

Network Working Group
Request for Comments: 2735
Category: Standards Track

B. Fox
Equipe Communications
B. Petri
Siemens AG
December 1999

NHRP Support for Virtual Private Networks

Status of this Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (1999). All Rights Reserved.

Abstract

The NBMA Next Hop Resolution Protocol (NHRP) is used to determine the NBMA subnetwork addresses of the "NBMA next hop" towards a public internetworking layer address (see [1]). This document describes the enhancements necessary to enable NHRP to perform the same function for private internetworking layer addresses available within the framework of a Virtual Private Network (VPN) service on a shared NBMA network.

1. Introduction

NHRP is a public internetworking layer based resolution protocol. There is an implicit understanding in [1] that a control message applies to the public address space.

Service Providers of Virtual Private Network (VPN) services will offer VPN participants specific service level agreements (SLA) which may include, for example, dedicated routing functions and/or specific QoS levels. A particularly important feature of a VPN service is the ability to use a private address space which may overlap with the address space of another VPN or the Public Internet. Therefore, such an internetworking layer address only has meaning within the VPN in which it exists. For this reason, it is necessary to identify the VPN in which a particular internetworking layer address has meaning, the "scope" of the internetworking layer address.

As VPNs are deployed on shared networks, NHRP may be used to resolve a private VPN address to a shared NBMA network address. In order to properly resolve a private VPN address, it is necessary for the NHRP device to be able to identify the VPN in which the address has meaning and determine resolution information based on that "scope".

As VPN services are added to an NBMA network using NHRP devices, it may be necessary to support the service with legacy NHRP devices that do not have VPN knowledge and so do not explicitly support VPNs. This document describes requirements for "VPN-aware" NHRP entities to support VPN services while communicating with both "VPN-aware" and "non-VPN-aware" NHRP entities.

2. Overview of NHRP VPN Support

2.1 Terminology

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 RFC 2119 [4].

In addition to the terminology specified in section 2.1 of [1], the following definitions and acronyms are used:

Default Routing Instance -- In the presence of VPNs, all packets are processed (e.g., routed) within the context of a specific VPN. In the case where no VPN is indicated, a packet is processed according to a default VPN, i.e., a Default Routing Instance. This routing instance may be the Public Internet, a particular VPN, etc. The term only has meaning for "VPN-aware" NHRP entities.

Virtual Private Network (VPN) -- in the context of this specification, this term is used as described in [3].

VPN-aware -- a "VPN-aware" NHRP entity is an NHRP entity that implements the NHRP enhancements for VPNs as defined in this document.

Non-VPN-aware -- a "non-VPN-aware" NHRP entity is an NHRP entity which is deployed as part of a single VPN, but is not VPN-aware. Restrictions applying to non-VPN-aware NHRP entities are outlined below. NHRP devices as specified in [1] are examples of non-VPN-aware entities.

VPN encapsulation -- An LLC/SNAP encapsulation of a PDU with an indication of the VPN to which the PDU belongs. In the case that the underlying NBMA network is an ATM network, VPN encapsulation is specified in section 8 of [2].

VPN identifier (VPN-ID) -- in the context of this specification, this term is used as specified in [3].

VPN signalling -- in the context of this specification, this term is used to denote a method to indicate the VPN-ID via control signalling or similar ways in the control path.

2.2 VPN Support Overview

When supporting NHRP for a VPN, it is necessary to specify to which VPN the NHRP message applies in order to comply with the VPN service level agreement applicable to that VPN.

On some NBMA networks, it is possible to establish a VPN-specific control path between NHRP devices. This is sufficient to identify the NHRP control packets as belonging to the "inherited" VPN. However, when that alternative is not used, the NHRP device must specify the VPN to which an NHRP packet applies in the PDU.

It is not useful to add a VPN extension to NHRP control messages because transit NHRP Servers are not required to process the extensions to an NHRP control message (see 5.3 in [1]). NHRP Servers already deployed might resolve the control packet within the scope of the public internetworking layer address space instead of the private address space causing problems in routing.

Instead, an LLC/SNAP header with a VPN indication (as specified in Section 4.1 below) will be prepended to the NHRP control message. This solution allows the same VPN-specific LLC/SNAP header to be prepended to PDUs in both the control and data paths.

3. NHRP VPN Operation

3.1 VPN-Aware NHRP Operation

When a VPN-aware NHRP device forwards a packet pertaining to a particular VPN, that device MUST be able to indicate the VPN either:

- a) explicitly through use of the VPN-specific LLC/SNAP header or
- b) implicitly through an indication via VPN signalling.

This applies to NHC-NHS, NHS-NHS, and NHS-NHC control messages as well as NHC-NHC shortcut traffic.

For case a), the indication of the VPN-ID is via a VPN-specific LLC/SNAP header specified in section 4.2 below. In the case of an underlying ATM network, see also section 8 of [2].

For case b), the method used to indicate the VPN-ID via VPN signalling depends on the mechanisms available in the underlying network and is outside the scope of this memo. A VPN-aware NHRP entity using VPN signalling SHOULD NOT also indicate the VPN-ID explicitly for any PDU on the related path.

In transiting an NHRP Server, the VPN identification MAY be forwarded in a different format than was received, however, the same VPN-ID MUST be indicated for the message. For example, a PDU received with an LLC/SNAP header containing a VPN identifier may be forwarded on a control path which was established with an indication of the same VPN without the VPN-specific LLC/SNAP header.

When a VPN capable NHRP entity receives an NHRP message from a VPN-aware NHRP device without a VPN indication via VPN encapsulation or VPN signalling, the message applies to the default routing instance supported by the shared infrastructure. The public Internet or a particular VPN routing realm may be configured as the default routing instance.

3.2 Interactions of VPN-aware and non-VPN-aware NHRP entities

A VPN-aware NHRP entity MUST be able to indicate the VPN-ID in one of the ways specified in section 3.1 above. It MAY participate in more than one VPN.

Because a non-VPN-aware NHRP device does not understand the concept of VPNs, it only supports a single routing instance. Therefore, a non-VPN-aware NHRP entity belongs to exactly one VPN without being aware of it. All internetworking packets sent by that entity are assumed to belong to that VPN (Note that if the current IPv4-based Internet is regarded as just one big VPN, attached IPv4 hosts may e.g. be regarded as being "contained" in that VPN).

In order for a non-VPN-aware NHRP entity to interact with a VPN-aware NHRP entity, the VPN-aware NHRP entity MUST be configured to associate the correct VPN-ID with information received from the non-VPN-aware entity. In other words, the VPN-aware NHRP entity acts as in the case of option b) from section 3.1 where the VPN-ID was indicated via VPN signalling. However, this association is provisioned using administrative means that are beyond the scope of this document instead of via VPN signalling. Further, it MUST be ensured by administrative means that non-VPN-aware NHRP entities only communicate either with other NHRP entities contained in the same VPN, or with VPN-aware NHRP entities with pre-configured information about the related VPN-ID of those non-VPN-aware entities.

VPN-aware NHRP entities SHALL only send information to non-VPN-aware NHRP entities if that information belongs to the VPN in which the non-VPN-aware entity is contained. Information sent to a non-VPN-aware NHRP entity MUST not include any indication of the VPN-ID.

In order to correctly transfer data packets, it is necessary for VPN-aware ingress NHRP clients to know whether their partner is also VPN-aware. If the egress is VPN-aware, the ingress NHC will also use the means described in section 3.1 on an NBMA shortcut to that egress NHC to specify the VPN to which the data packet belongs.

For this purpose, a further NHRP extension (in addition to those specified in section 5.3 of [1]) is specified which is called NHRP Device Capabilities extension (see section 4.2 below). This extension currently indicates the VPN capabilities of NHRP source and destination entities, but may also be used in the future for further additions to NHRP to indicate other capabilities as well.

3.3 Handling of the NHRP Device Capabilities extension

The NHRP Device Capabilities extension MUST be attached to all NHRP Resolution Requests generated by a VPN-aware source NHRP entity. The device SHOULD set the Source Capabilities field to indicate that it supports VPNs. The compulsory bit MUST be set to zero, so that a non-VPN-aware NHS may safely ignore the extension when forwarding the request. In addition, the A-bit (see section 5.2.1 of [1]) SHOULD be set to indicate that only authoritative next hop information is desired to avoid non-authoritative replies from non-VPN-aware NHRP servers.

Since a non-VPN-aware NHS is not able to process the NHRP Device Capability extension, Network Administrators MUST avoid configurations in which a VPN-aware NHRP Client is authoritatively served by a non-VPN-aware NHRP Server.

If an egress NHS receives an NHRP Resolution Request with an NHRP Device Capability Extension included, it returns an NHRP Resolution Reply with an indication of whether the destination is VPN-aware by correctly setting the target capabilities flag [see Section 4.2].

If an egress NHS receives an NHRP Resolution Request without an NHRP Device Capability Extension included or with the source capabilities flag indicating that the source NHRP device is non-VPN-aware, it MAY act in one of the following ways:

- It MAY reject the NHRP Resolution Request; this is because the VPN-aware destination will be unable to determine the context of information received on an NBMA shortcut from a non-VPN-aware NHRP source. This is the default case.
- If the destination is also non-VPN-aware, it MAY accept the request and return an NHRP Resolution Reply. By default, the two non-VPN-aware NHRP clients will interact correctly.
- It MAY offer itself as a destination and resolve the request using its own NBMA address, if it has the related capabilities.
- If the indicated VPN-ID identifies the default routing instance of the destination, the NHS MAY accept the request and send a corresponding NHRP Resolution Reply.

The NHRP Device Capabilities extension SHOULD NOT be included in the NHRP Register Request and Reply messages.

3.4 Error handling procedures

If an NHRP entity receives a PDU with a VPN-ID indicated via VPN encapsulation which is in conflict to a VPN-ID earlier allocated to that communication (e.g. via VPN signalling or administratively via configuration), it SHOULD send back an NHRP error indication (see 5.2.7 of [1]) to the sender indicating error code 16 (VPN mismatch). However, in order to avoid certain security issues, an NHRP entity MAY instead silently drop the packet.

If a VPN-aware NHRP entity receives a packet for a VPN that it does not support, it SHOULD send back an NHRP error indication to the sender with an error code 17 (VPN not supported). However, in order to avoid certain security issues, an NHRP entity MAY instead silently drop the packet.

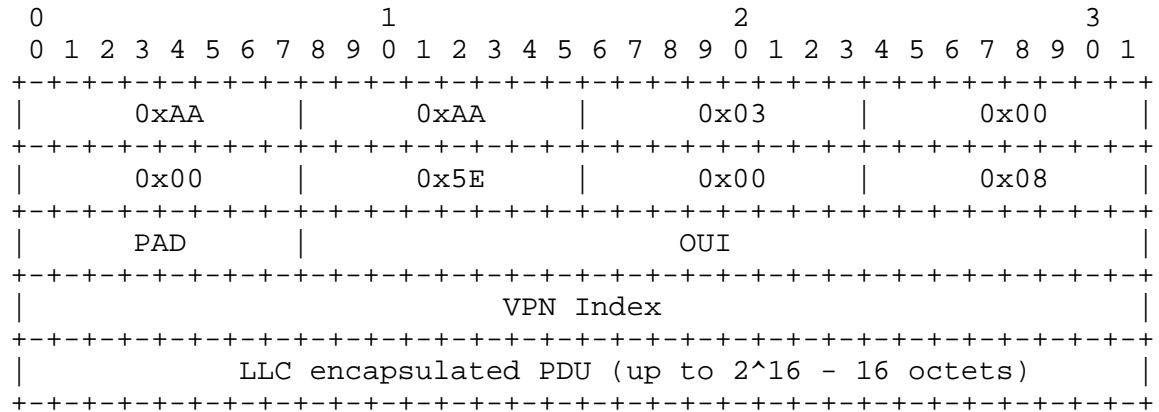
If a VPN-aware NHS cannot find a route to forward a VPN-related NHRP message, it SHOULD send back an NHRP error indication to the sender with error code 6 (protocol address unreachable). However, in order to avoid certain security issues, an NHRP entity MAY instead silently drop the packet.

In all cases, where an NHRP error indication is returned by a VPN-aware NHRP entity, the incorrect VPN-ID related to this indication SHALL be indicated via VPN encapsulation or VPN signalling, except when sending it to a non-VPN-aware NHRP device (see 3.1 / 3.2 above).

4. NHRP Packet Formats

4.1 VPN encapsulation

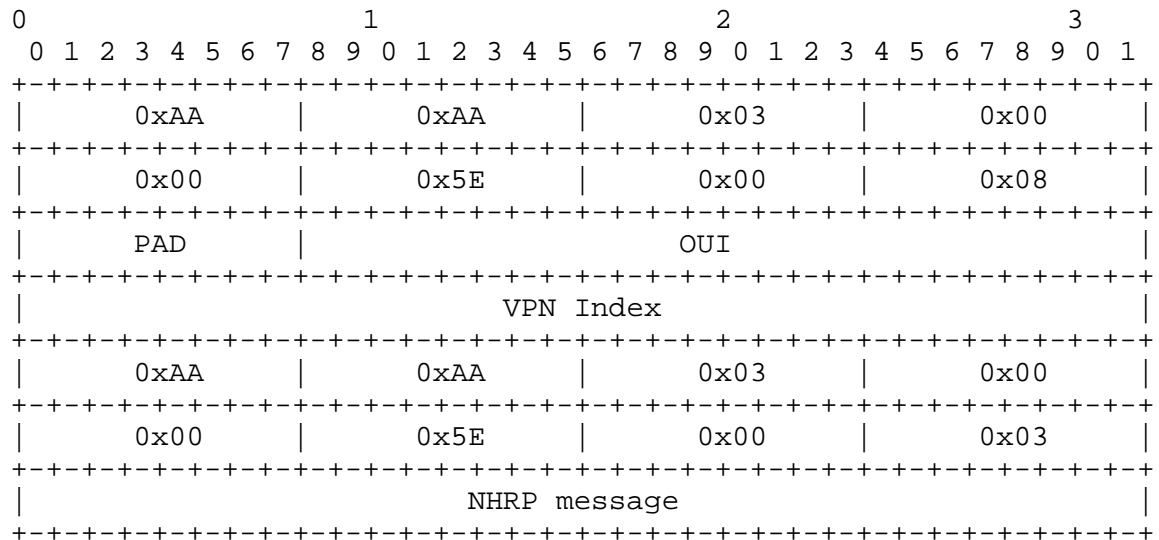
The format of the VPN encapsulation header is as follows:




It consists of the following parts:

- LLC/SNAP indication (0xAA-AA-03)
- OUI (of IANA) (0x00-00-5E)
- PID allocated by IANA for VPN encapsulation (0x00-08)
- PAD field (inserted for 32-bit alignment)
this field is coded as 0x00, and is ignored on receipt
- VPN related OUI (see [3])
- VPN Index (see [3]).

When this encapsulation header is used, the remainder of the PDU MUST be structured according to the appropriate LLC/SNAP format (i.e. that would have been used without the additional VPN encapsulation header). Correspondingly, the following figure shows how NHRP messages are transferred using VPN encapsulation:

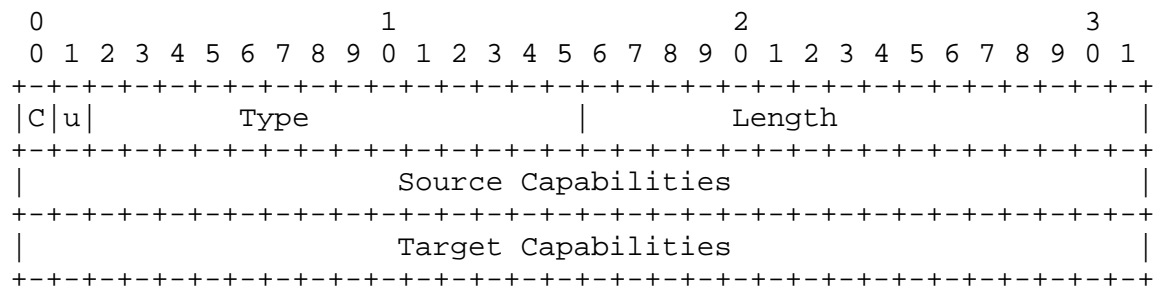


The following example shows how IP packets are transferred by VPN encapsulation:



4.2 NHRP device capabilities extension

The format of the NHRP device capabilities extension is as follows:



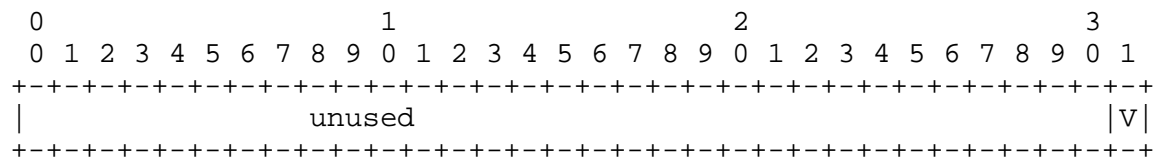
C: Compulsory = 0 (not a compulsory extension)

u: Unused and MUST be set to zero.

Type = 0x0009

Length = 0x0008

Source Capabilities field:



V bit:

0x0 - the source NHRP device is non-VPN-aware

0x1 - the source NHRP device is VPN-aware

The unused bits MUST be set to zero on transmission and ignored on receipt.

Target Capabilities field:

```

      0               1               2               3
    0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|               unused               |V|
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+

```

V bit:

- 0x0 - the destination NHRP device is non-VPN-aware
- 0x1 - the destination NHRP device is VPN-aware

The unused bits MUST be set to zero on transmission and ignored on receipt.

4.3 Error Codes

The following further Error Codes are defined in addition to those specified in section 5.2.7 of [1]):

16 - VPN mismatch

This error code is returned by a VPN-capable NHRP device, if it receives a PDU with a VPN-ID in the LLC/SNAP header different from the VPN-ID which had been specified earlier via VPN signalling.

17 - VPN not supported

This error code is returned by a VPN-capable NHRP device, if it receives an NHRP message for a VPN that it does not support.

5. Security Considerations

For any VPN application, it is important that VPN-related information is not misdirected to other VPNs and is not accessible when being transferred across a public or shared infrastructure. It is therefore RECOMMENDED to use the VPN support functions specified in this document in combination with NHRP authentication as specified in section 5.3.4 of [1]. Section 5.3.4.4 of [1] also provides further information on general security considerations related to NHRP.

In cases where the NHRP entity does not trust all of the NHRP entities, or is uncertain about the availability of the end-to-end NHRP authentication chain, it may use IPsec for confidentiality, integrity, etc.

6. IANA Considerations

The LLC/SNAP protocol ID 0x00-08 for VPN encapsulation had already been allocated by IANA in conjunction with [2]. This specification does not require the allocation of any additional LLC/SNAP protocol IDs beyond that.

It should be noted that IANA - as the owner of the VPN-related OUI: 0x00-00-5E - is itself also a VPN authority which may allocate VPN indices to identify VPNs. The use of these particular VPN indices within the context of this specification is reserved, and requires allocation and approval by the IESG in accordance with RFC 2434.

References

- [1] Luciani, J., Katz, D., Piscitello, D., Cole, B. and N. Doraswamy, "NMBA Next Hop Resolution Protocol (NHRP)", RFC 2332, April 1998.
- [2] Grossman, D. and J. Heinanen, "Multiprotocol Encapsulation over ATM Adaptation Layer 5", RFC 2684, September 1999.
- [3] Fox, B. and B. Gleeson, "Virtual Private Networks Identifier", RFC 2685, September 1999.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Authors' Addresses

Barbara A. Fox
Equipe Communications
100 Nagog Park
Acton, MA 01720

Phone: +1-978-795-2009
EMail: bfox@equipecom.com

Bernhard Petri
Siemens AG
Hofmannstr. 51
Munich, Germany, D-81359

Phone: +49 89 722-34578
EMail: bernhard.petri@icn.siemens.de

Full Copyright Statement

Copyright (C) The Internet Society (1999). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

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

