

Network Working Group
Request for Comments: 4039
Category: Standards Track

S. Park
P. Kim
SAMSUNG Electronics
B. Volz
Cisco Systems
March 2005

Rapid Commit Option for the
Dynamic Host Configuration Protocol version 4 (DHCPv4)

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 (2005).

Abstract

This document defines a new Dynamic Host Configuration Protocol version 4 (DHCPv4) option, modeled on the DHCPv6 Rapid Commit option, for obtaining IP address and configuration information using a 2-message exchange rather than the usual 4-message exchange, expediting client configuration.

Table of Contents

1. Introduction.....	2
2. Requirements.....	2
3. Client/Server Operations.....	2
3.1. Detailed Flow.....	3
3.2. Administrative Considerations.....	6
4. Rapid Commit Option Format.....	7
5. IANA Considerations.....	7
6. Security Considerations.....	7
7. References.....	7
7.1. Normative References.....	7
7.2. Informative References.....	8
8. Acknowledgements.....	8
Authors' Addresses.....	9
Full Copyright Statement.....	10

1. Introduction

In some environments, such as those in which high mobility occurs and the network attachment point changes frequently, it is beneficial to rapidly configure clients. And, in these environments it is possible to more quickly configure clients because the protections offered by the normal (and longer) 4-message exchange may not be needed. The 4-message exchange allows for redundancy (multiple DHCP servers) without wasting addresses, as addresses are only provisionally assigned to a client until the client chooses and requests one of the provisionally assigned addresses. The 2-message exchange may therefore be used when only one server is present or when addresses are plentiful and having multiple servers commit addresses for a client is not a problem.

This document defines a new Rapid Commit option for DHCPv4, modeled on the DHCPv6 Rapid Commit option [RFC3315], which can be used to initiate a 2-message exchange to expedite client configuration in some environments. A client advertises its support of this option by sending it in DHCPDISCOVER messages. A server then determines whether to allow the 2-message exchange based on its configuration information and can either handle the DHCPDISCOVER as defined in [RFC2131] or commit the client's configuration information and advance to sending a DHCPACK message with the Rapid Commit option as defined herein.

2. Requirements

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 [RFC2119].

3. Client/Server Operations

If a client that supports the Rapid Commit option intends to use the rapid commit capability, it includes a Rapid Commit option in DHCPDISCOVER messages that it sends. The client MUST NOT include it in any other messages. A client and server only use this option when configured to do so.

A client that sent a DHCPDISCOVER with Rapid Commit option processes responses as described in [RFC2131]. However, if the client receives a DHCPACK message with a Rapid Commit option, it SHOULD process the DHCPACK immediately (without waiting for additional DHCPOFFER or DHCPACK messages) and use the address and configuration information contained therein.

A server that supports the Rapid Commit option MAY respond to a DHCPDISCOVER message that included the Rapid Commit option with a DHCPACK that includes the Rapid Commit option and fully committed address and configuration information. A server MUST NOT include the Rapid Commit option in any other messages.

The Rapid Commit option MUST NOT appear in a Parameter Request List option [RFC2132].

All other DHCP operations are as documented in [RFC2131].

3.1. Detailed Flow

The following is revised from Section 3.1 of [RFC2131], which includes handling of the Rapid Commit option.

1. The client broadcasts a DHCPDISCOVER message on its local physical subnet. If the client supports the Rapid Commit option and intends to use the rapid commit capability, it includes a Rapid Commit option in the DHCPDISCOVER message. The DHCPDISCOVER message MAY include options that suggest values for the network address and lease duration. BOOTP relay agents may pass the message on to DHCP servers not on the same physical subnet.
2. Each server may respond with either a DHCPOFFER message or a DHCPACK message with the Rapid Commit option (the latter only if the DHCPDISCOVER contained a Rapid Commit option and the server's configuration policies allow use of Rapid Commit). These would include an available network address in the 'yiaddr' field (and other configuration parameters in DHCP options). Servers sending a DHCPOFFER need not reserve the offered network address, although the protocol will work more efficiently if the server avoids allocating the offered network address to another client. Servers sending the DHCPACK message commit the binding for the client to persistent storage before sending the DHCPACK. The combination of 'client identifier' or 'chaddr' and assigned network address constitute a unique identifier for the client's lease and are used by both the client and server to identify a lease referred to in any DHCP messages. The server transmits the DHCPOFFER or DHCPACK message to the client, if necessary by using the BOOTP relay agent.

When allocating a new address, servers SHOULD check that the offered network address is not already in use; e.g., the server may probe the offered address with an ICMP Echo Request.

Servers SHOULD be implemented so that network administrators MAY choose to disable probes of newly allocated addresses.

Figure 3 in [RFC2131] shows the flow for the normal 4-message exchange. Figure 1 below shows the 2-message exchange.

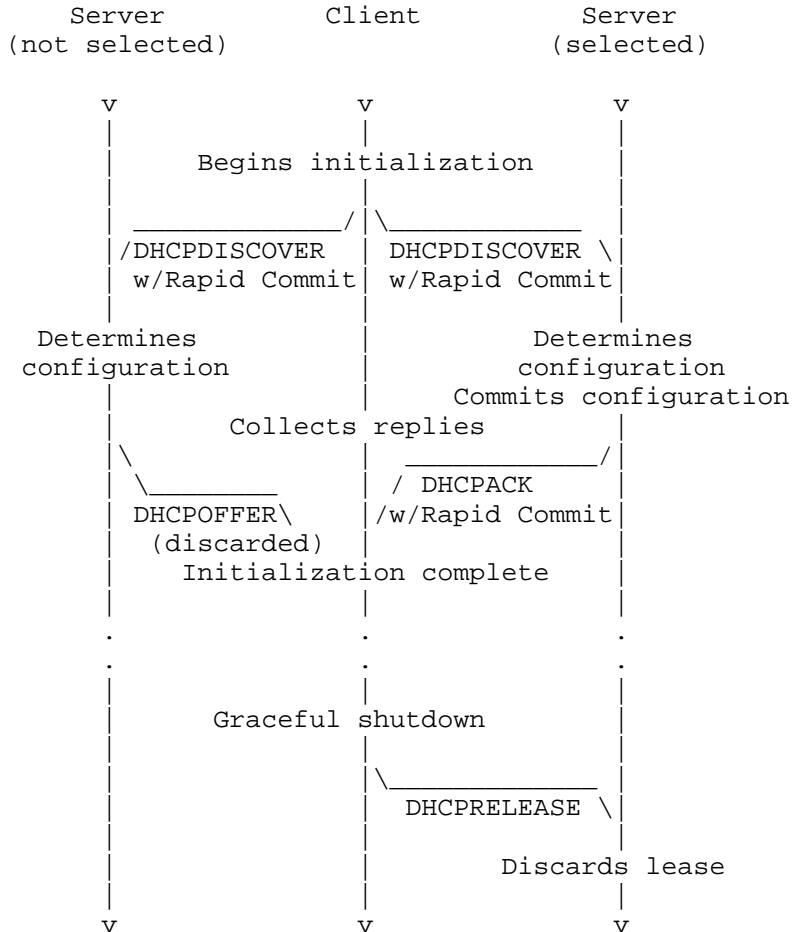


Figure 1 Timeline diagram when Rapid Commit is used

- The client receives one or more DHCP OFFER or DHCPACK (if the Rapid Commit option is sent in DHCPDISCOVER) messages from one or more servers. If a DHCPACK (with the Rapid Commit option) is received, the client may immediately advance to step 5 below if the offered configuration parameters are acceptable. The client may choose to wait for multiple responses. The client chooses one server from which to request or accept

configuration parameters, based on the configuration parameters offered in the DHCPOFFER and DHCPACK messages. If the client chooses a DHCPACK, it advances to step 5. Otherwise, the client broadcasts a DHCPREQUEST message that MUST include the 'server identifier' option to indicate which server it has selected, the message MAY include other options specifying desired configuration values. The 'requested IP address' option MUST be set to the value of 'yiaddr' in the DHCPOFFER message from the server. This DHCPREQUEST message is broadcast and relayed through DHCP/BOOTP relay agents. To help ensure that any BOOTP relay agents forward the DHCPREQUEST message to the same set of DHCP servers that received the original DHCPDISCOVER message, the DHCPREQUEST message MUST use the same value in the DHCP message header's 'secs' field and be sent to the same IP broadcast address as was the original DHCPDISCOVER message. The client times out and retransmits the DHCPDISCOVER message if the client receives no DHCPOFFER messages.

4. The servers receive the DHCPREQUEST broadcast from the client. Servers not selected by the DHCPREQUEST message use the message as notification that the client has declined those servers' offers. The server selected in the DHCPREQUEST message commits the binding for the client to persistent storage and responds with a DHCPACK message containing the configuration parameters for the requesting client. The combination of 'client identifier' or 'chaddr' and assigned network address constitute a unique identifier for the client's lease and are used by both the client and server to identify a lease referred to in any DHCP messages. Any configuration parameters in the DHCPACK message SHOULD NOT conflict with those in the earlier DHCPOFFER message to which the client is responding. The server SHOULD NOT check the offered network address at this point. The 'yiaddr' field in the DHCPACK messages is filled in with the selected network address.

If the selected server is unable to satisfy the DHCPREQUEST message (e.g., the requested network address has been allocated), the server SHOULD respond with a DHCPNAK message.

A server MAY choose to mark addresses offered to clients in DHCPOFFER messages as unavailable. The server SHOULD mark an address offered to a client in a DHCPOFFER message as available if the server receives no DHCPREQUEST message from that client.

5. The client receives the DHCPACK message with configuration parameters. The client SHOULD perform a final check on the parameters (e.g., ARP for allocated network address), and it notes the duration of the lease specified in the DHCPACK

message. At this point, the client is configured. If the client detects that the address is already in use (e.g., through the use of ARP), the client MUST send a DHCPDECLINE message to the server, and it restarts the configuration process. The client SHOULD wait a minimum of ten seconds before restarting the configuration process to avoid excessive network traffic in case of looping.

If the client receives a DHCPNAK message, the client restarts the configuration process.

The client times out and retransmits the DHCPREQUEST message if the client receives neither a DHCPACK nor a DHCPNAK message. The client retransmits the DHCPREQUEST according to the retransmission algorithm in section 4.1 of [RFC2131]. The client should choose to retransmit the DHCPREQUEST enough times to give an adequate probability of contacting the server without causing the client (and the user of that client) to wait too long before giving up; e.g., a client retransmitting as described in section 4.1 of [RFC2131] might retransmit the DHCPREQUEST message four times, for a total delay of 60 seconds, before restarting the initialization procedure. If the client receives neither a DHCPACK nor a DHCPNAK message after employing the retransmission algorithm, the client reverts to INIT state and restarts the initialization process. The client SHOULD notify the user that the initialization process has failed and is restarting.

6. The client may choose to relinquish its lease on a network address by sending a DHCPRELEASE message to the server. The client identifies the lease to be released with its 'client identifier' or 'chaddr' and network address in the DHCPRELEASE message. If the client used a 'client identifier' when it obtained the lease, it MUST use the same 'client identifier' in the DHCPRELEASE message.

3.2. Administrative Considerations

Network administrators MUST only enable the use of Rapid Commit on a DHCP server if one of the following conditions is met:

1. The server is the only server for the subnet.
2. When multiple servers are present, they may each commit a binding for all clients, and therefore each server must have sufficient addresses available.

A server MAY allow configuration for a different (likely shorter) initial lease time for addresses assigned when Rapid Commit is used to expedite reclaiming addresses not used by clients.

4. Rapid Commit Option Format

The Rapid Commit option is used to indicate the use of the two-message exchange for address assignment. The code for the Rapid Commit option is 80. The format of the option is:

```

      Code  Len
    +-----+-----+
    |  80  |  0  |
    +-----+-----+

```

A client MUST include this option in a DHCPDISCOVER message if the client is prepared to perform the DHCPDISCOVER-DHCPACK message exchange described earlier.

A server MUST include this option in a DHCPACK message sent in a response to a DHCPDISCOVER message when completing the DHCPDISCOVER-DHCPACK message exchange.

5. IANA Considerations

IANA has assigned value 80 for the Rapid Commit option code in accordance with [RFC2939].

6. Security Considerations

The concepts in this document do not significantly alter the security considerations for DHCP (see [RFC2131] and [RFC3118]). However, use of this option could expedite denial of service attacks by allowing a mischievous client to consume all available addresses more rapidly or to do so without requiring two-way communication (as injecting DHCPDISCOVER messages with the Rapid Commit option is sufficient to cause a server to allocate an address).

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC2131] Droms, R., "Dynamic Host Configuration Protocol", RFC 2131, March 1997.

7.2. Informative References

- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", RFC 2132, March 1997.
- [RFC2939] Droms, R., "Procedures and IANA Guidelines for Definition of New DHCP Options and Message Types", BCP 43, RFC 2939, September 2000.
- [RFC3118] Droms, R. and W. Arbaugh, "Authentication for DHCP Messages", RFC 3118, June 2001.
- [RFC3315] Droms, R., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, July 2003.

8. Acknowledgements

Special thanks to Ted Lemon and Andre Kostur for their many valuable comments. Thanks to Ralph Droms for his review comments during WGLC. Thanks to Noh-Byung Park and Youngkeun Kim for their supports on this work.

Particularly, the authors would like to acknowledge the implementation contributions by Minhoo Lee of SAMSUNG Electronics.

Authors' Addresses

SooHong Daniel Park
Mobile Platform Laboratory
SAMSUNG Electronics
416, Maetan-3dong, Yeongtong-Gu
Suwon, Korea

Phone: +82-31-200-4508
EMail: soohong.park@samsung.com

Pyungsoo Kim
Mobile Platform Laboratory
SAMSUNG Electronics
416, Maetan-3dong, Yeongtong-Gu
Suwon, Korea

Phone: +82-31-200-4635
EMail: kimps@samsung.com

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

Phone: +1-978-936-0382
EMail: volz@cisco.com

Full Copyright Statement

Copyright (C) The Internet Society (2005).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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