

NHRP Protocol Applicability Statement

Status of this Memo

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

Copyright Notice

Copyright (C) The Internet Society (1998). All Rights Reserved.

Abstract

As required by the Routing Protocol Criteria [RFC 1264], this memo discusses the applicability of the Next Hop Resolution Protocol (NHRP) in routing of IP datagrams over Non-Broadcast Multiple Access (NBMA) networks, such as ATM, SMDS and X.25.

1. Protocol Documents

The NHRP protocol description is defined in [1]. The NHRP MIB description is defined in [2].

2. Introduction

This document summarizes the key features of NHRP and discusses the environments for which the protocol is well suited. For the purposes of description, NHRP can be considered a generalization of Classical IP and ARP over ATM which is defined in [3] and of the Transmission of IP Datagrams over the SMDS Service, defined in [4]. This generalization occurs in 2 distinct directions.

Firstly, NHRP avoids the need to go through extra hops of routers when the Source and Destination belong to different Logical Internet Subnets (LIS). Of course, [3] and [4] specify that when the source and destination belong to different LISs, the source station must forward data packets to a router that is a member of multiple LISs, even though the source and destination stations may be on the same logical NBMA network. If the source and destination stations belong to the same logical NBMA network, NHRP provides the source station

with an inter-LIS address resolution mechanism at the end of which both stations can exchange packets without having to use the services of intermediate routers. This feature is also referred to as "short-cut" routing. If the destination station is not part of the logical NBMA network, NHRP provides the source with the NBMA address of the current egress router towards the destination.

The second generalization is that NHRP is not specific to a particular NBMA technology. Of course, [3] assumes an ATM network and [4] assumes an SMDS network at their respective subnetwork layers.

NHRP is specified for resolving the destination NBMA addresses of IP datagrams over IP subnets within a large NBMA cloud. NHRP has been designed to be extensible to network layer protocols other than IP, possibly subject to other network layer protocol specific additions.

As an important application of NHRP, the Multiprotocol Over ATM (MPOA) Working Group of the ATM Forum has decided to adopt and to integrate NHRP into its MPOA Protocol specification [5]. As such, NHRP will be used in resolving the ATM addresses of MPOA packets destined outside the originating subnet.

3. Key Features

NHRP provides a mechanism to obtain the NBMA network address of the destination, or of a router along the path to the destination. NHRP is not a routing protocol, but may make use of routing information. This is further discussed in Section 5.

The most prominent feature of NHRP is that it avoids extra router hops in an NBMA with multiple LISs. To this goal, NHRP provides the source with the NBMA address of the destination, if the destination is directly attached to the NBMA. If the destination station is not attached to the NBMA, then NHRP provides the source with the NBMA address of an exit router that has connectivity to the destination. In general, there may be multiple exit routers that have connectivity to the destination. If NHRP uses the services of a dynamic routing algorithm in fulfilling its function, which is necessary for robust and scalable operation, then the exit router identified by NHRP reflects the selection made by the network layer dynamic routing protocol. In general, the selection made by the routing protocol would often reflect a desirable attribute, such as identifying the exit router that induces the least number of hops in the original routed path.

NHRP is defined for avoiding extra hops in the delivery of IP packets with a single destination. As such, it is not intended for direct use in a point-to-multipoint communication setting. However, elements of NHRP may be used in certain multicast scenarios for the purpose of providing short cut routing. Such an effort is discussed in [6]. In this case, NHRP would avoid intermediate routers in the multicast path. The scalability of providing short-cut paths in a multicast environment is an open issue.

NHRP can be used in host-host, host-router and router-host communications. When used in router-router communication, NHRP (as defined in [1]) can produce persistent routing loops if the underlying routing protocol loses information critical to loop suppression. This may occur when there is a change in router metrics across the autonomous system boundaries. NHRP for router-router communication that avoids persistent forwarding loops will be addressed in a separate document.

A special case of router-router communication where loops will not occur is when the destination host is directly adjacent to the non-NBMA interface of the egress router. If it is believed that the adjacency of the destination station to the egress router is a stable topological configuration, then NHRP can safely be used in this router-router communication scenario. If the NHRP Request has the Q bit set, indicating that the requesting party is a router, and if the destination station is directly adjacent to the egress router as a stable topological configuration, then the egress router can issue a corresponding NHRP reply. If the destination is not adjacent to the egress router, and if Q bit is set in the Request, then a safe mode of operation for the egress router would be to issue a negative NHRP Reply (NAK) for this particular request, thereby enforce data packets to follow the routed path.

As a result of having inter-LIS address resolution capability, NHRP allows the communicating parties to exchange packets by fully utilizing the particular features of the NBMA network. One such example is the use of QoS guarantees when the NBMA network is ATM.

Here, due to short-cut routing, ATM provided QoS guarantees can be implemented without having to deal with the issues of re-assembling and re-segmenting IP packets at each network layer hop.

NHRP protocol can be viewed as a client-server interaction. An NHRP Client is the one who issues an NHRP Request. An NHRP Server is the one who issues a reply to an NHRP request, or the one who forwards a received NHRP request to another Server. Of course, an NHRP entity may act both as a Client and a Server.

4. Use of NHRP

In general, issuing an NHRP request is an application dependent action [7]. For applications that do not have particular QoS requirements, and that are executed within a short period of time, an NBMA short-cut may not be a necessity. In situations where there is a "cost" associated with NBMA short-cuts, such applications may be better served by network layer hop-by-hop routing. Here, "cost" may be understood in a monetary context, or as additional strain on the equipment that implements short-cuts. Therefore, there is a trade-off between the "cost" of a short-cut path and its utility to the user. Reference [7] proposes that this trade-off should be addressed at the application level. In an environment consisting of LANs and routers that are interconnected via dedicated links, the basic routing decision is whether to forward a packet to a router, or to broadcast it locally. Such a decision on local vs. remote is based on the destination address. When routing IP packets over an NBMA network, where there is potentially a direct Source to Destination connectivity with QoS options, the decision on local vs. remote is no longer as fundamentally important as in the case where packets have to traverse routers that are interconnected via dedicated links. Thus, in an NBMA network with QoS options, the basic decision becomes the one of short-cut vs. hop-by-hop network layer routing. In this case, the relevant criterion becomes applications' QoS requirements [7]. NHRP is particularly applicable for environments where the decision on local vs. remote is superseded by the decision on short-cut vs. hop-by-hop network layer routing.

Let us assume that the trade-off is in favor of a short-cut NBMA route. Generally, an NHRP request can be issued by a variety of NHRP aware entities, including hosts and routers with NBMA interfaces. If an IP packet traverses multiple hops before a short-cut path has been established, then there is a chance that multiple short-cut paths could be formed. In order to avoid such an undesirable situation, a useful operation rule is to authorize only the following entities to issue an NHRP request and to perform short-cut routing.

- i) The host that originates the IP packet, if the host has an NBMA interface.
- ii) The first router along the routing path of the IP packet such that the next hop is reachable through the NBMA interface of that particular router.
- iii) A policy router within an NBMA network through which the IP packet has to traverse.

5. Protocol Scalability

As previously indicated, NHRP is defined for the delivery of IP packets with a single destination. Thus, this discussion is confined to a unicast setting. The scalability of NHRP can be analyzed at three distinct levels:

- o Client level
- o LIS level
- o Domain level

At the the Client level, the scalability of NHRP is affected by the processing and memory limitations of the NIC that provides interface to the NBMA network. When the NBMA network is connection oriented, such as ATM, NIC limitations may bound the scalability of NHRP in certain applications. For example, a server that handles hundreds of requests per second using an ATM interface may be bounded by the performance characteristics of the corresponding NIC. Similarly, when the NHRP Client resides at an NBMA interface of a router, memory and processing limitations of router's NIC may bound the scalability of NHRP. This is because routers generally deal with an aggregation of traffic from multiple sources, which in turn creates a potentially large number of SVCCs out of the router's NBMA interface.

At the LIS level, the main issue is to maintain and deliver a sizable number of NBMA to Network layer address mappings within large LISs. To this goal, NHRP implementations can use the services of the Server Cache Synchronization Protocol (SCSP) [8] that allows multiple synchronized NHSS within an LIS, and hence resolve the associated scalability issue.

At the NHRP Domain level, network layer routing is used in resolving the NBMA address of a destination outside the LIS. As such, the scalability of NHRP is closely tied to the scalability of the network layer routing protocol used by NHRP. Dynamic network layer routing protocols are proven to scale well. Thus, when used in conjunction with dynamic routing algorithms, at the NHRP domain level, NHRP should scale in the same order as the routing algorithm, subject to the assumption that all the routers along the path are NHRP aware. If an NHRP Request is processed by a router that does not implement NHRP, it will be silently discarded. Then, short-cuts cannot be implemented and connectivity will be provided on a hop-by-hop basis.

Thus, when NHRP is implemented in conjunction with dynamic network layer routing, a scaling requirement for NHRP is that virtually all the routers within a logical NBMA network should be NHRP aware.

One can also use static routing in conjunction with NHRP. Then, not all the routers in the NBMA network need to be NHRP aware. That is, since the routers that need to process NHRP control messages are specified by static routing, routers that are not included in the manually defined static paths do not have to be NHRP aware. Of course, static routing does not scale, and if the destination is off the NBMA network, then the use of static routing could result in persistently suboptimal routes. Use of static routing also has fairly negative failure modes.

6. Discussion

NHRP does not replace existing routing protocols. In general, routing protocols are used to determine the proper path from a source host or router, or intermediate router, to a particular destination. If the routing protocol indicates that the proper path is via an interface to an NBMA network, then NHRP may be used at the NBMA interface to resolve the destination IP address into the corresponding NBMA address. Of course, the use of NHRP is subject to considerations discussed in Section 4.

Assuming that NHRP is applicable and the destination address has been resolved, packets are forwarded using the particular data forwarding and path determination mechanisms of the underlying NBMA network. Here, the sequence of events are such that route determination is performed by IP routing, independent of NHRP. Then, NHRP is used to create a short-cut track upon the path determined by the IP routing protocol. Therefore, NHRP "shortens" the routed path. NHRP (as defined in [1]) is not sufficient to suppress persistent forwarding loops when used for router-router communication if the underlying routing protocol loses information critical to loop suppression [9]. Work is in progress [10] to augment NHRP to enable its use for the router-router communication without persistent forwarding loops.

When the routed path keeps changing on some relatively short time scale, such as seconds, this situation will have an effect on the operation of NHRP. In certain router-router operations, changes in the routed path could create persistent routing loops. In host-router, or router-host communications, frequent changes in routed paths could result in inefficiencies such as frequent creation of short-cut paths which are short lived.

7. Security Considerations

NHRP is an address resolution protocol, and SCSP is a database synchronization protocol. As such, they are possibly subject to server (for NHRP) or peer (for SCSP) spoofing and denial of service attacks. They both provide authentication mechanisms to allow their

use in environments in which spoofing is a concern. Details can be found in sections 5.3.4 in [1] and B.3.1 in [8]. There are no additional security constraints or concerns raised in this document that are not already discussed in the referenced sections.

References

- [1] Luciani, J., Katz, D., Piscitello, D., Cole, B., and N. Doraswamy, "NBMA Next Hop Resolution Protocol (NHRP)", RFC 2332, April 1998.
- [2] Greene, M., and J. Luciani, "NHRP Management Information Base", Work in Progress.
- [3] Laubach, M., and J. Halpern, "Classical IP and ARP over ATM", RFC 2225, April 1998.
- [4] Lawrance, J., and D. Piscitello, "The Transmission of IP datagrams over the SMDS service", RFC 1209, March 1991.
- [5] Multiprotocol Over ATM Version 1.0, ATM Forum Document af-mpoa-0087.000
- [6] Rekhter, Y., and D. Farinacci, "Support for Sparse Mode PIM over ATM", Work in Progress.
- [7] Rekhter, Y., and D. Kandlur, "Local/Remote" Forwarding Decision in Switched Data Link Subnetworks", RFC 1937, May 1996.
- [8] Luciani, J., Armitage, G., Halpern, J., and N. Doraswamy, "Server Cache Synchronization Protocol (SCSP) - NBMA", RFC 2334, April 1998.
- [9] Cole, R., Shur, D., and C. Villamizar, "IP over ATM: A Framework Document", RFC 1932, April 1996.
- [10] Rekhter, Y., "NHRP for Destinations off the NBMA Subnetwork", Work in Progress.

Acknowledgements

The author acknowledges valuable contributions and comments from many participants of the ION Working Group, in particular from Joel Halpern of Newbridge Networks, David Horton of Centre for Information Technology Research, Andy Malis of Nexion, Yakov Rekhter and George Swallow of Cisco Systems and Curtis Villamizar of ANS.

Author's Address

Derya H. Cansever
GTE Laboratories Inc.
40 Sylvan Rd. MS 51
Waltham MA 02254

Phone: +1 617 466 4086
EMail: dcansever@gte.com

Full Copyright Statement

Copyright (C) The Internet Society (1998). All Rights Reserved.

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

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

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

