dhc Group
Internet-Draft
Intended status: Informational
Expires: February 10, 2007

B. Volz
R. Droms
Cisco Systems, Inc.
August 9, 2006

DHCPv6 Server Reply Sequence Number Option
draft-volz-dhc-dhcpv6-srsn-option-00.txt

Status of this Memo

Copyright Notice

Abstract

This memo defines the Server Reply Sequence Number option for the Dynamic Host Configuration Protocol for IPv6 (DHCPv6). This option is sent from a DHCPv6 server to a DHCPv6 relay agent to allow a relay agent to detect proper sequencing of Relay-Reply messages that could be delivered out of order.

Table of Contents

1. Introduction

When a DHCPv6 server sends a Reply, there is no guarantee as to the order of delivery of those datagrams sent by a server. A DHCPv6 client is protected against delivery of old Reply messages because of the transaction-id in the message. However, a relay agent receiving Relay-Reply messages and maintaining client state information (e.g., [5]) has no such technique. Thus a delayed earlier Relay-Reply may arrive after other Relay-Reply messages. As an example, suppose a client sends a Request, the Reply (encapsulated in a Relay-Reply) is delayed between server and relay agent. The client retransmits the Request, the retransmitted Reply is processed through the relay agent and then by the client. The client next transacts a Release/Reply sequence, which causes the relay agent to remove the client's state information when relaying the Relay-Reply. However, now the delayed first Request's Reply arrives at the relay agent; if the relay agent were to update the client's state based on this out of order message (e.g., [5]), it would add client state that is no longer valid. The Server Reply Sequence Number (SRSN) option can be used to prevent this as the relay agent can detect and discard out of order message.

To allow a relay agent to detect and discard out of order messages, the relay agent requests the server to include a SRSN option in Relay-Reply messages. The SRSN option contains a monotonically increasing sequence number that the relay agent can use to re-sequence (or discard) out of order Relay-Reply messages from the server.

2. Terminology

The key words MUST, MUST NOT, REQUIRED, SHALL, SHALL NOT, SHOULD, SHOULD NOT, RECOMMENDED, MAY, and OPTIONAL in this document are to be interpreted as described in RFC 2119 [1].

Additional terms used in the description of DHCPv6 and DHCPv6 prefix delegation are defined in [2] and [3].

3. The Server Reply Sequence Number Option

The SRSN option is used by a server to indicate the order in which the server has generated replies and therefore allows a relay agent receiving Relay-Reply messages to determine the order in which those Relay-Reply messages were originally sent. The Relay-Reply messages are sent via UDP and, therefore, may be delivered out of order.

The server's sequence number in the SRSN is a monotonically

increasing series. This property is maintained by the server even if the server loses internal state; e.g., if the server is restarted. The value of the sequence number in the SRSN option is not related to the contents of any options in the Relay-Reply message; the existence of this sequence number does not indicate that any data at the server has necessarily changed.

The DHCPv6 Server Reply Sequence Number Option has the following format:

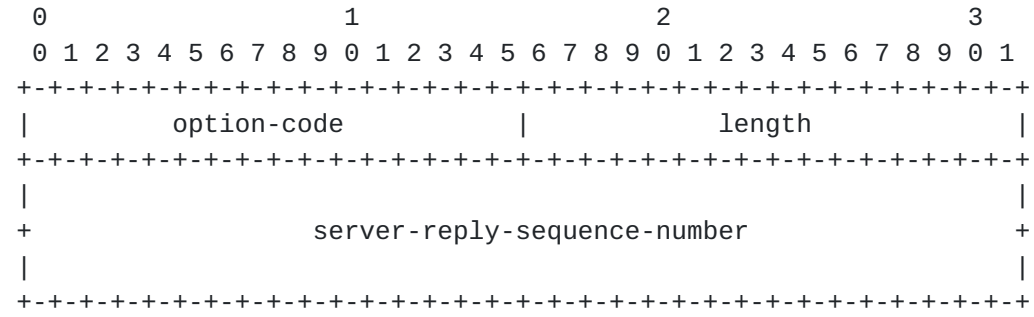

option-code OPTION_SRSN (TBD).

length 8.

server-reply-sequence-number
The 64-bit monotonically increasing server reply sequence number.

4. DHCPv6 Relay Agent Behavior

If a relay agent requires the server to provide the SRSN option, it MUST include an Option Request option, requesting the SRSN option, as described in section 22.7 of [2].

A relay agent MUST save the received server-reply-sequence-number (and the server's Server Identifier Option, OPTION_SERVERID) with any client state information extracted from a Relay-Reply if it needs to assure it does not use out of date information.

A relay agent that uses the SRSN option needs do the following when maintaining client state information:

1. Only update existing client state information if the received server-reply-sequence-number (if from the same server) is greater than the stored server-reply-sequence-number for the information;

the received server-reply-sequence-number and Server Identifier must be stored with the client state information.

2. Delay removing expired client state information from its storage for at least the maximum lifetime of a datagram. This assures that any undelivered Relay-Reply datagrams will have expired and been dropped from the network; and thus the server-reply-sequence-number checking will prevent outdated information from being used. A value of 2 minutes is the recommended value for the maximum datagram lifetime, based on the maximum segment lifetime used by the Transmission Control Protocol (TCP) [4].

A change in the server-reply-sequence-number MUST NOT be used to assume a client's state has changed, as a server may be retransmitting the same information but with a different server-reply-sequence-number.

5. DHCPv6 Server Behavior

If a relay agent has requested the SRSN option in an ORO, the server SHOULD return the option with a monotonically increasing sequence number. And, the server MUST also include a Server Identifier Option (OPTION_SERVER_ID) in the Relay-Reply if it includes the SRSN option.

The server MUST monotonically increase the sequence number for any Relay-Reply messages it transmits which include a SRSN option. A server MAY increase the sequence number for each message it transmits, even those that do not include a SRSN option.

One technique for a server to provide the monotonically increasing sequence number is by splitting the 64-bit number into two 32-bit values (minding network/host byte ordering) - a major (most significant bits) and minor sequence number. When the server starts, the major sequence number is read from stable storage if available, incremented, and saved to stable storage. If no sequence number is in stable storage, the current time (in seconds since Jan 1, 1970) is used to seed the sequence number and write it to stable storage. The minor sequence number is set to 0 and only it is incremented while the server is running (except if it rolls over, in which case the major sequence number MUST be updated); there is no need to commit the minor sequence number to stable storage.

6. IANA considerations

IANA is requested to assign an option code from the "DHCPv6 and DHCPv6 options" registry,

http://www.iana.org/assignments/dhcpv6-parameters, to OPTION_SRSN.

7. Security considerations

Security issues related to DHCP are described in [2] and [3].

The DHCPv6 Server Reply Sequence Number option may be used to mount a denial of service attack by causing a relay agent to incorrectly record a very high server-reply-sequence-number and thus preventing legitimate Relay-Reply messages from a server from being processed. Communication between a server and a relay agent, and communication between relay agents, can be secured through the use of IPSec, as described in section 21.1 of [2].

8. References

8.1. Normative References

[1] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[2] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

[3] Troan, O. and R. Droms, "IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6", RFC 3633, December 2003.

8.2. Informative References

[4] Postel, J., "Transmission Control Protocol", STD 7, RFC 793, September 1981.

[5] Droms, R., Volz, B., and O. Troan, "DHCP Relay Agent Assignment Notification Option (draft-ietf-dhc-dhcpv6-agentopt-delegate-*)", August 2006.

Authors' Addresses

Bernie Volz
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone: +1 978.936.0382
Email: volz@cisco.com

Ralph Droms
Cisco Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA

Phone: +1 978.936.1674
Email: rdroms@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).