

A Distributed NHRP Service Using SCSP

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

This document describes a method for distributing an NHRP service within a LIS [1]. This method uses the Server Cache Synchronization Protocol (SCSP) [2] to synchronize the client information databases held by NHRP Servers (NHSS) within a LIS.

1. Introduction

The keywords MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL, when they appear in this document, are to be interpreted as described in [4].

NHRP Clients (NHCs) register their existence and reachability information with NHRP Servers (NHSS). There may be multiple NHSS in a given Logical IP Subnet (LIS). NHCs do not necessarily register with all NHSS in a LIS; however, all NHCs need to be able to query at least one NHS about any NHC within the LIS. Thus, the contents of the NHS databases in a LIS need to be synchronized across the LIS. The Server Cache Synchronization Protocol (SCSP) solves the generalized server synchronization/cache-replication problem for distributed databases and thus SCSP may be applied to the NHS database synchronization problem within the LIS.

SCSP is defined in two parts: the protocol independent part and the client/server protocol specific part. The protocol independent part is defined in [2] whereas this document will specify the client/server protocol specific part where NHRP is the client/server protocol.

This document is separate from [2] because it was felt that it was desirable to allow the client/server protocol specific part specification for NHRP to progress independently from the protocol independent specification.

2. Overview

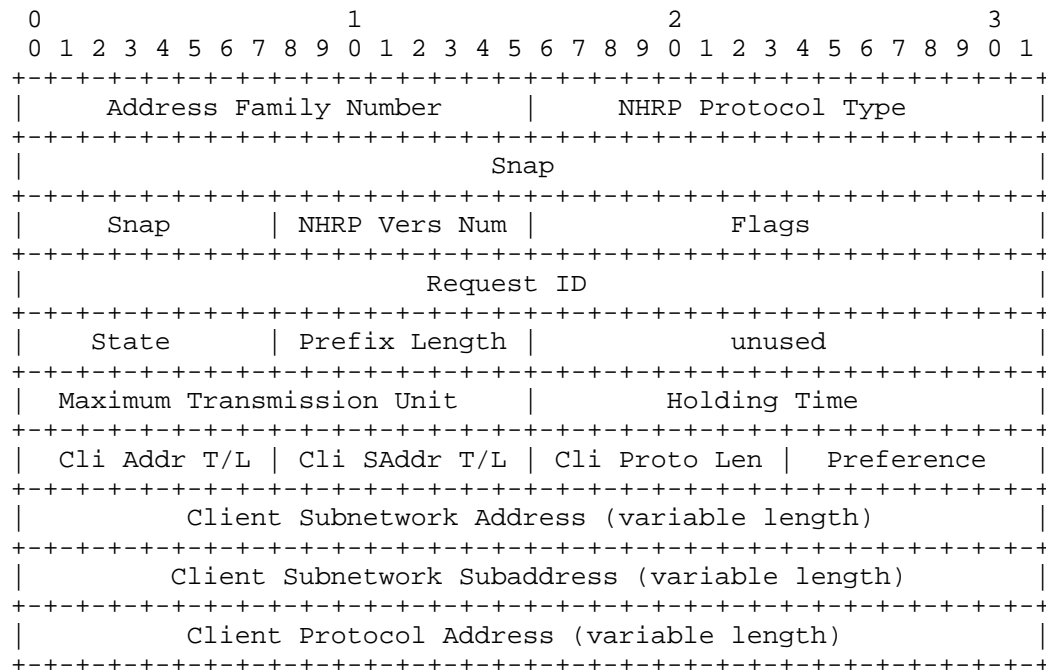
All NHSs belonging to a Logical IP Subnet (LIS) [1] are said to belong to a Server Group (SG). An SG is identified by, not surprisingly, its SGID which is contained in a field in all SCSP packets. All SCSP packets contain a Protocol ID (PID) field as well. This PID field is set to 0x0002 to signify that SCSP synchronizing NHS databases as opposed to synchronizing some other protocol's databases (see Section B.2.0.1 of [2] for more details). In general, PIDs for SCSP will be assigned by IANA as described in Section C of [2]. In the case of NHRP, the client/server protocol specific specification was initially written at the same time as SCSP, and thus a PID=0x0002 was assigned by the author.

SCSP places no topological requirements upon an NHRP SG. Obviously, however, the resultant graph of NHSs must span the set of NHSs to be synchronized. For more information about the client/server protocol independent part of SCSP, the reader is encouraged to see [2].

When a SG is using SCSP for synchronization, an NHC will register with only one NHS, but the NHC MAY use any NHS in the SG. When an NHC wishes to leave a SG, the NHC MUST do one of the following: 1) the NHC MUST send an NHRP Purge Request for itself requesting a reply, and it MUST wait for an NHRP Purge Reply, 2) the NHC MUST keep the Request ID it used when registering itself in non-volatile RAM and use a Request ID larger than the one saved when re-registering, or 3) the NHC MUST not re-register for a time equal to the Holding Time specified in the previous registration. It is necessary to do one of the previous in order to prevent the unlikely case of race conditions from occurring during updated. In the case where method 2 is used, the NHS with which the NHC registered uses its ID as the OID and the Request ID from the NHC as the CSA Sequence Number in the CSA(S) Record.

3. Format of the CSA Record NHRP Specific Part

CSA Records in SCSP contain a "Client/Server Protocol Specific Part" which contains the non-protocol independent information for a given server's cache entry.



The following six fields contain values specified in the common header of the mandatory part of an NHRP Registration Request or NHRP Purge Request packet which caused the creation/deletion/modification/update/etc. of an NHS's cache entry.

Address Family Number

Defines the type of "link layer" addresses being carried. This number is taken from the 'address family number' list specified in [3]. This field is the same field which would be supplied in an NHRP packet in the ar\$afn field.

NHRP Protocol Type

This field is the same field which would be supplied in an NHRP packet in the ar\$pro.type field.

Snap

This field is the same field which would be supplied in an NHRP packet in the ar\$pro.snap field.

NHRP Vers Num

This field indicates what version of generic address mapping and management protocol that is represented by this message. This field contains 0x01 for the NHRP protocol version 1. This field is the same field which would be supplied in an NHRP packet in the ar\$op.version field.

Flags

Defined flags are as follows:

```

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

```

U

This is the Uniqueness bit.

A

When set, this bit specifies that the cache entry was created as a result of ATMARP client interaction with the NHS.

Request ID

This field contains the Request ID value placed in the cache entry of the NHS as a result of an NHRP Registration Request. This NHS is the NHS causing a synchronization event.

State

This field contains a value which represents the new state of the client.

- 0 - Client is registered and available.
- 1 - Client reregistered.
- 2 - Client has been purged.
- 3 - No such client data in server cache

Note that a time-out of a cache entry does not cause a CSA Record to be sent because, if everything is working properly then all NHSs have the cache entry timing out at the same time. Thus, the individual NHSs would take the appropriate actions necessary.

The following ten fields contain values specified in or derived from the CIE of an NHRP Registration Request or NHRP Purge Request packet which caused the creation/deletion/modification/update/etc. of an NHS's cache entry.

Prefix Length

This field contains the internetwork layer address prefix length value covered by the cache entry being synchronized.

Maximum Transmission Unit

This field contains a value supplied by or derived from information in the CIE of the NHRP Registration Request packet.

Holding Time

The Holding Time field specifies the number of seconds remaining for which the Next Hop NBMA information specified in the CIE of the NHRP Registration Request is considered to be valid by the NHS initiating the synchronization event.

Cli Addr T/L

Type & length of next hop NBMA address (see [1]).

Cli SAddr T/L

Type & length of next hop NBMA subaddress (see [1]).

Cli Proto Len

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

Preference

This field specifies the preference value for use of the next hop NBMA information specified.

Client NBMA Address

This is the client's NBMA address.

Client NBMA SubAddress

This is the client's NBMA subaddress.

Client Protocol Address

This is the client's internetworking layer address.

4. Values for SCSIP Protocol Independent Part

The following sections give values for fields of the SCSIP Protocol Independent Part of the various SCSIP messages.

4.1 Values for the SCSIP "Mandatory Common Part"

Protocol ID = 0x0002

Sender ID Len = 0x04

Recv ID Len = 0x04

See Section B.2.0.1 of [2] for a detailed description of these fields.

4.2 Values for the SCSP "CSAS Record"

Cache Key Len = 0x04
Orig ID Len = 0x04

See Section B.2.0.2 of [2] for a detailed description of these fields.

5. Security Considerations

Relevant security considerations are documented in [1] and [2].

References

- [1] Luciani, J., Katz, D., Piscitello, D., Cole, B., and N. Doraswamy, "NMBA Next Hop Resolution Protocol (NHRP)", RFC 2332, April 1998.
- [2] Luciani, J., Armitage, G., Halpern, J., and N. Doraswamy, "Server Cache Synchronization Protocol (SCSP)", RFC 2334, April 1998.
- [3] Reynolds, J., and J. Postel, "Assigned Numbers", STD 2, RFC 1700, October 1994.
- [4] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

Acknowledgments

I would like to thank (in no particular order) Maxine Burns of ISR and Joel Halpern of Newbridge. I would also like to thank the members of the ION working group of the IETF, whose review and discussion of this document has been invaluable.

Author's Address

James V. Luciani
Bay Networks, Inc.
3 Federal Street, BL3-03
Billerica, MA 01821

Phone: +1-978-916-4734
EMail: luciani@baynetworks.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.

