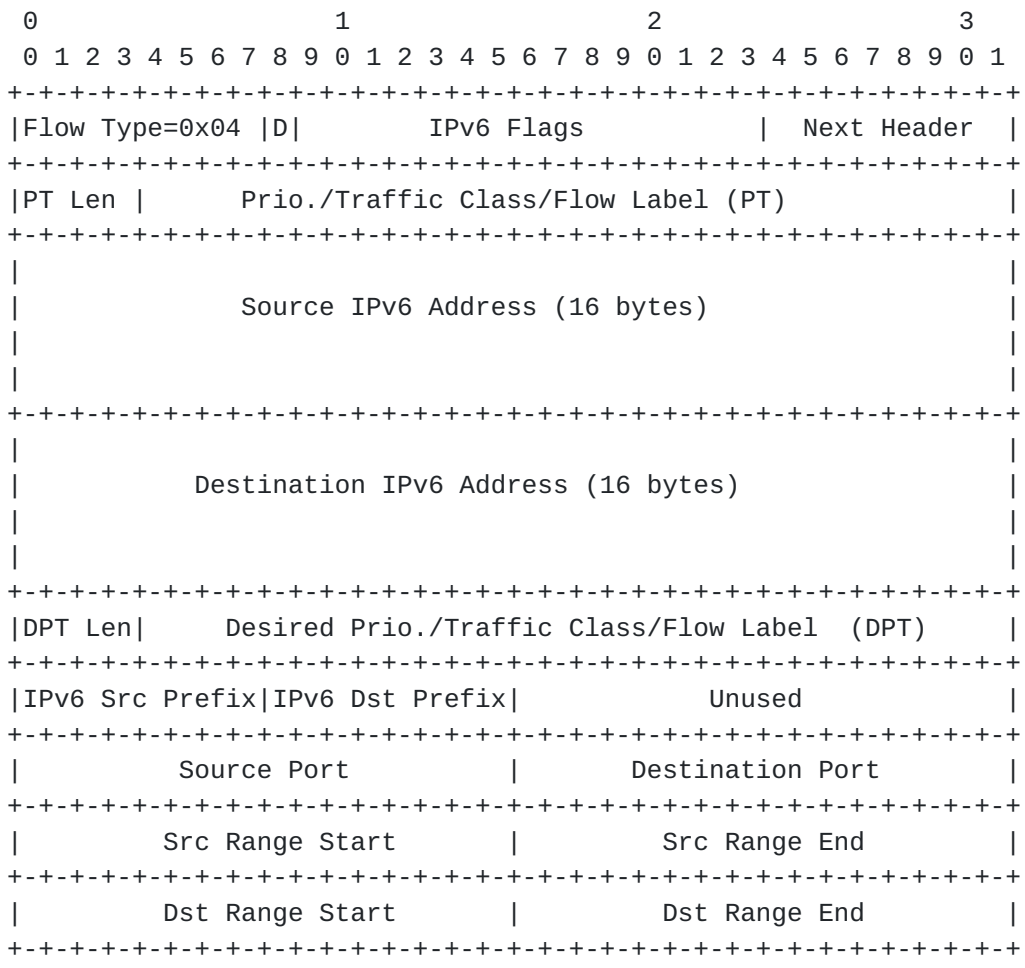

IPv6 Flags

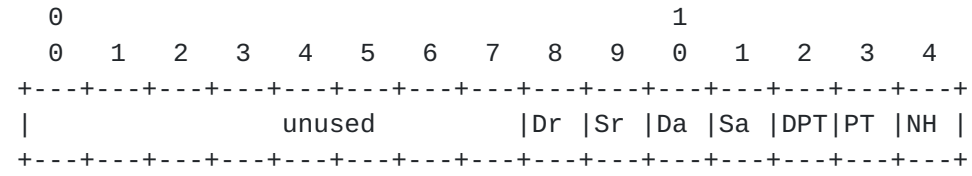

3.5 NHRP Flow Extension Format: IPv4-IPSEC

This format is used to describe IPv4 data packets carrying an IPSEC transport protocol. There are currently two such protocols defined: Authentication Header (AH), for authentication[RFC1826]; and Encapsulating Security Payload (ESP), for integrity and confidentiality[RFC1827]. Flows are identified for both protocols using the protocol's IPSEC Security Parameter Index, or SPI, field.

Fields and flags not listed have the same meaning as defined in previous sections.

Extension Length = 20

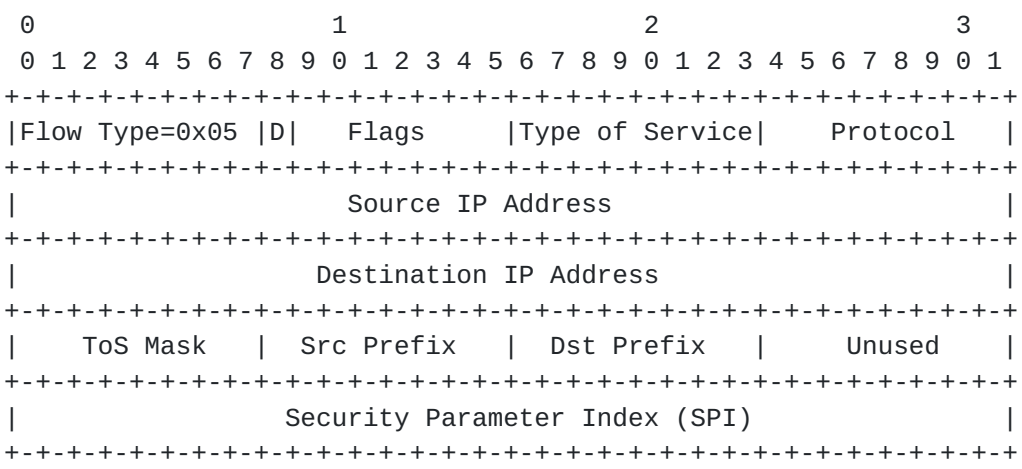

Flags

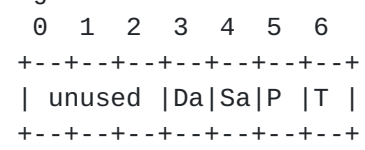

Security Parameter Index

The Security Parameter Index (SPI) field value associated with the flow. The SPI field is 32-bits long and located at the start of the transport header for ESP, and at bit 32 of the transport header for AH. This value is set by the source NHC and MAY NOT be modified by transit or serving NHSs.

3.6 NHRP Flow Extension Format: IPv6-IPSEC

This format is used to describe IPv6 data packets carrying an IPSEC transport protocol.

Fields and flags not listed have the same meaning as defined in previous sections.

Extension Length = 56

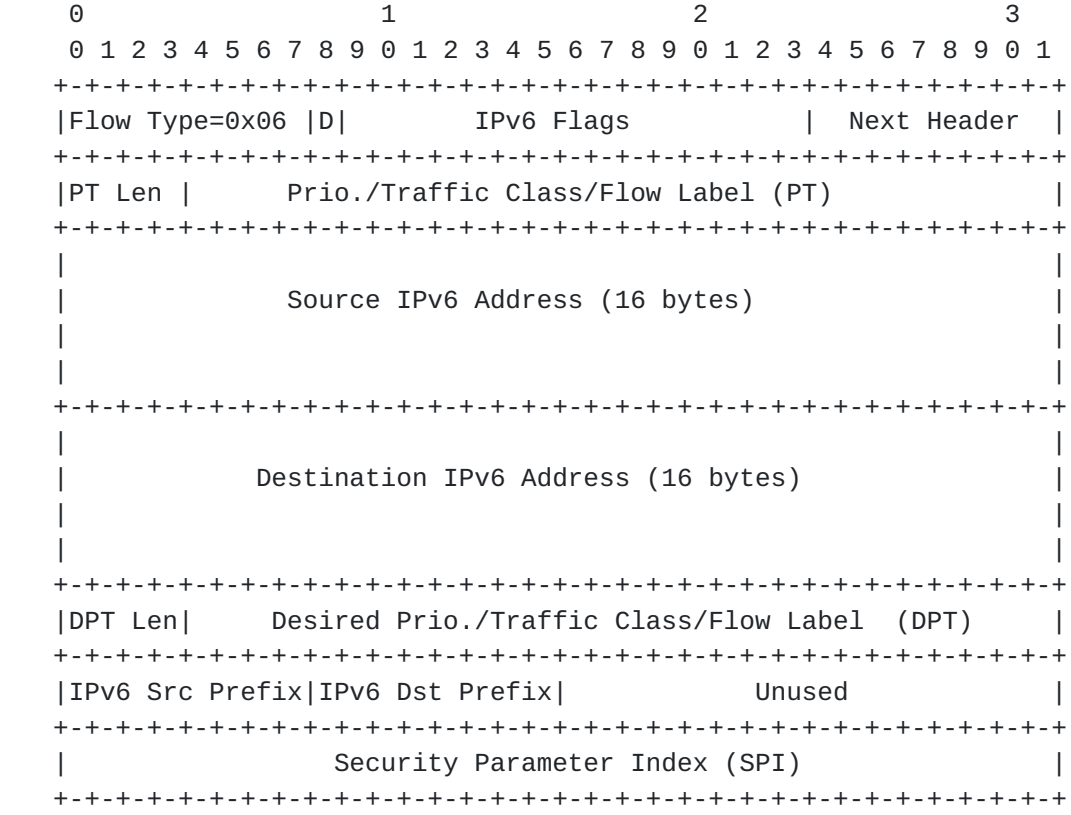

IPv6 Flags

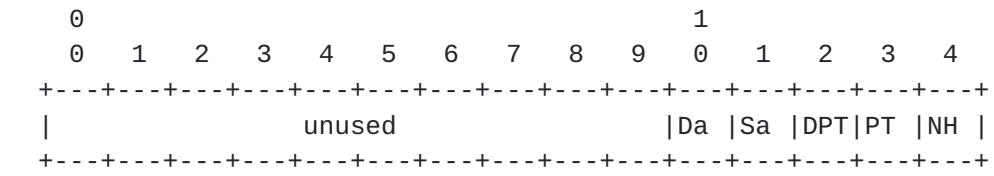

4. Message Processing

If a requester wishes to obtain flow information, it SHALL include the Flow Extension in the NHRP Resolution Request. The requester SHOULD include the most specific flow type that it supports. The requester MAY specify a particular flow or subset of flows by setting some or all of the fields in the flow detail to non-zero values.

When processing NHRP Resolution Request messages, NHSs that support the Flow Extension MUST verify that the Request is following the routed path. The NHS checks if the sender of the message is an NHC. When the sender is an NHC, the NHS verifies that the NHC is one of its served NHCs. If the sending NHC is not a served NHC, then the NHS MUST generate an NHRP Error Indication with an Error Code of Protocol Error.

If the protocol check passes, NHSs processing Request messages containing a Flow Extension MUST examine the described flow to verify that it is acceptable. Both the request and policy information related fields MUST be examined. The request portion of the extension is unacceptable if data packets matching the flow would not be forwarded. The policy information is unacceptable if the Flow Extension cannot be updated to reflect an NHS' policy information. If both portions are acceptable, the NHS MUST update the policy information related fields as needed to reflect its policy and then handle the Request message per [RFC2332]. Note that an NHS MAY NOT change any policy information field to be less restrictive.

If either the request or policy information portions of the Flow Extension is unacceptable, the NHS MUST issue a NAK Resolution Reply of type Administratively Prohibited. The NHS MUST update the Flow Extension to reflect the matching policy and SHOULD set the "D" bit. When updating the policy information portion of the Flow Extension, an NHS SHOULD describe the matching (failed or acceptable) policy in the broadest possible terms rather than simply replaying the requested flow. An NHS responder MAY choose to reply with limited information for security reasons.

On receipt of a Resolution Reply message that contains a Flow Extension, a requester that supports the Flow Extension SHOULD follow the policy information relayed in the extension. If the requester does not follow the policy, the requester MUST NOT forward data packets based on the information contained in the reply message. A requester may chose not to follow a relayed policy because the policy is unacceptable or due to resource limitations.

In the expected normal case, the requester will follow the policy information relayed in the extension. When the message is a NAK Resolution Reply of type Administratively Prohibited the requester SHOULD check the "D" bit in the Flow Extension. If the bit is set, data packets matching the flow information SHOULD be silently discarded. If the requester forwards matching data packets, such packets MUST be forwarded via the routed path.

When the received Resolution Reply message does not contain a NAK, the requester SHOULD forward data packets that match the flow specified in the Flow Extension using the information contained in the Resolution Reply. The requester MUST NOT forward data packets that do not match specified flow using the information contained in the Resolution Reply.

5. IANA Considerations

IANA is requested to manage and assign the range of Flow Type field values from 0x07 to 0x7F. New assignments should be made with the guidance of the relevant IESG Area Director or their designee. Additionally, all requests for assignments should be honored when the usage of the requested value is documented in a publicly available and unrestricted (including time) form. Preferably the document will be available via a standards organization's web site.

6. References

[RFC1826] Atkinson, R., "IP Authentication Header", RFC 1826, NRL, August 1995.

[RFC1827] Atkinson, R., "IP Encapsulating Security Payload", RFC 1827, NRL, August 1995.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," RFC 2119.

[RFC2332] Luciani, J., Katz, D., Piscitello, D., Cole, B., Doraswamy, N., "NBMA Next Hop Resolution Protocol (NHRP)," RFC 2332.

7. Acknowledgments

The authors thank Rajesh Varadarajan, Ravi Shekhar and Jim Luciani for their valuable comments and corrections.

8. Authors' Addresses

Lou Berger
FORE Systems
1595 Spring Hill Road
Vienna, VA 22182 USA
Phone: +1 703-245-4544
Email: lberger@fore.com

Rob Enns
Berkeley Networks
1805 McCandless Drive
Milpitas, CA 95035
Phone: 408-719-3059
Email: rpe@berkeleynet.com