

Routing over Large Clouds Working Group Dave
Katz
INTERNET-DRAFT (cisco
Systems)
<[draft-ietf-rolc-nhrp-07.txt](#)> David
Piscitello
Inc.) (Core Competence,
Bruce
Cole
Systems) (cisco
Luciani James V.
Nexion) (Ascom
Expires June
1996

NBMA Next Hop Resolution Protocol (NHRP)

Status of this Memo

This document is an Internet-Draft. Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as ``work in progress.''

To learn the current status of any Internet-Draft, please check the ``[id-abstracts.txt](#)'' listing contained in the Internet-Drafts Shadow

Directories on [ds.internic.net](#) (US East Coast), [nic.nordu.net](#) (Europe), [ftp.isi.edu](#) (US West Coast), or [munnari.oz.au](#) (Pacific Rim).

Abstract

This document describes the NBMA Next Hop Resolution Protocol (NHRP).

NHRP can be used by a source station (host or router) connected to a Non-Broadcast, Multi-Access (NBMA) subnetwork to determine the internetworking layer address and NBMA subnetwork addresses of the "NBMA next hop" towards a destination station. If the destination is connected to the NBMA subnetwork, then the NBMA next hop is the destination station itself. Otherwise, the NBMA next hop is the egress router from the NBMA subnetwork that is "nearest" to the destination station. NHRP is intended for use in a multiprotocol

internetworking layer environment over NBMA subnetworks.

This document is intended to be a functional superset of the NBMA Address Resolution Protocol (NARP) documented in [[1](#)].

Katz, Piscitello, Cole, Luciani
1]

[Page

Operation of NHRP as a means of establishing a transit path across an NBMA subnetwork between two routers will be addressed in a separate document.

1. Introduction

The NBMA Next Hop Resolution Protocol (NHRP) allows a source station (a host or router), wishing to communicate over a Non-Broadcast, Multi-Access (NBMA) subnetwork, to determine the internetworking layer addresses and NBMA addresses of suitable "NBMA next hops" toward a destination station. A subnetwork can be non-broadcast either because it technically doesn't support broadcasting (e.g., an X.25 subnetwork) or because broadcasting is not feasible for one reason or another (e.g., an SMDS multicast group or an extended Ethernet would be too large). If the destination is connected to the

NBMA subnetwork, then the NBMA next hop is the destination station itself. Otherwise, the NBMA next hop is the egress router from the NBMA subnetwork that is "nearest" to the destination station.

An NBMA subnetwork may, in general, consist of multiple Local Address

Groups (LAGs). In the case of IP, a logically independent IP subnet (LIS) is an example of a LAG. LISs, as defined in [3] and [4], have the following properties:

- 1) All members of a LIS have the same IP network/subnet number and address mask.
- 2) All members within a LIS are directly connected to the same NBMA subnetwork.
- 3) All members outside of the LIS are accessed via a router.

Address resolution as described in [3] and [4] only resolves the next hop address if the destination station is a member of the same LIS as

the source station; otherwise, the source station must forward packets to a router that is a member of multiple LIS's. In multi-LIS

configurations, hop-by-hop address resolution may not be sufficient to resolve the "NBMA next hop" toward the destination station, and

IP packets may traverse the NBMA subnetwork more than once.

NHRP describes a next hop resolution method that relaxes the forwarding restrictions of the LIS model. With NHRP when the internetwork layer address is of type IP, once the NBMA next hop has been resolved, the source may either start sending IP packets to the

destination (in a connectionless NBMA subnetwork such as SMDS) or
may
first establish a connection to the destination with the desired
bandwidth and QOS characteristics (in a connection-oriented NBMA

Katz, Piscitello, Cole, Luciani
2]

[Page

subnetwork such as ATM).

NHRP in its most basic form provides a simple internetworking layer to NBMA subnetwork layer address binding service. This may be sufficient for hosts which are directly connected to an NBMA subnetwork, allowing for straightforward implementations in NBMA stations. NHRP also has the capability of determining the egress point from an NBMA subnetwork when the destination is not directly connected to the NBMA subnetwork and the identity of the egress router is not learned by other methods (such as routing protocols). Optional extensions to NHRP provide additional robustness and diagnosability.

Address resolution techniques such as those described in [3] and [4] may be in use when NHRP is deployed. ARP servers and services over NBMA subnetworks may be required to support hosts that are not capable of dealing with any model for communication other than the LIS model, and deployed hosts may not implement NHRP but may continue

to support ARP variants such as those described in [3] and [4].

NHRP

is intended to reduce or eliminate the extra router hops required by the LIS model, and can be deployed in a non-interfering manner alongside existing ARP services.

The operation of NHRP to establish transit paths across NBMA subnetworks between two routers requires additional mechanisms to avoid stable routing loops, and will be described in a separate document.

2. Overview

2.1 Terminology

The term "network" is highly overloaded, and is especially confusing in the context of NHRP. We use the following terms:

Internetwork layer--the media-independent layer (IP in the case of TCP/IP networks).

Subnetwork layer--the media-dependent layer underlying the internetwork layer, including the NBMA technology (ATM, X.25, SMDS, etc.)

2.2 Protocol Overview

In this section, we briefly describe how a source S (which potentially can be either a router or a host) uses NHRP to determine

the "NBMA next hop" to destination D.

For administrative and policy reasons, a physical NBMA subnetwork may be partitioned into several, disjoint "Logical NBMA subnetworks". A Logical NBMA subnetwork is defined as a collection of hosts and routers that share unfiltered subnetwork connectivity over an NBMA subnetwork. "Unfiltered subnetwork connectivity" refers to the absence of closed user groups, address screening or similar features that may be used to prevent direct communication between stations connected to the same NBMA subnetwork. (Hereafter, unless otherwise specified, we use the term "NBMA subnetwork" to mean *logical* NBMA subnetwork.)

Placed within the NBMA subnetwork are one or more entities that implement the NHRP protocol. Such stations which are capable of answering Next Hop Resolution Requests are known as "Next Hop Servers" (NHSs). Each NHS serves a set of destination hosts, which may or may not be directly connected to the NBMA subnetwork. NHSs cooperatively resolve the NBMA next hop within their logical NBMA subnetwork. In addition to NHRP, NHSs may participate in protocols used to disseminate routing information across (and beyond the boundaries of) the NBMA subnetwork, and may support "classical" ARP service as well.

An NHS maintains a "next-hop resolution" cache, which is a table of address mappings (internetwork layer address to NBMA subnetwork layer address). This table can be constructed from information gleaned from NHRP Register packets (see [Section 5.2.3](#) and 5.2.4), extracted from Next Hop Resolution Requests/Replies that traverse the NHS as they are forwarded, or through mechanisms outside the scope of this document (examples of such mechanisms include ARP [[2](#), [3](#), [4](#)] and pre-configured tables). [Section 6.2](#) further describes cache management issues.

A host or router that is not an NHRP server must be configured with the identity of the NHS which serves it (see Configuration, [Section 4](#)).

[Note: for NBMA subnetworks that offer group or multicast addressing features, it may be desirable to configure stations with a group identity for NHSs, i.e., addressing information that would solicit a response from "all NHSs". The means whereby a group of NHSs divide responsibilities for next hop resolution are not described here.]

Whether or not a particular station within the NBMA subnetwork which is making use of the NHRP protocol needs to be able to act as an NHS is a local matter. For a station to avoid providing NHS functionality, there must be one or more NHSs within the NBMA

subnetwork which are providing authoritative NBMA information on its behalf. If NHRP is to be able to resolve the NBMA address for stations that lack NHS functionality, these serving NHSs must exist along all routed paths between Next Hop Resolution Requesters and the station which cannot answer Next Hop Resolution Requests.

The protocol proceeds as follows. An event occurs triggering station

S to want to resolve the NBMA address of a path to D. This is most likely to be when a data packet addressed to station D is to be emitted from station S (either because station S is a host, or station S is a transit router), but the address resolution could also

be triggered by other means (a routing protocol update packet, for example). Station S first determines the next hop to station D through normal routing processes (for a host, the next hop may simply

be the default router; for routers, this is the "next hop" to the destination internetwork layer address). If the next hop is reachable through its NBMA interface, S constructs a Next Hop Resolution Request packet (see [Section 5.2.1](#)) containing station D's internetwork layer address as the (target) destination address, S's own internetwork layer address as the source address (Next Hop Resolution Request initiator), and station S's NBMA addressing information. Station S may also indicate that it prefers an authoritative Next Hop Resolution Reply (i.e., station S only wishes to receive a Next Hop Resolution Reply from the NHS-speaker that maintains the NBMA-to-internetwork layer address mapping for this destination). Station S emits the Next Hop Resolution Request packet

towards the destination, using the NBMA address of the next routed hop.

If the Next Hop Resolution Request is triggered by a data packet, station S may choose to dispose of the data packet while awaiting a Next Hop Resolution Reply in one of the following ways:

- (a) Drop the packet
- (b) Retain the packet until the Next Hop Resolution Reply arrives and a more optimal path is available
- (c) Forward the packet along the routed path toward D

The choice of which of the above to perform is a local policy matter,

though option (c) is the recommended default, since it may allow data

to flow to the destination while the NBMA address is being resolved.

Note that an Next Hop Resolution Request for a given destination MUST

NOT be triggered on every packet, though periodically retrying a Next

Hop Resolution Request is permitted.

When the NHS receives an Next Hop Resolution Request, a check is made to see if it "serves" station D, i.e., the NHS checks to see if there is a "next hop" entry for D in its next-hop resolution cache. If the

Katz, Piscitello, Cole, Luciani
5]

[Page

NHS does not serve D, the NHS forwards the Next Hop Resolution Request to another NHS. (Mechanisms for determining how to forward the Next Hop Resolution Request are discussed in [Section 3](#), Deployment.) Note that NHSs must be next hops to one another in order for forwarding of NHRP packets to be possible.

If this NHS serves D, the NHS resolves station D's NBMA address, and generates a positive Next Hop Resolution Reply on D's behalf. (Next Hop Resolution Replies in this scenario are always marked as "authoritative".) The Next Hop Resolution Reply packet contains the next hop internetwork layer address and the NBMA address for station D and is sent back to S. (Note that if station D is not on the NBMA subnetwork, the next hop internetwork layer address will be that of the egress router through which packets for station D are forwarded.)

An NHS receiving a Next Hop Resolution Reply may cache the NBMA next hop information contained therein. To a subsequent Next Hop Resolution Request, this NHS may respond with the cached, non-authoritative, NBMA next hop information or with cached negative information, if the NHS is allowed to do so, see [section 6.2](#). Non-authoritative Next Hop Resolution Replies are distinguished from authoritative Next Hop Resolution Replies so that if a communication attempt based on non-authoritative information fails, a source station can choose to send an authoritative Next Hop Resolution Request. NHSs MUST NOT respond to authoritative Next Hop Resolution Requests with cached information.

[Note: An Next Hop Resolution Reply can be returned directly to the Next Hop Resolution Request initiator, i.e., without traversing the list of NHSs that forwarded the Next Hop Resolution Request, if all of the following criteria are satisfied:

- (a) Direct communication is available via datagram transfer (e.g., SMDS) or the NHS has an existing virtual circuit connection to the Next Hop Resolution Request initiator or is permitted to open one.
- (b) The Next Hop Resolution Request initiator has not included the NHRP Reverse NHS record Extension (see [Section 5.3.5](#)).
- (c) The authentication policy in force permits direct communication between the NHS and the Next Hop Resolution Request initiator.

The purpose of allowing an NHS to send a Next Hop Resolution Reply directly is to reduce response time. A consequence of allowing a direct Next Hop Resolution Reply is that NHSs that would under

normal circumstances be traversed by the Next Hop Resolution Reply would not cache next hop information contained therein.]

Katz, Piscitello, Cole, Luciani
6]

[Page

The process of forwarding the Next Hop Resolution Request is repeated until the Next Hop Resolution Request is satisfied, or an error occurs (e.g., no NHS in the NBMA subnetwork can resolve the Next Hop Resolution Request.) If the determination is made that station D's next hop cannot be resolved, a negative Next Hop Resolution Reply (NAK) is returned. This occurs when (a) no next-hop resolution information is available for station D from any NHS, or (b) an NHS is unable to forward the Next Hop Resolution Request (e.g., connectivity is lost).

NHRP Registration Requests, NHRP Registration Replies, NHRP Purge Requests, NHRP Purge Replies, and NHRP Error Indications follow the routed path from sender to receiver in the same fashion that Next Hop Resolution Requests and Next Hop Resolution Replies do. That is, "requests" and "indications" follow the routed path from Source Protocol Address (which is the address of the station initiating the communication) to the Destination Protocol Address. "Replies", on the other hand, follow the routed path from the Destination Protocol Address back to the Source Protocol Address with the exceptions mentioned above where a direct VC may be created.

Next Hop Resolution Requests and Next Hop Resolution Replies MUST NOT cross the borders of a logical NBMA subnetwork (an explicit NBMA subnetwork identifier may be included as an extension in the Next Hop Resolution Request, see [section 5.3.2](#)). Thus, the internetwork layer traffic out of and into a logical NBMA subnetwork always traverses an internetwork layer router at its border. Internetwork layer filtering can then be implemented at these border routers.

NHRP optionally provides a mechanism to send a Next Hop Resolution Reply which contains aggregated NBMA next hop information. Suppose that router X is the NBMA next hop from station S to station D. Suppose further that X is an egress router for all stations sharing an internetwork layer address prefix with station D. When an Next Hop Resolution Reply is generated in response to a Next Hop Resolution Request, the responder may augment the internetwork layer address of station D with a prefix length (see [Section 5.3.1](#)). A subsequent (non-authoritative) Next Hop Resolution Request for some destination that shares an internetwork layer address prefix (for the number of bits specified in the prefix length) with D may be satisfied with this cached information. See [section 6.2](#) regarding caching issues.

To dynamically detect subnetwork-layer filtering in NBMA subnetworks (e.g., X.25 closed user group facility, or SMDS address screens), as well as to provide loop detection and diagnostic capabilities, a "Route Record" may be included in NHRP packets (see Sections [5.3.4](#) and 5.3.5). The Route Record extensions contain the internetwork

(and subnetwork layer) addresses of all intermediate NHSs between source and destination (in the forward direction) and between destination and source (in the reverse direction). When a source station is unable to communicate with the responder (e.g., an attempt to open an SVC fails), it may attempt to do so successively with other subnetwork layer addresses in the Route Record until it succeeds (if authentication policy permits such action). This approach can find a suitable egress point in the presence of subnetwork-layer filtering (which may be source/destination sensitive, for instance, without necessarily creating separate logical NBMA subnetworks) or subnetwork-layer congestion (especially in connection-oriented media).

3. Deployment

Next Hop Resolution Requests traverse one or more hops within an NBMA subnetwork before reaching the station that is expected to generate a response. Each station, including the source station, chooses a neighboring NHS to which it will forward the Next Hop Resolution Request. The NHS selection procedure typically involves performing a routing decision based upon the network layer destination address of the Next Hop Resolution Request. Ignoring error situations, the Next Hop Resolution Request eventually arrives at a station that is to generate an Next Hop Resolution Reply. This responding station either serves the destination, or is the destination itself if both NHRP client and server functionality are co-resident in the same station. The responding station generates a Next Hop Resolution Reply using the source address from within the NHRP packet to determine where the Next Hop Resolution Reply should be sent.

The Next Hop Resolution Request packet is carried at the NBMA layer, with a destination NBMA address set to that of the locally determined NHS. If the addressed NHS does not serve the destination address specified in the Next Hop Resolution Request, the Next Hop Resolution Request packet is routed at the network layer based upon the Next Hop Resolution Requester's destination address, and forwarded to the neighboring NHS determined by the routing decision. Alternately, the NHS may use static configuration information in order to determine to which neighboring NHSs to forward the Next Hop Resolution Request packet. Each NHS/router examines the Next Hop Resolution Request packet on its way toward the destination, optionally modifying it on

the way (such as updating the Forward Record extension), and continues to forward it until it reaches the NHS that serves the destination network layer address.

In order to forward NHRP packets to a neighboring NHS, NHRP clients must nominally be configured with the NBMA address of at least one

NHS. In practice, a client's default router should also be its NHS. A client may be able to derive the NBMA address of its NHS from the configuration that was already required for the client to be able to communicate with its next hop router.

Forwarding of NHRP packets within an NBMA subnetwork requires a contiguous deployment of NHRP capable stations. During migration to NHRP, it cannot be expected that all stations within the NBMA subnetwork are NHRP capable. NHRP traffic which would otherwise need to be forwarded through such stations can be expected to be dropped due to the NHRP packet being unrecognized. In this case, NHRP will be unable to establish any transit paths whose discovery requires the traversal of the non-NHRP speaking stations. If the client has tried and failed to acquire a cut through path the the client should use the network layer routed path as a default.

The path taken by Next Hop Resolution Requests will normally be the same as the path taken by data packets which are routed at the network layer to the desired destination. (The paths may be different in situations where NHSs have been statically configured to forward traffic by other means. For example, an Next Hop Resolution Request may be forwarded to a group multicast address.)

NHSs should acquire knowledge about destinations other NHSs serve as a direct consequence of participating in intra-domain and inter-domain routing protocol exchange. In this case, the NHS serving a particular destination must lie along the routed path to that destination. In practice, this means that all egress routers must double as NHSs serving the destinations beyond them, and that hosts on the NBMA subnetwork are served by routers that double as NHSs.

NHSs (and end stations) may alternately be statically configured with the NBMA addresses of their neighbors, the identities of the destinations that each of them serves, and optionally a logical NBMA subnetwork identifier. Such static configurations may be necessary in cases where NHSs do not contain network layer routing protocol implementations.

If the NBMA subnetwork offers a link layer group addressing or multicast feature, the client (station) may be configured with a group address assigned to the group of next-hop servers. The client might then submit Next Hop Resolution Requests to the group address, eliciting a response from one or more NHSs, depending on the response strategy selected. Note that the constraints described in [Section 2](#) regarding directly sending Next Hop Resolution Reply may apply.

NHSs may also be deployed with the group or multicast address of their peers, and an NHS might use this as a means of forwarding Next

Katz, Piscitello, Cole, Luciani
9]

[Page

Hop Resolution Requests it cannot satisfy to its peers. This might elicit a response (to the NHS) from one or more NHSs, depending on the response strategy. The NHS would then forward the Next Hop Resolution Reply to the Next Hop Resolution Request originator. The purpose of using group addressing or a similar multicast mechanism in this scenario would be to eliminate the need to preconfigure each NHS in a logical NBMA subnetwork with both the individual identities of other NHSs as well as the destinations they serve. It reduces the number of NHSs that might be traversed to process an Next Hop Resolution Request (in those configurations where NHSs either respond or forward via the multicast, only two NHSs would be traversed), and allows the NHS that serves the Next Hop Resolution Request originator to cache next hop information associated with the Next Hop Resolution Reply (again, within the constraints described in [Section 2](#)).

4. Configuration

Stations

To participate in NHRP, a station connected to an NBMA subnetwork should be configured with the NBMA address(es) of its NHS(s) (alternatively, it should be configured with a means of acquiring them, i.e., the group address that members of a NHS group use for the purpose of address or next-hop resolution.) The NHS(s) will likely also represent the station's default or peer routers, so their NBMA addresses may be obtained from the station's existing configuration. If the station is attached to several subnetworks (including logical NBMA subnetworks), the station should also be configured to receive routing information from its NHS(s) and peer routers so that it can determine which internetwork layer networks are reachable through which subnetworks.

Next Hop Servers

An NHS is configured with knowledge of its own internetwork layer and NBMA addresses, a set of internetwork layer address prefixes that correspond to the internetwork layer addresses of the stations it serves, and a logical NBMA subnetwork identifier (see [Section 5.3.2](#)). If a served station is attached to several subnetworks, the NHS may also need to be configured to advertise routing information to such stations.

If an NHS acts as an egress router for stations connected to other subnetworks than the NBMA subnetwork, the NHS must, in addition to the above, be configured to exchange routing information between

the NBMA subnetwork and these other subnetworks.

Katz, Piscitello, Cole, Luciani
10]

[Page

In all cases, routing information is exchanged using conventional intra-domain and/or inter-domain routing protocols.

The NBMA addresses of the stations served by the NHS may be learned via NHRP Register packets or manual configuration.

5. NHRP Packet Formats

This section describes the format of NHRP packets.

An NHRP packet consists of a Fixed Part, a Mandatory Part, and an Extensions Part. The Fixed Part is common to all NHRP packet types. The Mandatory Part **MUST** be present, but varies depending on packet type. The Extensions Part also varies depending on packet type, and need not be present.

The length of the Fixed Part is fixed at 20 octets. The length of the Mandatory Part is determined by the contents of the extensions offset field (ar\$extoff). If ar\$extoff=0x0 then the mandatory part length is equal to total packet length (ar\$pktsiz) minus 20 otherwise the mandatory part length is equal to ar\$extoff minus 20. The length of the Extensions Part is implied by ar\$pktsiz minus ar\$extoff minus 20. NHSs may increase the size of an NHRP packet as a result of extension processing, but not beyond the offered maximum SDU size of the NBMA network.

NHRP packets are encapsulated using the native formats used on the particular NBMA network over which NHRP is carried. For example, SMDS networks always use LLC/SNAP encapsulation at the NBMA layer, and an NHRP packet is preceded by the following LLC/SNAP encapsulation:

```
[0xAA-AA-03] [0x00-00-5E] [0x00-03]
```

The first three octets are LLC, indicating that SNAP follows. The SNAP OUI portion is the IANA's OUI, and the SNAP PID portion identifies NHRP (see [4]).

ATM uses either LLC/SNAP encapsulation of each packet (including NHRP), or uses no encapsulation on VCs dedicated to a single protocol (see [7]). Frame Relay and X.25 both use NLPID/SNAP encapsulation or identification of NHRP, using a NLPID of 0x0080 and the same SNAP contents as above (see [8], [9]).

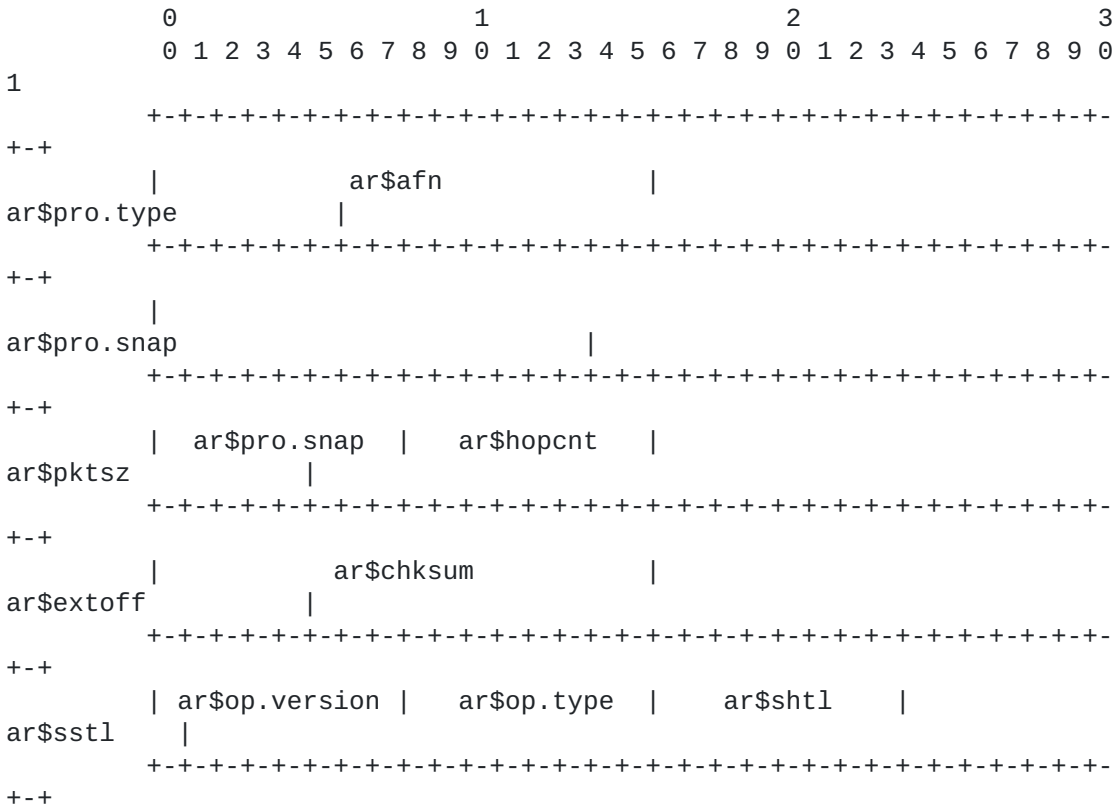
Fields marked "unused" **MUST** be set to zero on transmission, and ignored on receipt.

Most packet types (ar\$op.type) have both internetwork layer protocol-independent fields and protocol-specific fields. The

protocol-independent fields always come first in the packet, and the protocol type/snap fields (ar\$pro.type/snap) qualify the format of the protocol-specific fields.

5.1 NHRP Fixed Header

The Fixed Part of the NHRP packet contains those elements of the NHRP packet which are always present and do not vary in size with the type of packet.



ar\$afn
Defines the type of "link layer" addresses being carried. This number is taken from the 'address family number' list specified in [6]. This field has implications to the coding of ar\$shtl and ar\$sstl as described below.

ar\$pro.type
field is a 16 bit unsigned integer representing the following number space:

0x0000 to 0x00FF Protocols defined by the equivalent NLPIDs.

0x0100 to 0x03FF Reserved for future use by the IETF.
0x0400 to 0x04FF Allocated for use by the ATM Forum.
0x0500 to 0x05FF Experimental/Local use.
0x0600 to 0xFFFF Protocols defined by the equivalent
Ethertypes.

(based on the observations that valid Ethertypes are never smaller
than 0x600, and NLPIDs never larger than 0xFF.)

ar\$pro.snap

When ar\$pro.type has a value of 0x0080, a SNAP encoded extension
is

being used to encode the protocol type. This snap extension is placed in the ar\$pro.snap field. This is termed the 'long form' protocol ID. If ar\$pro != 0x0080 then the ar\$pro.snap field MUST

be

zero on transmit and ignored on receive. The ar\$pro.type field itself identifies the protocol being referred to. This is termed the 'short form' protocol ID.

In all cases, where a protocol has an assigned number in the ar\$pro.type space (excluding 0x0080) the short form MUST be used when transmitting NHRP messages. Additionally, where a protocol

has

valid short and long forms of identification, receivers MAY choose to recognize the long form.

ar\$hopcnt

The Hop count indicates the maximum number of NHSS that an NHRP packet is allowed to traverse before being discarded.

ar\$pktsz

The total length of the NHRP packet, in octets (excluding link layer encapsulation).

ar\$chksum

The standard IP checksum over the entire NHRP packet (starting

with

the fixed header). If only the hop count field is changed, the checksum is adjusted without full recomputation. The checksum is completely recomputed when other header fields are changed.

ar\$extoff

This field identifies the existence and location NHRP extensions. If this field is 0 then no extensions exist otherwise this field represents the offset from the beginning of the NHRP packet (i.e., starting from the ar\$afn field) of the first extension.

ar\$op.version

This field is set to 0x0001 for NHRP version 1.

ar\$op.type

This is the NHRP packet type: NHRP Next Hop Resolution Request(1), NHRP Next Hop Resolution Reply(2), NHRP Registration Request(3), NHRP Registration Reply(4), NHRP Purge Request(5), NHRP Purge Reply(6), or NHRP Error Indication(7).

ar\$shtl

Type & length of source NBMA address interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for NSAP, ar\$afn=8 for E.164). When ar\$afn=0x000F (E.164 address plus NSAP subaddress) then both ar\$shtl and ar\$sstl must be coded appropriately (see below).

ar\$sstl

Type & length of source NBMA subaddress interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x000F for NSAP). When an NBMA technology has no concept


of

a subaddress, the subaddress length is always coded ar\$sstl = 0

and

no storage is allocated for the subaddress in the appropriate mandatory part.

ar\$shl, ar\$sstl, subnetwork layer addresses, and subnetwork layer subaddresses fields are coded as follows:



The most significant bit is reserved and MUST be set to zero. The second most significant bit (x) is a flag indicating whether the address being referred to is in:

- NSAP format (x = 0).
- Native E.164 format (x = 1).

For NBMA technologies that use neither NSAP nor E.164 format addresses, x = 0 SHALL be used to indicate the native form for the particular NBMA technology.

In the case where the NBMA is ATM, if a subaddress is to be included then ar\$afn MUST be set to 0x000F which means that if a subaddress exists then it is of type NSAP.

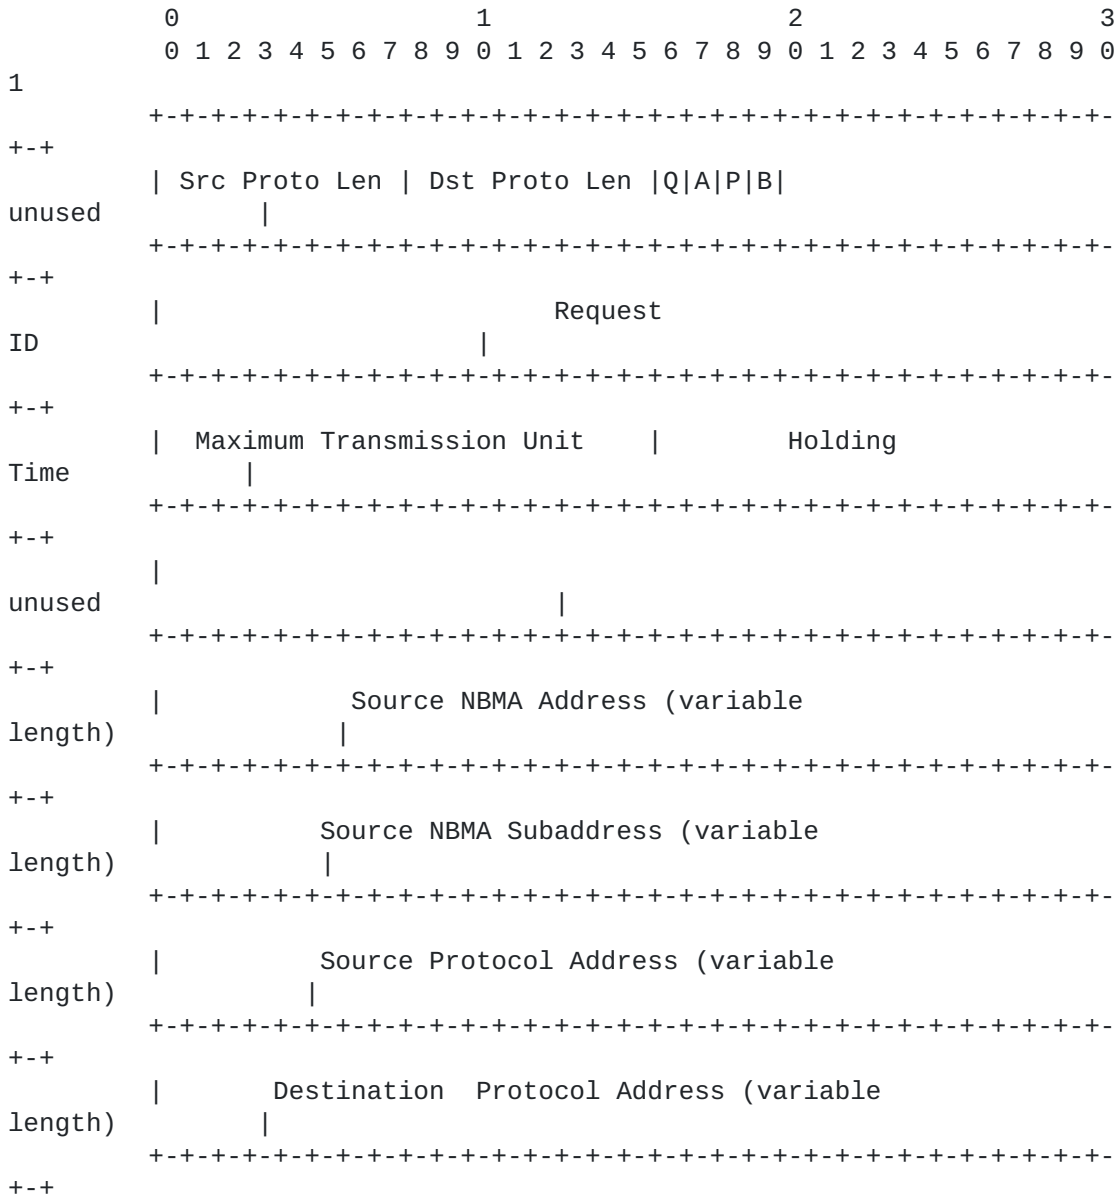
The bottom 6 bits is an unsigned integer value indicating the length of the associated NBMA address in octets. If this value is zero the flag x is ignored.

5.2 Mandatory Part

The Mandatory Part of the NHRP packet contains the operation specific information (e.g., Next Hop Resolution Request/Reply, etc.) and variable length data which is pertinent to the packet type.

5.2.1 NHRP Next Hop Resolution Request

The NHRP Next Hop Resolution Request packet has a Type code of 1. The Mandatory Part has the following format:



Src Proto Len
This field holds the length in octets of the Source Protocol Address.

Dst Proto Len
This field holds the length in octets of the Destination Protocol Address.

NHRP Next Hop Resolution Request/Reply Flags

Q
Set if the station sending the Next Hop Resolution Request is a router; clear if the it is a host.

A

bit A response to a Next Hop Resolution Request may contain cached information. If an authoritative answer is desired, then this ("Authoritative") should be set. If non-authoritative (cached) information is acceptable, this bit should be clear.

P

Unused (clear on transmit)

B

Unused (clear on transmit)

Request ID

A value which, when coupled with the address of the source,

provides a unique identifier for the information contained in a Next Hop Resolution Request and its associated Next Hop Resolution Reply, and any subsequent Purge. This value can be used by the source to aid in matching a Next Hop Resolution Request with a

Next

Hop Resolution Reply. This value could also be sent across a virtual circuit (in SVC environments) to aid in matching NHRP transactions with virtual circuits (this use is for further study).

The value is taken from a 32 bit counter that is incremented each time a new Next Hop Resolution Request is transmitted. The same value MUST be used when sending another Next Hop Resolution

Request

for the same destination when a previous Next Hop Resolution Request is still active or pending, i.e., when retransmitting a Next Hop Resolution Request because a Next Hop Resolution Reply

was

not received, or when refreshing an existing entry to avoid

holding

timer expiration. A new value MUST be used when sending a Next

Hop

Resolution Request when no cache entry is present, or a previous cache entry was deleted for any reason.

Maximum Transmission Unit

This field gives the maximum transmission unit for the target station. This field is ignored in Next Hop Resolution Requests

and

should be set to 0. A possible use of this field in the Next Hop Resolution Request packet is for the Next Hop Resolution Requester to ask for a target MTU. This use is for further study.

Holding Time

The Holding Time field specifies the number of seconds for which the client NBMA information (the information of the client issuing the Next Hop Resolution Request) is considered to be valid. The contents of this field along with the source address information MAY be cached by transit NHRPs. The holding time should be set to the remaining time left in the client's registration with its server. If this field is set to 0 then transit NHRPs MUST not

cache

the client's NBMA information.

Source NBMA Address

The Source NBMA address field is the address of the source station which is sending the Next Hop Resolution Request. If the field's length as specified in ar\$sh1 is 0 then no storage is allocated for this address at all.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the source

station which is sending the Next Hop Resolution Request. If the field's length as specified in ar\$sstl is 0 then no storage is allocated for this address at all.

Katz, Piscitello, Cole, Luciani
16]

[Page


```

+--+
length) |           Destination Protocol Address (variable
          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+--+
length) |           Next Hop NBMA Address (variable
          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+--+
length) |           Next Hop NBMA Subaddress (variable
          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+--+
length) |           Next Hop Protocol Address (variable
          |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
+--+

```

Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Next Hop Resolution Request/Reply Flags

Q

Hop

Copied from the Next Hop Resolution Request. Set if the Next Resolution Requester is a router; clear if it is a host.

A

Set if the next hop in the Next Hop Resolution Reply is authoritative; clear if the Next Hop Resolution Reply is non-authoritative.

P

Set if the Next Hop Resolution Reply is positive; clear if the Next Hop Resolution Reply is negative.

B

Hop

Set if the association between the destination and the next hop information is guaranteed to be stable for the lifetime of the information (the holding time). This is the case if the Next protocol address identifies the destination (though it may be different in value than the Destination address if the destination system has multiple addresses) or if the destination is not connected directly to the NBMA subnetwork but the egress router to that destination is guaranteed to be stable (such as when the destination is immediately adjacent to the egress

router

through a non-NBMA interface). This information affects caching strategies (see [section 6.2](#)).

An NHS is not allowed to send a Next Hop Resolution Reply to an Next Hop Resolution Request for authoritative information with cached information, but may do so for an NHRP Next Hop Resolution Request which indicates a request for non-authoritative information. An NHS may send an Next Hop Resolution Reply to an Next Hop Resolution Request for non-authoritative information with authoritative information.

Request ID

Next

A value which, when coupled with the address of the source, provides a unique identifier for the information contained in a Next Hop Resolution Request and its associated Next Hop Resolution Reply, and any subsequent Purge. This value can be used by the source to aid in matching a Next Hop Resolution Request with a Hop Resolution Reply. This value could also be sent across a virtual circuit (in SVC environments) to aid in matching NHRP

transactions with virtual circuits (this use is for further study).

The value is taken from a 32 bit counter that is incremented each time a new Next Hop Resolution Request is transmitted. The same value MUST be used when sending another Next Hop Resolution Request for the same destination when a previous Next Hop Resolution Request is still active or pending, i.e., when retransmitting a Next Hop Resolution Request because a Next Hop Resolution Reply was not received, or when refreshing an existing entry to avoid holding timer expiration. A new value MUST be used when sending a Next Hop Resolution Request when no cache entry is present, or a previous cache entry was deleted for any reason.

Maximum Transmission Unit

This field gives the maximum transmission unit for the Next Hop information supplied in the mandatory part of the packet. If this value is 0 then either the default MTU is used or the MTU negotiated via signaling is used if such negotiation is possible for the given NBMA.

Holding Time

The Holding Time field specifies the number of seconds for which the Next Hop NBMA information specified in the mandatory part of the packet is considered to be valid. Cached information SHALL be discarded when the holding time expires. This field must be set to 0 on a NAK.

NH Addr T/L

Type & length of next hop NBMA address specified in the mandatory part of the packet. This field is interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for ATM).

NH SAddr T/L

Type & length of next hop NBMA subaddress specified in the mandatory part of the packet. This field is interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0015 for ATM makes the address an E.164 and the subaddress an ATM Forum NSAP address). When an NBMA technology has no concept of a subaddress the subaddress is always null with a length of 0. When the address length is specified as 0 no storage is allocated for the address.

NH Proto Len

This field holds the length in octets of the Next Hop Protocol Address specified in the mandatory part of the packet (additional next hop entries may be specified in the Additional Next Hop Entries Extension (see [Section 5.2.9](#))).

Preference

This field specifies the preference of the Next Hop entry specified in the mandatory part of the packet. This preference value is relative to other Next Hop entries in this NHRP Next Hop

Resolution

Reply packet which may be by the Additional Next Hop Entries Extension (see [Section 5.3.9](#)) for the given internetworking protocol. Higher values indicate higher preference. Action taken when multiple next hop entries have the highest preference value

is

a local matter. Set to 0 on a NAK.

Source NBMA Address

The Source NBMA address field is the address of the source station which sent the Next Hop Resolution Request. If the field's length as specified in ar\$shtl is 0 then no storage is allocated for this address at all.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the source station which sent the Next Hop Resolution Request. If the field's length as specified in ar\$sstl is 0 then no storage is allocated for this address at all.

Source Protocol Address

This is the protocol address of the station which sent the Next Hop Resolution Request.

Destination Protocol Address

This is the protocol address of the station for which the NBMA hop is desired.

(The NBMA address/subaddress form allows combined E.164/NSAPA form of NBMA addressing. For NBMA technologies without a subaddress concept, the subaddress field is always ZERO length and ar\$sstl = 0.)

The following is the Next Hop entry as specified in the Mandatory Part of the packet:

Next Hop NBMA Address

This is the NBMA address of the station that is the next hop for packets bound for the internetworking layer address specified.

Next Hop NBMA SubAddress

This is the NBMA subaddress of the station that is the next hop for packets bound for the internetworking layer address specified.

Next Hop Protocol Address

Katz, Piscitello, Cole, Luciani
20]

[Page

This internetworking layer address specifies the next hop. This will be the address of the destination host if it is directly attached to the NBMA subnetwork, or the egress router if it is not directly attached.

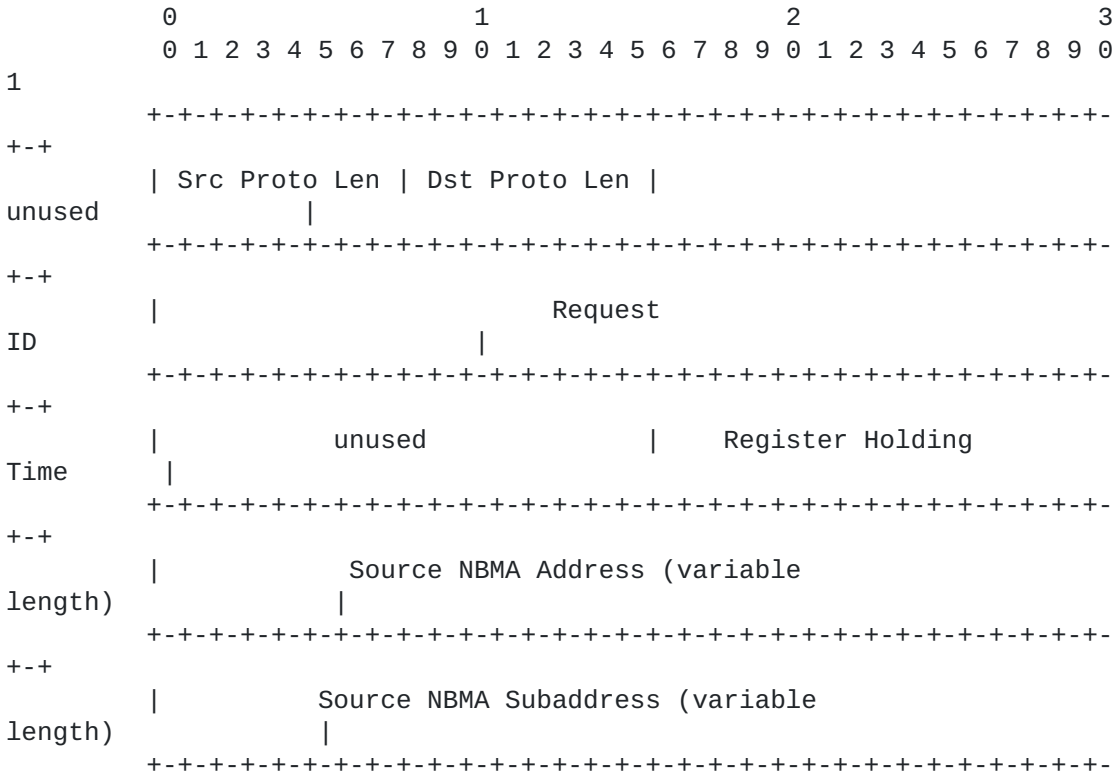
There may be multiple Next Hop entries returned in the Next Hop Resolution Reply by including the Additional Next Hop Entries Extension. See [Section 5.3.9](#) for use of these entries. The most preferable Next Hop must be specified in the mandatory part of the Next Hop Resolution Reply.

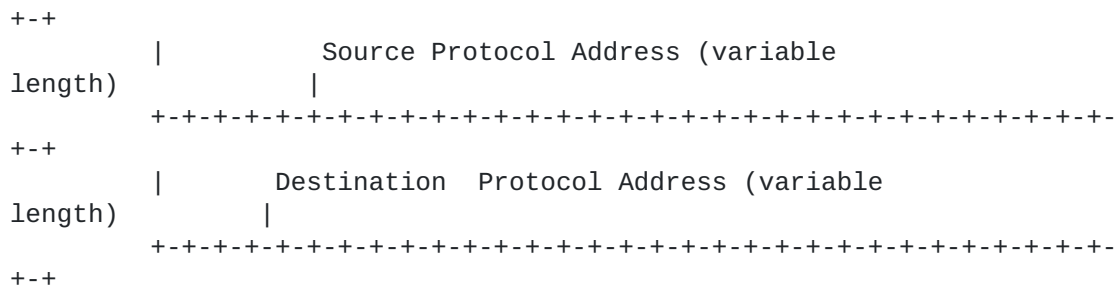
Any extensions present in the Next Hop Resolution Request packet MUST be present in the NHRP Next Hop Resolution Reply packet, except for the case of unrecognized non-Compulsory extensions.

If an unsolicited NHRP Next Hop Resolution Reply packet is received, an Error Indication of type Invalid Next Hop Resolution Reply Received SHOULD be sent in response.

5.2.3 NHRP Registration Request

The NHRP Registration Request is sent from a station to an NHS to notify the NHS of the station's NBMA information. It has a Type code of 3. The Mandatory Part has the following format:





Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Request ID

A value which, when coupled with the address of the source, provides a unique identifier for the information contained in an NHRP Registration Request packet. This value is copied directly from an NHRP Registration Request packet into the associated Registration Reply. This value could also be sent across a virtual circuit (in SVC environments) to aid in matching the NHRP transactions with virtual circuits (this use is for further study).

Register Holding Time

The Register Holding Time field specifies the number of seconds for which the registered NBMA information is considered to be valid. Cached information SHALL be discarded when the holding time expires.

Source NBMA Address

The Source NBMA address field is the address of the source station which is sending the NHRP Registration Request.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the source station which is sending the NHRP Registration Request. If the field's length as specified in ar\$sstl is 0 then no storage is allocated for this address at all.

Source Protocol Address

This is the protocol address of the station which is sending the NHRP Registration Request.

Destination Protocol Address

This is the protocol address of the NHS for which the source NBMA next hop information is being registered.

This packet is used to register a station's Protocol and NBMA addresses with its NHSs, as configured or known through conventional routing means. This allows static configuration information to be reduced; the NHSs need not be configured with the identities of all of the stations that they serve. If an NHS receives an NHRP Registration Request packet for a station that it does not serve and that packet has a Destination Protocol Address which is not the protocol address of the NHS that is currently inspecting the packet then the NHS inspecting the packet MUST forward the registration along the routed path to the Destination Protocol Address.

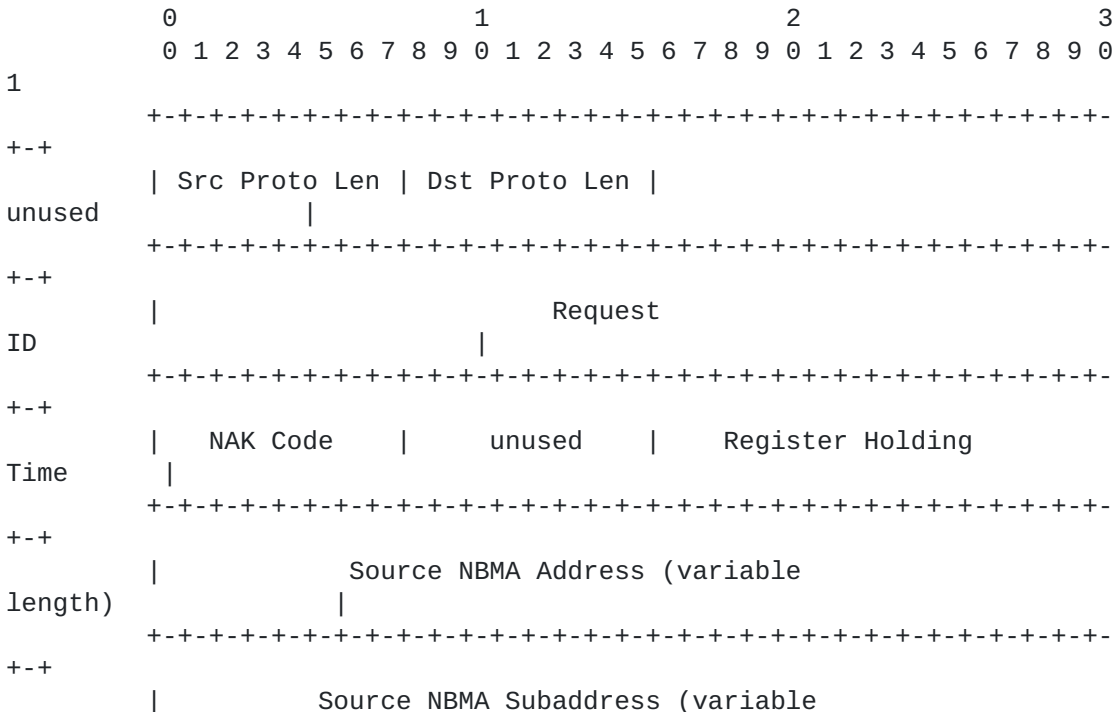
It is possible that a misconfigured station will attempt to register with the wrong NHS (i.e., one that cannot serve it due to policy constraints or routing state). If this is the case, the NHS MUST reply with a NAK-ed Registration Reply of type Can't Serve This Address.

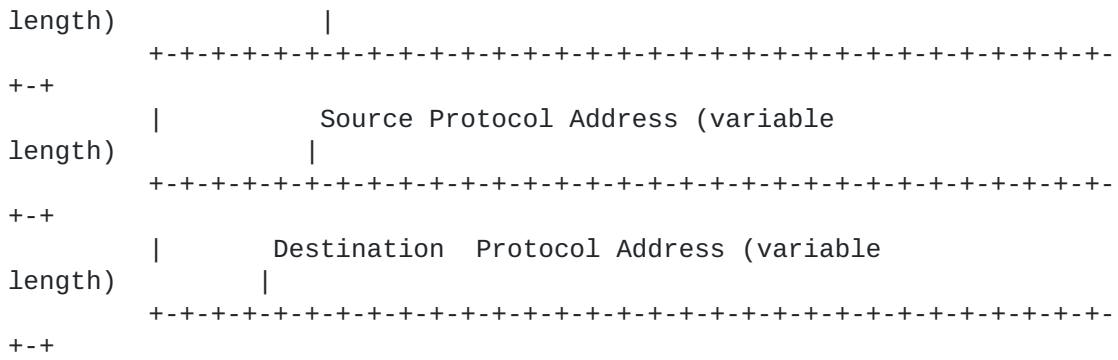
If an NHS cannot serve a station due to a lack of resources, the NHS MUST reply with a NAK-ed Registration Reply of type Registration Overflow.

In order to keep the registration entry from being discarded, the station MUST re-send the NHRP Registration Request packet often enough to refresh the registration, even in the face of occasional packet loss. It is recommended that the NHRP Registration Request packet be sent at an interval equal to one-third of the Holding Time specified therein.

5.2.4 NHRP Registration Reply

The NHRP Registration Reply is sent by an NHS to a client in response to that client's NHRP Registration Request. If the NAK Code field has anything other than 0 zero in it then the NHRP Registration Reply is a NAK otherwise the reply is an ACK. The NHRP Registration Reply has a Type code of 4. Its mandatory part has the following format:





Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Request ID

A value which, when coupled with the address of the source, provides a unique identifier for the information contained in an NHRP Registration Reply packet. This value is copied directly from an NHRP Registration Request packet into the associated NHRP Registration Reply. This value could also be sent across a virtual circuit (in SVC environments) to aid in matching NHRP transactions with virtual circuits (this use is for further study).

NAK Code

If this field is set to zero then this packet contains a positively acknowledged NHRP Registration Reply. If this field contains any other value then this contains an NHRP Registration Reply NAK which means that the internetworking layer to NBMA address binding was not stored at the client's NHS. Currently defined NAK Codes are as follows:

4 - Can't Serve This Address

An NHS may refuse an NHRP Registration Request attempt for administrative reasons. If so, the NHS MUST send an NHRP Registration Reply which contains a NAK code of 4.

5 - Registration Overflow

If an NHS cannot serve a station due to a lack of resources, the NHS MUST reply with a NAKed NHRP Registration Reply which contains a NAK code of 5.

Register Holding Time

The Register Holding Time field specifies the number of seconds for which the registered NBMA information is considered to be valid. Cached information SHALL be discarded when the holding time expires.

Source NBMA Address

The Source NBMA address field is the address of the source station which sent the Next Hop Registration Request.

Source NBMA SubAddress

The Source NBMA subaddress field is the subaddress of the source station which sent the Next Hop Registration Request. If the

field's length as specified in ar\$sstl is 0 then no storage is allocated for this address at all.

Katz, Piscitello, Cole, Luciani
24]

[Page

Source Protocol Address

This is the protocol address of the station which sent the NHRP Registration Request.

Destination Protocol Address

This is the protocol address of the NHS in which the client is attempting to register the client's NBMA information.

This packet is used to register a station's Protocol and NBMA addresses with its neighboring NHSs, as configured or known through conventional routing means. This allows static configuration information to be reduced; the NHSs need not be configured with the identities of all of the stations that they serve. If an NHS receives

an NHRP Registration Request packet for a station that it does not serve and that packet has a Destination Protocol Address which is not

the protocol address of the NHS that is currently inspecting the packet then the NHS inspecting the packet MUST forward the registration along the routed path to the Destination Protocol Address.

It is possible that a misconfigured station will attempt to send a Next Hop Registration Request to the wrong NHS (i.e., one that cannot

serve it due to policy constraints or routing state). If this is the

case, the NHS MUST reply with a NAK-ed NHRP Registration Reply of type Can't Serve This Address.

If an NHS cannot serve a station due to a lack of resources, the NHS MUST reply with a NAK-ed NHRP Registration Reply of type Registration Overflow.


In order to keep the client's registration entry in the client's NHS from being timed out, the station MUST re-send the NHRP Registration Request packet often enough to refresh the registration entry, even in the face of occasional packet loss. It is recommended that the NHRP Registration Request packet be sent at an interval equal to one-third of the Holding Time specified therein.

5.2.5 NHRP Purge Request

The NHRP Purge Request packet is sent in order to invalidate cached information in a station. The NHRP Purge Request packet has a type code of 5. The Mandatory Part has the following format:

Katz, Piscitello, Cole, Luciani
25]

[Page



Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Trgt Proto Len

This field holds the length in octets of the Target Protocol Address.

Source NBMA Address

The Source NBMA address field is the address of the source station which is sending the NHRP Purge Request.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the source station which is sending the NHRP Purge Request. If its length as specified in ar\$stl is 0 then no storage is allocated for this address at all.

Source Protocol Address

The address of the station which is sending the NHRP Purge Request.

Destination Protocol Address

The address of the station that will be receiving the NHRP Purge Request.

Target Protocol Address

The address which is to be purged from the receiver's database.

An NHRP Purge Request packet is sent from an NHS to a station to cause it to delete previously cached information. This is done when the information may be no longer valid (typically when the NHS has previously provided next hop information for a station that is not directly connected to the NBMA subnetwork, and the egress point to that station may have changed).

An NHRP Purge Request packet may also be sent from a client to an NHS with which the client had previously registered. This allows for a client to invalidate it's registration with NHRP before it would otherwise expire via the holding timer.

The station sending the NHRP Purge Request MAY periodically retransmit the NHRP Purge Request until either NHRP Purge Request is acknowledged or until the holding time of the information being purged has expired. Retransmission strategies are for further investigation.

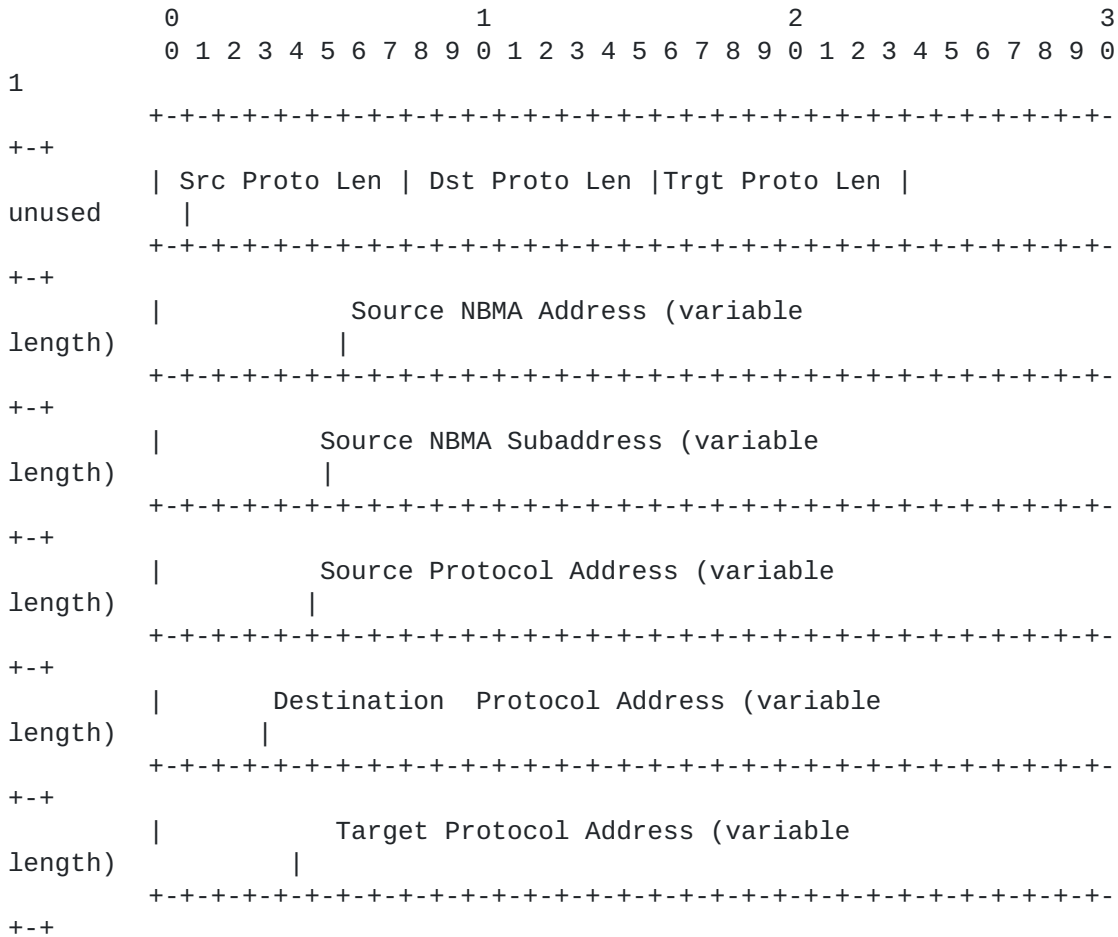
When a station receives an NHRP Purge Request, it MUST discard any previously cached information that matches the Target Protocol Address.

An NHRP Purge Reply MUST be returned for the NHRP Purge Request even if the station does not have a matching cache entry.

If the station wishes to reestablish communication with the destination shortly after receiving an NHRP Purge Request, it should make an authoritative Next Hop Resolution Request in order to avoid any stale cache entries that might be present in intermediate NHSs (See [section 6.2.2.](#)). It is recommended that authoritative Next Hop Resolution Requests be made for the duration of the holding time of the old information.

5.2.6 NHRP Purge Reply

The NHRP Purge Reply packet is sent in order to assure the sender of an NHRP Purge Request that all cached information of the specified type has been purged from the station sending the reply. The NHRP Purge packet has a type code of 6. The Mandatory Part has the following format:



Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Trgt Proto Len

This field holds the length in octets of the Target Protocol Address.

Source NBMA Address

The Source NBMA address field is the address of the source station which sent the NHRP Purge Request.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the source station which sent the NHRP Purge Request. If its length as specified in ar\$sst1 is 0 then no storage is allocated for this address at all.

Source Protocol Address

The address of the station which sent the NHRP Purge Request.

Destination Protocol Address

The address of the station which is sending the NHRP Purge Reply.

Target Protocol Address

The address which is to be purged from the receiver's database.

An NHRP Purge Request packet is sent from an NHS to a station to cause it to delete previously cached information. This is done when the information may be no longer valid (typically when the NHS has previously provided next hop information for a station that is not directly connected to the NBMA subnetwork, and the egress point to that station may have changed).

An NHRP Purge Request packet may also be sent from a client to an NHS with which the client had previously registered. This allows for a client to invalidate it's registration with NHRP before it would otherwise expire via the holding timer.

The station sending the NHRP Purge Request MAY periodically retransmit the NHRP Purge Request until it is acknowledged, or until the holding time of the information being purged has expired. Retransmission strategies are for further investigation.

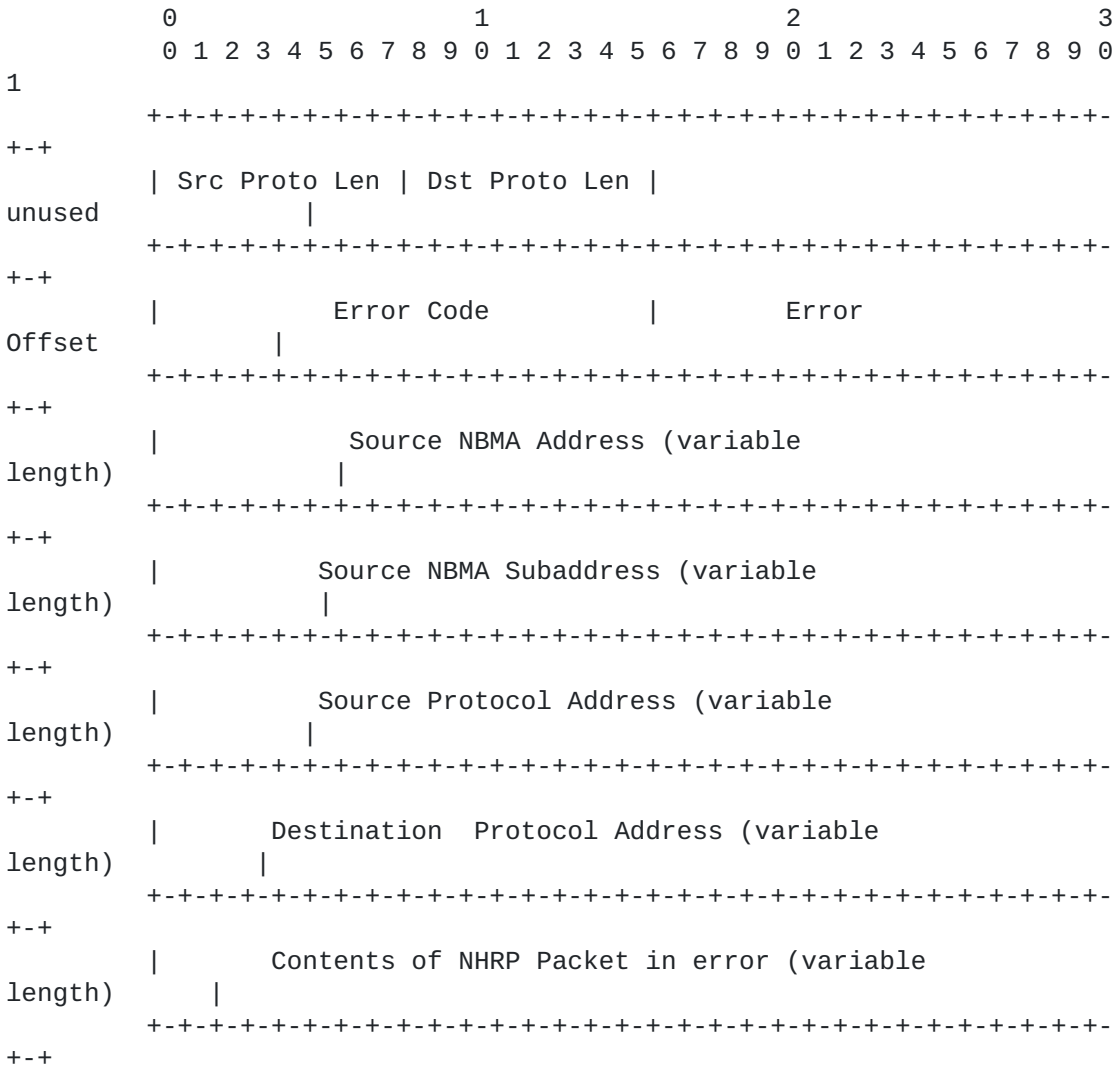
When a station receives an NHRP Purge Request, it MUST discard any previously cached information that matches the Target Protocol Address.

An NHRP Purge Reply MUST be returned as a result of receiving an NHRP Purge Request even if the station does not have a matching cache entry.

If the station wishes to reestablish communication with the destination shortly after receiving an NHRP Purge Request, it should make an authoritative Next Hop Resolution Request in order to avoid any stale cache entries that might be present in intermediate NHSs. (See [section 6.2.2.](#)) It is recommended that authoritative Next Hop Resolution Requests be made for the duration of the holding time of the old information.

5.2.7 NHRP Error Indication

The NHRP Error Indication is used to convey error indications to the sender of an NHRP packet. It has a type code of 6. The Mandatory Part has the following format:



Src Proto Len

This field holds the length in octets of the Source Protocol Address.

Dst Proto Len

This field holds the length in octets of the Destination Protocol Address.

Error Code

An error code indicating the type of error detected, chosen from the following list:

- 1 - Unrecognized Extension

When the Compulsory bit of an extension in NHRP packet is set, the NHRP packet cannot be processed unless the extension has been processed. The responder MUST return an NHRP Error Indication of type Unrecognized Extension if it is incapable

of

processing the extension. However, if a transit NHS (one which is not going to generate a reply) detects an unrecognized extension, it SHALL ignore the extension.

2 - Subnetwork ID Mismatch

This error occurs when the current station receives an NHRP packet whose NBMA subnetwork identifier matches none of the locally known identifiers for the NBMA subnetwork on which the packet is received.

3 - NHRP Loop Detected

A Loop Detected error is generated when it is determined that an NHRP packet is being forwarded in a loop.

8 - NHRP SDU Size Exceeded

size
If the SDU size of the NHRP packet exceeds the maximum SDU of the NBMA network, this error is returned.

9 - Invalid Extension

inappropriate
If an NHS finds an extension in a packet which is inappropriate for the packet type, an error is sent back to the sender with Invalid Extension as the code.

10- Invalid Next Hop Resolution Reply Received

Hop
If a client receives a Next Hop Resolution Reply for a Next Hop Resolution Request which it believes it did not make then an error packet is sent to the station making the reply with an error code of Invalid Reply Received.

Error Offset

The offset in octets into the NHRP packet, starting at the NHRP Fixed Header, at which the error was detected.

Source NBMA Address

The Source NBMA address field is the address of the station which observed the error.

Source NBMA SubAddress

The Source NBMA subaddress field is the address of the station which observed the error. If the field's length as specified in ar\$stl is 0 then no storage is allocated for this address at all.

Source Protocol Address

This is the protocol address of the station which issued the Error packet.

Destination Protocol Address

This is the protocol address of the station which sent the packet which was found to be in error.

An NHRP Error Indication packet SHALL NEVER be generated in response to another NHRP Error Indication packet. When an NHRP Error Indication packet is generated, the offending NHRP packet SHALL be discarded. In no case should more than one NHRP Error Indication packet be generated for a single NHRP packet.

If an NHS sees its own Protocol and NBMA Addresses in the Source NBMA and Source Protocol address fields of a transiting NHRP Error Indication packet then the NHS will quietly drop the packet and do nothing (this scenario would occur when the NHRP Error Indication packet was itself in a loop).


Note that no extensions may be added to an NHRP Error Indication.

5.3 Extensions Part

In the following, unless otherwise stated explicitly, the term "request" refers generically to any of the NHRP packet types which are "requests". Also, unless otherwise stated explicitly, the term "reply" refers generically to any of the NHRP packet types which are "replies".

The Extensions Part, if present, carries one or more extensions in {Type, Length, Value} triplets. Extensions are only present in a "reply" if they were present in the corresponding "request"; therefore, minimal NHRP station implementations that do not act as an NHS and do not transmit extensions need not be able to receive extensions. An implementation that is incapable of processing extensions SHALL return an NHRP Error Indication of type Unrecognized Extension when it receives an NHRP packet containing extensions.

Extensions have the following format:



C "Compulsory." If clear, and the NHS does not recognize the type code, the extension may safely be ignored. If set, and the NHS does not recognize the type code, the NHRP "request" is considered to be in error. (See below for details.)

u
Unused and must be set to zero.

Type
The extension type code (see below). The extension type is not

qualified by the Compulsory bit, but is orthogonal to it.

Length

The length in octets of the value (not including the Type and Length fields; a null extension will have only an extension header and a length of zero).

When extensions exist, the extensions list is terminated by the Null TLV, having Type = 0 and Length = 0.

Extensions may occur in any order, but any particular extension type (except for the vendor-private extension) may occur only once in an NHRP packet. The vendor-private extension may occur multiple times in a packet in order to allow for extensions which do not share the same vendor ID to be represented.

The Compulsory bit provides for a means to add to the extension set. If the bit is set, the NHRP message cannot be properly processed by the station responding to the message (e.g., the station that would issue a Next Hop Resolution Reply in response to a Next Hop Resolution Request) without processing the extension. As a result, the responder MUST return an NHRP Error Indication of type Unrecognized Extension. If the Compulsory bit is clear then the extension can be safely ignored; however, if an ignored extension is in a "request" then it MUST be returned, unchanged, in the corresponding "reply" packet type.

If a transit NHS (one which is not going to generate a "reply") detects an unrecognized extension, it SHALL ignore the extension.

If

the Compulsory bit is set, the transit NHS MUST NOT cache the information contained in the packet and MUST NOT identify itself as an egress router (in the Forward Record or Reverse Record extensions). Effectively, this means, if a transit NHS encounters

an

extension which it cannot process and which has the Compulsory bit set then that NHS MUST NOT participate in any way in the protocol exchange other than acting as a forwarding agent.

5.3.0 The End Of Extensions


Compulsory = 1
Type = 0
Length = 0

When extensions exist, the extensions list is terminated by the End Of Extensions/Null TLV.

5.3.1 Destination Prefix Length

Compulsory = 0
Type = 1
Length = 1

This extension is used to indicate that the information carried in an NHRP packet pertains to an equivalence class of internetwork layer addresses rather than just a single internetwork layer address specified. All internetwork layer addresses that match the first "Prefix Length" bit positions for the specific internetwork layer address are included in the equivalence class.



In the case of an Next Hop Resolution Request, if equivalence information is desired from the Next Hop Resolution Reply then the Destination Prefix Length extension is included in the Next Hop Resolution Request and the Prefix Length value is coded as 0xffff. For the Next Hop Resolution Reply, the Prefix Length is set to the length of the prefix of the Next Hop Protocol Address present in the mandatory part of the packet.

In the case of an NHRP Registration Request, if equivalence information is desired to be registered then the Destination Prefix Length extension is included in the NHRP Registration Request with the Prefix Length value set to the length of the prefix of the equivalence information for the Source Protocol Address. In Next Hop Registration Reply, the Destination Prefix Length extension is merely copied unchanged.

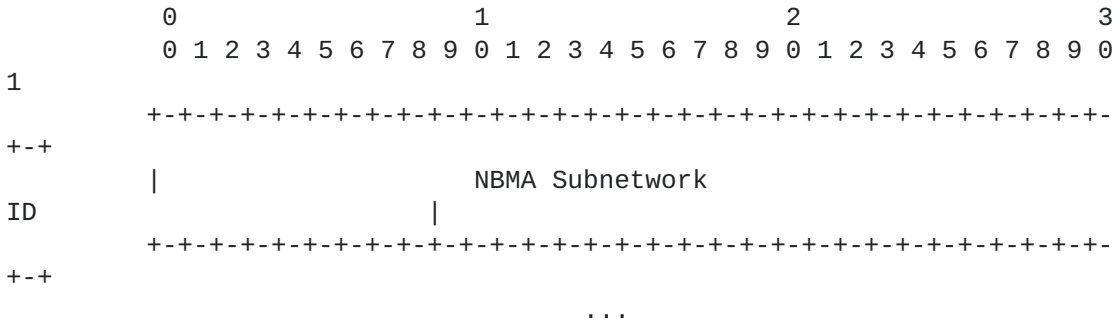
In the case of an Next Hop Purge Request, if equivalence information is desired then the Prefix Length value is set to the length of the prefix of the Target Protocol Address which represents the equivalence information to be purged. In Next Hop Purge Reply, the Destination Prefix Length extension is merely copied unchanged.

5.3.2 NBMA Subnetwork ID Extension

Compulsory = 1
Type = 2
Length = variable

This extension is used to carry one or more identifiers for the NBMA subnetwork. This can be used as a validity check to ensure that an NHRP packet does not leave a particular NBMA subnetwork. The extension is placed in a "request" packet with an ID value of zero. The first NHS along the routed path fills in the field with the identifier(s) for the NBMA subnetwork.

Multiple NBMA Subnetwork IDs may be used as a transition mechanism while NBMA Subnetworks are being split or merged.



Each identifier consists of a 32 bit globally unique value assigned to the NBMA subnetwork. This value may be chosen from the internetworking layer address space administered by the operators of the NBMA subnetwork if such an address can fit into a 32 bit field. This value is used for identification only, not for routing or any other purpose.

Each NHS processing a "request" or "reply" SHALL verify these values.

If the value is not zero and none of the values matches the NHS's NBMA Subnetwork ID, the NHS SHALL return an NHRP Error Indication to the entity identified in Source Protocol Address if the packet type is a "request" and to the Destination Protocol Address if the packet type is a "reply". The error indicated in this case is "Subnetwork ID Mismatch". The packet is discarded by the station sending the NHRP Error Indication.

When an NHS is building a "reply" and the NBMA Subnetwork ID extension is present in the correspond "request" then the NBMA Subnetwork ID extension SHALL be copied from the "request" to the "reply".

5.3.3 Responder Address Extension

- Compulsory = 1
- Type = 3
- Length = 4

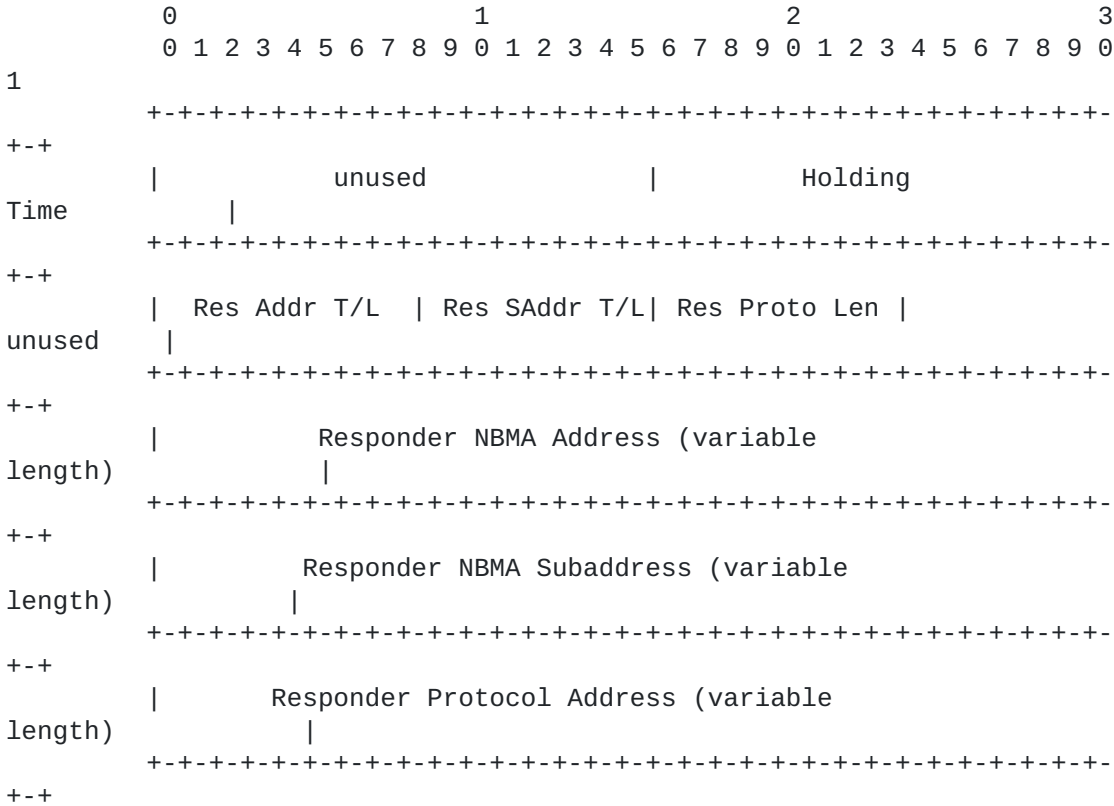
This extension is used to determine the address of the NHRP responder; i.e., the entity that generates the appropriate "reply"

packet for a given "request" packet. In the case of an Next Hop

Katz, Piscitello, Cole, Luciani
35]

[Page

Resolution Request, the station responding may be different (in the case of cached replies) than the system identified in the Next Hop field of the Next Hop Resolution Reply. Further, this extension may aid in detecting loops in the NHRP forwarding path.



Holding Time

The Holding Time field specifies the number of seconds for which the NBMA information is considered to be valid. Cached information SHALL be discarded when the holding time expires.

Res Addr T/L

Type & length of the responder NHS's NBMA address interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for ATM). When the address length is specified as 0 no storage is allocated for the address.

Res SAddr T/L

Type & length of responder NHS's NBMA subaddress interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0015 for ATM makes the address an E.164 and the subaddress an ATM Forum NSAP address). When an NBMA technology has no concept of a subaddress, the subaddress is always null with a length of 0. When the address length is specified as 0 no storage

is allocated for the address.

Res Proto Len

This field holds the length in octets of responding NHS's Protocol Address.

Responder NBMA Address

This is the NBMA address of the responding NHS.

Responder NBMA SubAddress

This is the NBMA subaddress of the responding NHS.

Responder Protocol Address

This is the Protocol Address of responding NHS.

If a "requester" desires this information, the "requester" SHALL include this extension with a value of zero. Note that this implies that no storage is allocated for the Holding Time and Type/Length fields until the "Value" portion of the extension is filled out.

If an NHS is generating a "reply" packet in response to a "request" containing this extension, the NHS SHALL include this extension, containing its protocol address in the "reply". If an NHS has more than one protocol address, it SHALL use the same protocol address consistently in all of the Responder Address, Forward NHS Record, and Reverse NHS Record extensions. The choice of which of several protocol address to include in this extension is a local matter.

If an NHRP Next Hop Resolution Reply packet being forwarded by an NHS contains a protocol address of that NHS in the Responder Address Extension then that NHS SHALL generate an NHRP Error Indication of type "NHRP Loop Detected" and discard the Next Hop Resolution Reply.

If an NHRP Next Hop Resolution Reply packet is being returned by an intermediate NHS based on cached data, it SHALL place its own address in this extension (differentiating it from the address in the Next Hop field).

5.3.4 NHRP Forward Transit NHS Record Extension

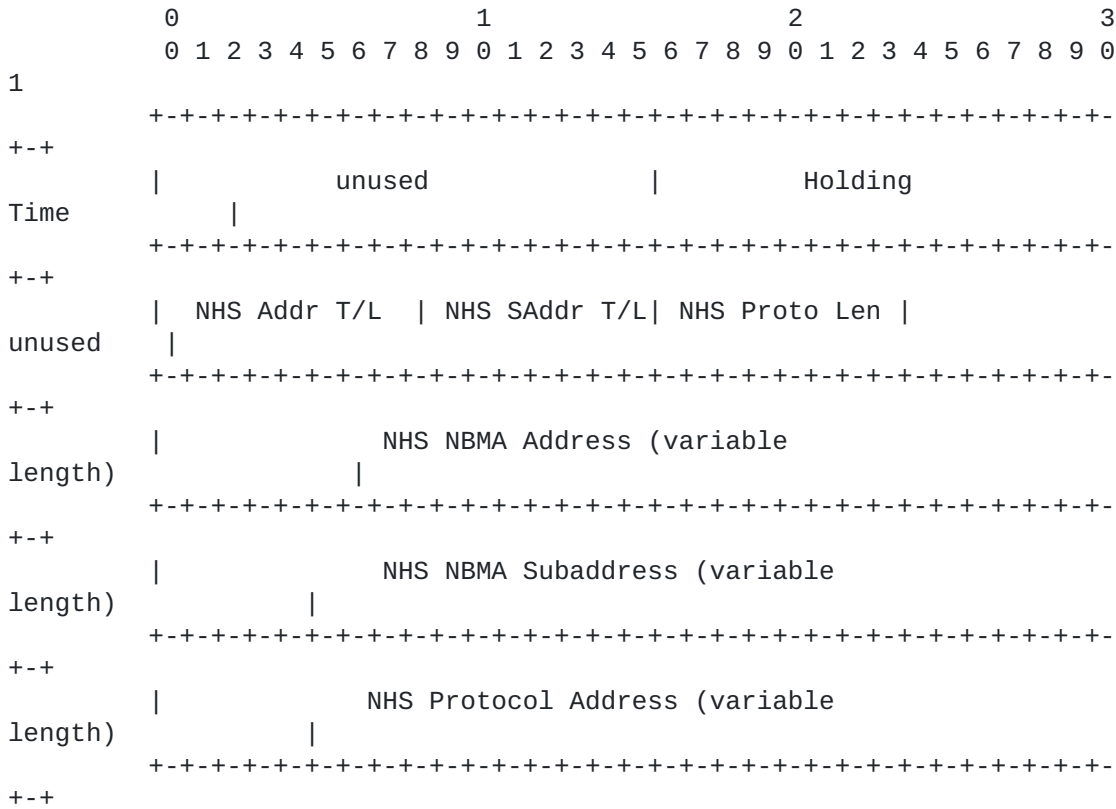
Compulsory = 1
Type = 4
Length = variable

The NHRP Forward Transit NHS record contains a list of transit NHSs through which a "request" has traversed. Each NHS SHALL append to the extension a Forward Transit NHS element (as specified below) containing its Protocol address The extension length field and the ar\$checksum fields SHALL be adjusted appropriately.

The responding NHS, as described in [Section 5.3.3](#), SHALL NOT update this extension.

In addition, NHSs that are willing to act as egress routers for packets from the source to the destination SHALL include information about their NBMA Address.

The Forward Transit NHS element has the following form:



Holding Time

The Holding Time field specifies the number of seconds for which the NBMA information is considered to be valid. Cached information SHALL be discarded when the holding time expires.

NHS Addr T/L

Type & length of the transit NHS's NBMA address interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for ATM). When the address length is specified as 0 no storage is allocated for the address.

NHS SAddr T/L

Type & length of the transit NHS's NBMA subaddress interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0015 for ATM makes the address an E.164 and the subaddress an ATM Forum NSAP address). When an NBMA technology has no concept of a subaddress the subaddress is always null with a length of 0. When the address length is specified as 0 no storage is allocated for the address.

NHS Proto Len

This field holds the length in octets of the transit NHS's Protocol Address.

NHS NBMA Address

This is the NBMA address of the transit NHS.

NHS NBMA SubAddress

This is the NBMA subaddress of the transit NHS.

NHS Protocol Address

This is the Protocol Address of the transit NHS.

If a "requester" wishes to obtain this information, it SHALL include

this extension with a length of zero. Note that this implies that no storage is allocated for the Holding Time and Type/Length fields until the "Value" portion of the extension is filled out.

If an NHS has more than one Protocol address, it SHALL use the same Protocol address consistently in all of the Responder Address, Forward NHS Record, and Reverse NHS Record extensions. The choice of which of several Protocol addresses to include in this extension is a local matter.

If a "request" that is being forwarded by an NHS contains the Protocol Address of that NHS in one of the Forward Transit NHS elements then the NHS SHALL generate an NHRP Error Indication of type "NHRP Loop Detected" and discard the "request".

5.3.5 NHRP Reverse Transit NHS Record Extension


Compulsory = 1
Type = 5
Length = variable

The NHRP Reverse Transit NHS record contains a list of transit NHSs through which a "reply" has traversed. Each NHS SHALL append a Reverse Transit NHS element (as specified below) containing its Protocol address to this extension. The extension length field and ar\$chksum SHALL be adjusted appropriately.

The responding NHS, as described in [Section 5.3.3](#), SHALL NOT update this extension.

In addition, NHSs that are willing to act as egress routers for packets from the source to the destination SHALL include information about their NBMA Address.

The Reverse Transit NHS element has the following form:



Holding Time

The Holding Time field specifies the number of seconds for which the NBMA information is considered to be valid. Cached information SHALL be discarded when the holding time expires.

NHS Addr T/L

Type & length of the responding NHS's NBMA address interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for ATM). When the address length is specified as 0 no storage is allocated for the address.

NHS SAddr T/L

Type & length of the responding NHS's NBMA subaddress interpreted in the context of the 'address family number'[6] indicated by ar\$afn (e.g., ar\$afn=0x0015 for ATM makes the address an E.164 and the subaddress an ATM Forum NSAP address). When an NBMA technology has no concept of a subaddress the subaddress is always null with a length of 0. When the address length is specified as 0 no storage is allocated for the address.

NHS Proto Len

This field holds the length in octets of the transit NHS's Protocol

Address.

NHS NBMA Address

This is the NBMA address of the transit NHS.

NHS NBMA SubAddress

This is the NBMA subaddress of the transit NHS.

NHS Protocol Address

This is the Protocol Address of the transit NHS.

If a "requester" wishes to obtain this information, it SHALL include

this extension with a length of zero. Note that this implies that no storage is allocated for the Holding Time and Type/Length fields until the "Value" portion of the extension is filled out.

If an NHS has more than one Protocol address, it SHALL use the same Protocol address consistently in all of the Responder Address, Forward NHS Record, and Reverse NHS Record extensions. The choice of which of several Protocol addresses to include in this extension is a local matter.

If a "reply" that is being forwarded by an NHS contains the Protocol Address of that NHS in one of the Reverse Transit NHS elements then the NHS SHALL generate an NHRP Error Indication of type "NHRP Loop Detected" and discard the "reply".

Note that this information may be cached at intermediate NHSs; if so, the cached value SHALL be used when generating a reply.

5.3.6 NHRP QoS Extension

Compulsory = 0
Type = 6
Length = variable

The NHRP QoS Extension is carried in Next Hop Resolution Request packets to indicate the desired QoS of the path to the indicated destination. This information may be used to help select the appropriate NBMA Next Hop.

It may also be carried in NHRP Register Request packets to indicate the QoS to which the registration applies.

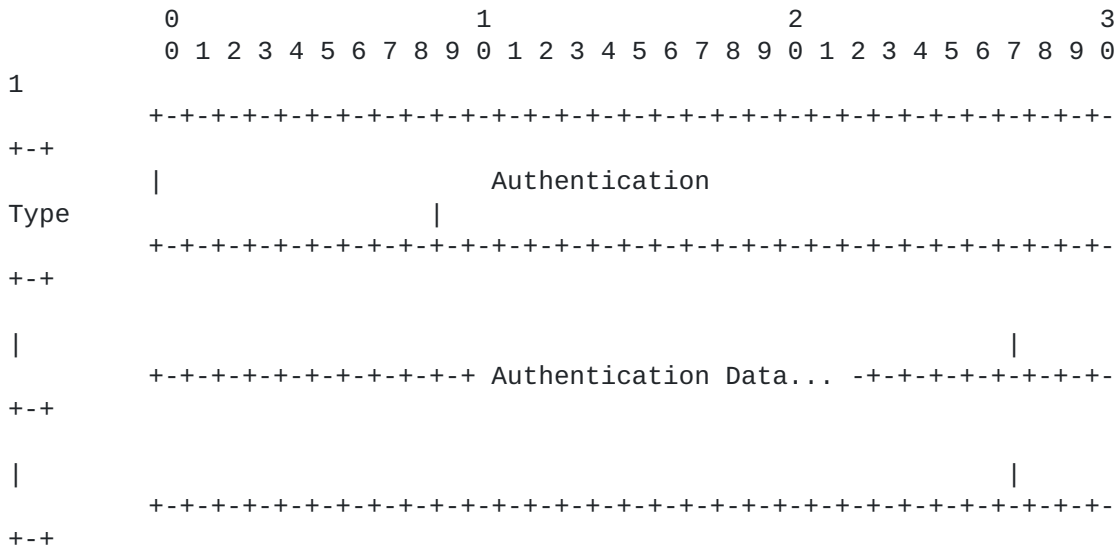
The syntax and semantics of this extension are TBD; alignment with resource reservation may be useful.

5.3.7 NHRP Authentication Extension

Compulsory = 1
Type = 7
Length = variable

The NHRP Authentication Extension is carried in NHRP packets to convey authentication information between NHRP speakers. The Authentication Extension may be included in any NHRP "request" or "reply".

Except in the case of an NHRP Registration Request/Reply Authentication is done pairwise on an NHRP hop-by-hop basis; i.e., the authentication extension is regenerated at each hop. In the case of an NHRP Registration Request/Reply, the Authentication is checked on an end-to-end basis rather than hop-by-hop. If a received packet fails the authentication test, the station SHALL generate an Error Indication of type "Authentication Failure" and discard the packet. Note that one possible authentication failure is the lack of an Authentication Extension; the presence or absence of the Authentication Extension is a local matter.



The Authentication Type field identifies the authentication method in use. Currently assigned values are:

- 1 - Cleartext Password
- 2 - Keyed MD5

All other values are reserved.

The Authentication Data field contains the type-specific authentication information.

In the case of Cleartext Password Authentication, the Authentication Data consists of a variable length password.

In the case of Keyed MD5 Authentication, the Authentication Data contains the 16 byte MD5 digest of the entire NHRP packet, including the encapsulated protocol's header, with the authentication key appended to the end of the packet. The authentication key is not transmitted with the packet.

Distribution of authentication keys is outside the scope of this document.

5.3.8 NHRP Vendor-Private Extension

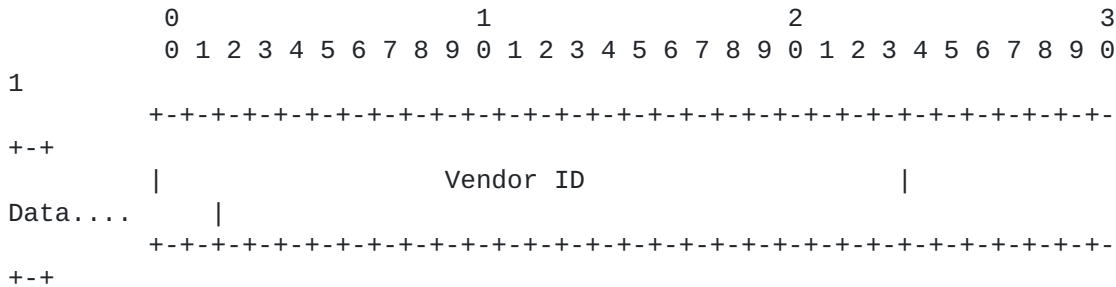
Compulsory = 0

Katz, Piscitello, Cole, Luciani
42]

[Page

Type = 8
Length = variable

The NHRP Vendor-Private Extension is carried in NHRP packets to convey vendor-private information or NHRP extensions between NHRP speakers.



Vendor ID
802 Vendor ID as assigned by the IEEE [6]

Data
The remaining octets after the Vendor ID in the payload are vendor-dependent data.

This extension may be added to any "request" or "reply" packet and it is the only extension that may be included multiple times. If the receiver does not handle this extension, or does not match the Vendor ID in the extension then the extension may be completely ignored by the receiver. If a Vendor Private Extension is included in a "request" then it must be copied in the corresponding "reply".

5.3.9 Additional Next Hop Entries Extension

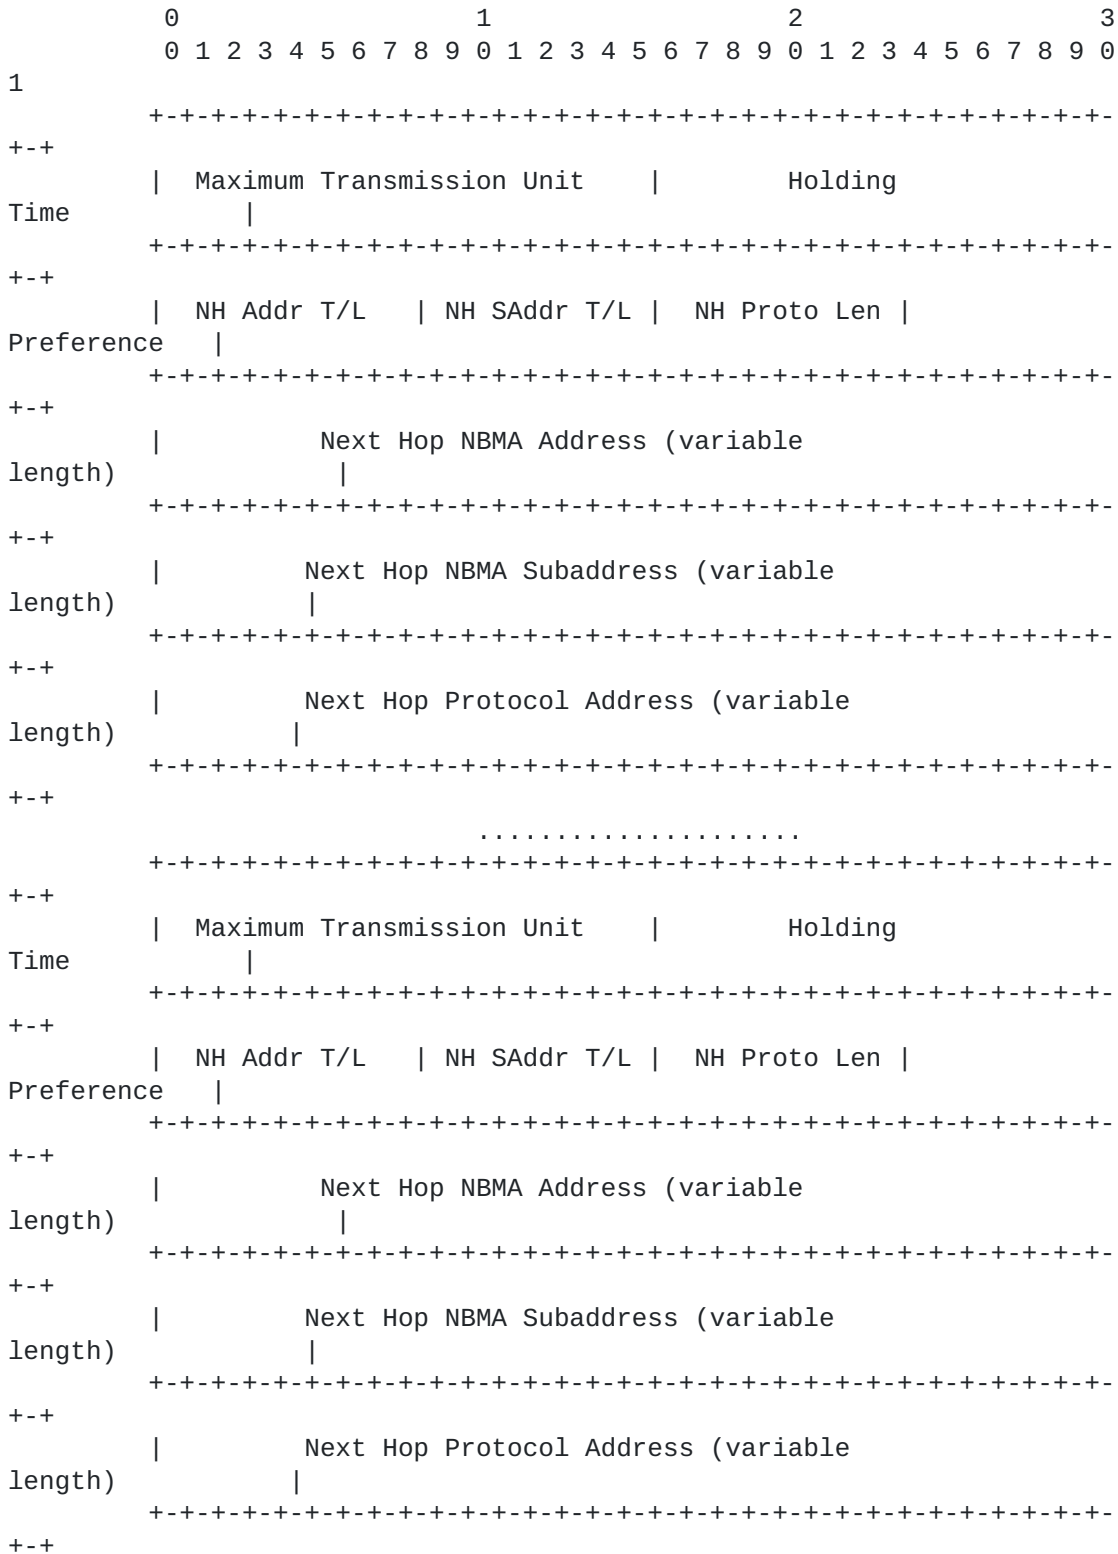
Compulsory = 0
Type = 9
Length = variable

This extension may be used to return multiple Next Hop entries in a single NHRP Reply packet. This extension MUST only be used for positive replies. The preference values are used to specify the relative preference of the entries contained in the extension. The same next Hop Protocol address may be associated with multiple NBMA addresses. Load-splitting may be performed over the addresses, given equal preference values, and the alternative addresses may be used in case of connectivity failure in the NBMA subnetwork (such as a failed call attempt in connection-oriented NBMA subnetworks).

The following shows the format for additional Next Hop Entries:

Katz, Piscitello, Cole, Luciani
43]

[Page



An NHS is not allowed to reply to an NHRP request for authoritative information with cached information, but may do so for an NHRP Request which indicates a request for non-authoritative information.

An NHS may reply to an NHRP request for non-authoritative information with authoritative information.

Maximum Transmission Unit

This field gives the maximum transmission unit for the target station. If this value is 0 then either the default MTU is used or the MTU negotiated via signaling is used if such negotiation is possible for the given NBMA.

Holding Time

The Holding Time field specifies the number of seconds for which the Next Hop NBMA information specified in the Next Hop Entry is considered to be valid. Cached information SHALL be discarded when the holding time expires. This field must be set to 0 on a NAK.

NH Addr T/L

Type & length of next hop NBMA address specified in the Next Hop Entry. This field is interpreted in the context of the 'address family number' [6] indicated by ar\$afn (e.g., ar\$afn=0x0003 for

ATM).

NH SAddr T/L

Hop Type & length of next hop NBMA subaddress specified in the Next Entry. This field is interpreted in the context of the 'address family number' [6] indicated by ar\$afn (e.g., ar\$afn=0x0015 for ATM makes the address an E.164 and the subaddress an ATM Forum NSAP address). When an NBMA technology has no concept of a subaddress the subaddress is always null with a length of 0. When the address length is specified as 0 no storage is allocated for the address.

NH Proto Len

This field holds the length in octets of the Next Hop Protocol Address specified in the Next Hop Entry.

Preference

This field specifies the preference of the specific Next Hop Entry relative to other Next Hop entries in this Next Hop Resolution Reply mandatory part or in the Additional Next Hop Entries Extension for the given internetworking protocol. Higher values indicate higher preference. Action taken when multiple next hop entries have the highest preference value is a local matter.

Next Hop NBMA Address

This is the NBMA address of the station that is the next hop for packets bound for the internetworking layer address specified.

Next Hop NBMA SubAddress

for This is the NBMA subaddress of the station that is the next hop packets bound for the internetworking layer address specified.

Next Hop Protocol Address

This internetworking layer address specifies the next hop. This will be the address of the destination host if it is directly attached to the NBMA subnetwork, or the egress router if it is not directly attached.

6. Protocol Operation

In this section, we discuss certain operational considerations of NHRP.

6.1 Router-to-Router Operation

In practice, the initiating and responding stations may be either hosts or routers. However, there is a possibility under certain conditions that a stable routing loop may occur if NHRP is used

between two routers. In particular, attempting to establish an NHRP path across a boundary where information used in route selection is lost may result in a routing loop. Such situations include the loss of BGP path vector information, the interworking of multiple routing protocols with dissimilar metrics (e.g, RIP and OSPF), etc. In such circumstances, NHRP should not be used. This situation can be avoided if there are no "back door" paths between the entry and egress router outside of the NBMA subnetwork. Protocol mechanisms to relax these restrictions are under investigation.

In general it is preferable to use mechanisms, if they exist, in routing protocols to resolve the egress point when the destination lies outside of the NBMA subnetwork, since such mechanisms will be more tightly coupled to the state of the routing system and will probably be less likely to create loops.

6.2 Cache Management Issues

The management of NHRP caches in the source station, the NHS serving the destination, and any intermediate NHSs is dependent on a number of factors.

6.2.1 Caching Requirements

Source Stations

Source stations MUST cache all received Next Hop Resolution Replies that they are actively using. They also must cache "incomplete" entries, i.e., those for which a Next Hop Resolution Request has been sent but which a Next Hop Resolution Reply has not been received. This is necessary in order to preserve the Request ID for retries, and provides the state necessary to avoid triggering Next Hop Resolution Requests for every data packet sent to the destination.

Source stations MUST purge expired information from their caches. Source stations MUST purge the appropriate cached information upon receipt of an NHRP Purge Request packet.

When a station has a co-resident client and NHS, the station may reply to Next Hop Resolution Requests with information which the station cached as a result of the station making its own Next Hop Resolution Requests and receiving its own Next Hop Resolution Replies as long as the station follows the rules for Transit NHSs as seen below.

Serving NHSs

The NHS serving the destination (the one which responds authoritatively to Next Hop Resolution Requests) SHOULD cache information about all Next Hop Resolution Requests to which it has responded if the information in the Next Hop Resolution Reply has the possibility of changing during its lifetime (so that an NHRP Purge Request packet can be sent). The NBMA information provided by the source station in the Next Hop Resolution Request may be cached for the duration of its holding time. This information is considered to be stable, since it identifies a station directly attached to the NBMA subnetwork. An example of unstable information is NBMA information derived from a routing table, where that routing table information has not been guaranteed to be stable through administrative means.

Transit NHSs

A Transit NHS (lying along the NHRP path between the source station and the responding NHS) may cache information contained in Next Hop Resolution Request packets that it forwards. A Transit NHS may cache information contained in Next Hop Resolution Reply packets that it forwards only if that Next Hop Resolution Reply has the Stable (B) bit set. It MUST discard any cached information whose holding time has expired. It may return cached information in response to non-authoritative Next Hop Resolution Requests only.

6.2.2 Dynamics of Cached Information

NBMA-Connected Destinations

NHRP's most basic function is that of simple NBMA address resolution of stations directly attached to the NBMA subnetwork. These mappings are typically very static, and appropriately chosen holding times will minimize problems in the event that the NBMA address of a station must be changed. Stale information will cause a loss of connectivity, which may be used to trigger an authoritative Next Hop Resolution Request and bypass the old data. In the worst case, connectivity will fail until the cache entry times out.

This applies equally to information marked in Next Hop Resolution Replies as being "stable" (via the "B" bit).

This also applies equally well to source stations that are routers as well as those which are hosts.

Note that the information carried in the Next Hop Resolution Request packet is always considered "stable" because it represents a station that is directly connected to the NBMA subnetwork.

Destinations Off of the NBMA Subnetwork

If the source of a Next Hop Resolution Request is a host and the destination is not directly attached to the NBMA subnetwork, and the route to that destination is not considered to be "stable," the destination mapping may be very dynamic (except in the case of a subnetwork where each destination is only singly homed to the NBMA subnetwork). As such the cached information may very likely become stale. The consequence of stale information in this case will be a suboptimal path (unless the internetwork has partitioned or some other routing failure has occurred).

Strategies for maintaining NHRP cache information in the presence of dynamic routing changes will be discussed in a separate document.

6.3 Use of the Destination Prefix Length Extension

A certain amount of care needs to be taken when using the Destination Prefix Length Extension, in particular with regard to the prefix length advertised (and thus the size of the equivalence class specified by it). Assuming that the routers on the NBMA subnetwork are exchanging routing information, it should not be possible for an NHS to create a black hole by advertising too large of a set of destinations, but suboptimal routing (e.g., extra internetwork layer hops through the NBMA) can result. To avoid this situation an NHS that wants to send the Destination Prefix Length Extension MUST obey the following rule:

The NHS examines the Network Layer Reachability Information (NLRI) associated with the route that the NHS would use to forward towards the destination (as specified by the Destination internetwork layer address in the Next Hop Resolution Request), and extracts from this NLRI the shortest address prefix such that: (a) the Destination internetwork layer address (from the Next Hop Resolution Request) is covered by the prefix, (b) the NHS does not have any routes with NLRI that forms a subset of what is covered by the prefix. The prefix may then be used for the Destination Prefix Length Extension.

The NHRP Destination Prefix Length Extension should be used with restraint, in order to avoid NHRP stations choosing suboptimal transit paths when overlapping prefixes are available. This

extension SHOULD only be used in a Next Hop Resolution Reply when either:

- (a) All destinations covered by the prefix are on the NBMA network, or
- (b) All destinations covered by the prefix are directly attached to the NHRP responding station.

For other cases, there may be no single optimal transit path for destinations encompassed by the address prefix, and an NHRP station may fail to choose the optimal transit path simply because it is not aware of all such paths. So for cases not covered by (a) and (b), an Next Hop Resolution Reply packet should not include the NHRP Destination Prefix Length Extension.

6.4 Domino Effect

One could easily imagine a situation where a router, acting as an ingress station to the NBMA subnetwork, receives a data packet, such that this packet triggers an Next Hop Resolution Request. If the router forwards this data packet without waiting for an NHRP transit path to be established, then when the next router along the path receives the packet, the next router may do exactly the same - originate its own Next Hop Resolution Request (as well as forward the packet). In fact such a data packet may trigger Next Hop Resolution Request generation at every router along the path through an NBMA subnetwork. We refer to this phenomena as the NHRP "domino" effect.

The NHRP domino effect is clearly undesirable. At best it may result in excessive NHRP traffic. At worst it may result in an excessive number of virtual circuits being established unnecessarily. Therefore, it is important to take certain measures to avoid or suppress this behavior. NHRP implementations for NHSSs MUST provide a mechanism to address this problem. It is recommended that implementations provide one or more of the following solutions.

Possibly the most straightforward solution for suppressing the domino effect would be to require transit routers to be preconfigured not to originate Next Hop Resolution Requests for data traffic which is simply being forwarded (not originated). In this case the routers avoid the domino effect through an administrative policy.

A second possible solution would be to require that when a router forwards an Next Hop Resolution Request, the router instantiates a (short-lived) state. This state consists of the route that was used

to forward the Next Hop Resolution Request. If the router receives
a data packet, and the packet triggers an Next Hop Resolution Request
generation by the router, the router checks whether the route to
forward the Next Hop Resolution Request was recently used to forward

some other Next Hop Resolution Request. If so, then the router suppresses generation of the new Next Hop Resolution Request (but still forwards the data packet). This solution also requires that when a station attempts to originate an Next Hop Resolution Request the station should send the Next Hop Resolution Request before the data packet that triggered the origination of the Next Hop

Resolution

Request. Otherwise, unnecessary Next Hop Resolution Requests may still be generated.

A third possible strategy would be to configure a router in such a way that Next Hop Resolution Request generation by the router would be driven only by the traffic the router receives over its non-NBMA interfaces (interfaces that are not attached to an NBMA subnetwork). Traffic received by the router over its NBMA-attached interfaces would not trigger NHRP Next Hop Resolution Requests. Just as in the first case, such a router avoids the NHRP domino effect through administrative means.

Lastly, rate limiting of Next Hop Resolution Requests may help to avoid the NHRP domino effect. Intermediate routers which would otherwise generate unnecessary Next Hop Resolution Requests may instead suppress such Next Hop Resolution Requests due to the measured Next Hop Resolution Request rate exceeding a certain threshold. Of course, such rate limiting would have to be very aggressive in order to completely avoid the domino effect. Further work is needed to analyze this solution.

7. Security Considerations

As in any routing protocol, there are a number of potential security attacks possible. Plausible examples include denial-of-service attacks, and masquerade attacks using register and purge packets.

The use of authentication on all packets is recommended to avoid such attacks.

The authentication schemes described in this document are intended to allow the receiver of a packet to validate the identity of the sender; they do not provide privacy or protection against replay attacks.

Detailed security analysis of this protocol is for further study.

8. Discussion

The result of an Next Hop Resolution Request depends on how routing is configured among the NHSs of an NBMA subnetwork. If the destination station is directly connected to the NBMA subnetwork and

the the routed path to it lies entirely within the NBMA subnetwork, the Next Hop Resolution Replies always return the NBMA address of the destination station itself rather than the NBMA address of some egress router. On the other hand, if the routed path exits the NBMA subnetwork, NHRP will be unable to resolve the NBMA address of the destination, but rather will return the address of the egress router.

For destinations outside the NBMA subnetwork, egress routers and routers in the other subnetworks should exchange routing information so that the optimal egress router may be found.

In addition to NHSs, an NBMA station could also be associated with one or more regular routers that could act as "connectionless servers" for the station. The station could then choose to resolve the NBMA next hop or just send the packets to one of its connectionless servers. The latter option may be desirable if communication with the destination is short-lived and/or doesn't require much network resources. The connectionless servers could, of course, be physically integrated in the NHSs by augmenting them with internetwork layer switching functionality.

References

- [1] NBMA Address Resolution Protocol (NARP), Juha Heinanen and Ramesh Govindan, [draft-ietf-rolc-nbma-arp-00.txt](#).
- [2] Address Resolution Protocol, David C. Plummer, [RFC 826](#).
- [3] Classical IP and ARP over ATM, Mark Laubach, [RFC 1577](#).
- [4] Transmission of IP datagrams over the SMDS service, J. Lawrence and D. Piscitello, [RFC 1209](#).
- [5] Protocol Identification in the Network Layer, ISO/IEC TR 9577:1990.
- [6] Assigned Numbers, J. Reynolds and J. Postel, [RFC 1700](#).
- [7] Multiprotocol Encapsulation over ATM Adaptation Layer 5, J. Heinanen, [RFC1483](#).
- [8] Multiprotocol Interconnect on X.25 and ISDN in the Packet Mode, A. Malis, D. Robinson, and R. Ullmann, [RFC1356](#).
- [9] Multiprotocol Interconnect over Frame Relay, T. Bradley, C. Brown, and A. Malis, [RFC1490](#).

Acknowledgments

We would like to thank Juha Heinenan of Telecom Finland and Ramesh Govidan of ISI for their work on NBMA ARP and the original NHRP draft, which served as the basis for this work. John Burnett of Adaptive, Dennis Ferguson of ANS, Joel Halpern of Newbridge, Paul Francis of NTT, Tony Li and Yakov Rekhter of cisco, and Grenville Armitage of Bellcore should also be acknowledged for comments and suggestions that improved this work substantially. We would also like to thank the members of the Routing Over Large Clouds working group of the IETF, whose review and discussion of this document have been invaluable.

Authors' Addresses

Dave Katz
cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134 USA

Phone: +1 408 526 8284
Email: dkatz@cisco.com

Bruce Cole
cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134 USA

Phone: +1 408 526 4000
Email: bcole@cisco.com

David Piscitello
Core Competence
1620 Tuckerstown Road
Dresher, PA 19025 USA

Phone: +1 215 830 0692
Email: dave@corecom.com

James V. Luciani
Ascom Nexion
289 Great Road
Acton, MA 01720-4379 USA

Phone: +1 508 266 3450
Email: luciani@nexen.com

