

DHC WG
Internet-Draft
Intended status: Standards Track
Expires: November 29, 2015

Y. Cui
Q. Sun
Tsinghua University
I. Farrer
Deutsche Telekom AG
Y. Lee
Comcast
Q. Sun
China Telecom
M. Boucadair
France Telecom
May 28, 2015

Dynamic Allocation of Shared IPv4 Addresses
draft-ietf-dhc-dynamic-shared-v4allocation-08

Abstract

This memo describes the dynamic allocation of shared IPv4 addresses to clients using DHCPv4. Address sharing allows a single IPv4 address to be allocated to multiple active clients simultaneously, each client being differentiated by a unique set of transport layer source port numbers. The necessary changes to existing DHCPv4 client and server behavior are described and a new DHCPv4 option for provisioning clients with shared IPv4 addresses is included.

Due to the nature of IP address sharing, some limitations to its applicability are necessary. This memo describes these limitations and recommends suitable architectures and technologies where address sharing may be utilized.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

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."

This Internet-Draft will expire on November 29, 2015.

Copyright Notice

Copyright (c) 2015 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	2
2.	Applicability Statement	3
3.	Requirements Language	3
4.	Terminology	3
5.	Functional Overview	4
6.	Client-Server Interaction	4
7.	Client Behavior	5
7.1.	Restrictions to Client Usage of a Shared IPv4 Address . .	6
8.	Server Behavior	7
8.1.	Leasing Shared and Non-Shared IPv4 Addresses from a Single DHCP 4o6 Server	8
9.	DHCPv4 Port Parameters Option	8
10.	Security Considerations	9
10.1.	Port Randomization	10
11.	IANA Considerations	10
12.	Acknowledgements	10
13.	References	11
13.1.	Normative References	11
13.2.	Informative References	11
	Authors' Addresses	12

[1.](#) Introduction

The shortage of available public IPv4 addresses means that it is not always possible for operators to allocate a full IPv4 address to every connected device. This problem is particularly acute whilst an operator is migrating from their existing, native IPv4 network to a native IPv6 network with IPv4 provided as an overlay service. During this phase, public IPv4 addresses are needed to provide for both existing and transition networks.

Two main types of solutions have emerged to address the problem (see [Appendix A of \[RFC6269\]](#)):

1. Deploying Carrier Grade Network Address Translation devices (CGNAT, [\[RFC6888\]](#)).
2. Distributing the same public IPv4 address to multiple clients differentiated by non-overlapping layer 4 port sets.

This memo focuses on the second category of solutions.

[RFC7341] introduces a "DHCP 4o6 Server", which offers dynamic leasing for IPv4 addresses to clients as in DHCPv4 [\[RFC2131\]](#) but transported within a DHCPv6 message flow. This memo specifies a new DHCPv4 option: `OPTION_V4_PORTPARAMS`, and describes how it can be used for the dynamic leasing of shared IPv4 addresses.

Although DHCPv4 over DHCPv6 is used as the underlying DHCPv4 transport mechanism throughout this document, `OPTION_V4_PORTPARAMS` as a DHCPv4 option may also be used in other solutions, if required.

2. Applicability Statement

The solution allows multiple hosts to be simultaneously allocated the same IP address. As the IP address is no longer a unique identifier for a host, this extension is only suitable for specific architectures based on the Address plus Port model (A+P) [\[RFC6346\]](#). Specifically, this document presents a solution that applies to [\[I-D.ietf-softwire-lw4over6\]](#) and certain configurations of [\[I-D.ietf-softwire-map\]](#) (e.g., EA-bit length set to 0).

The solution should only be used on point-to-point links, tunnels, and/or in environments where authentication at the link layer is performed before IP address assignment. It is not suitable for network access over shared media, including Ethernet, WLAN, cable, etc..

3. Requirements Language

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 [\[RFC2119\]](#).

4. Terminology

This document makes use of the following terms:

Shared IPv4 address: An IPv4 address with a restricted layer 4 port set.

Port Set ID (PSID): Identifier for a range of ports assigned to a DHCP client.

5. Functional Overview

Functionally, the dynamic allocation of shared IPv4 addresses by the DHCP 4o6 Server is similar to the dynamic allocation process for 'full' IPv4 addresses described in [\[RFC2131\]](#). The essential difference is that the DHCP 4o6 Server can allocate the same IPv4 address to more than one DHCP 4o6 client simultaneously, providing that each shared address allocation also includes a range of layer 4 source ports unique to that address (i.e., the combined tuple of IPv4 address and Port Set ID is to be unique for each active lease).

The DHCP 4o6 client implements `OPTION_V4_PORTPARAMS` (described below), which is a DHCPv4 option containing PSID (Port Set ID) information. The client includes this option within the Parameter Request List option [\[RFC2132\]](#) in its DHCPv4 DHCPDISCOVER and DHCPREQUEST messages, indicating its support for shared, dynamic address leasing to the DHCP 4o6 server.

`OPTION_V4_PORTPARAMS` is also implemented by the server to identify clients that support shared, dynamic address leasing. With this option, the server can dynamically allocate PSIDs to clients and maintain shared IPv4 address leases. The server then manages unique client leases based the IPv4 address and PSID tuple, instead of using only the IPv4 address.

In the event that a dynamic, shared addressing capable client receives more than one DHCP 4o6 offer, where a received offer does not contain `OPTION_V4_PORTPARAMS` (i.e., is an offer for a full IPv4 address), then the client SHOULD prefer the full IPv4 offer over the shared IPv4 address offer(s), unless specifically configured otherwise.

6. Client-Server Interaction

The following DHCPv4 message flow is transported within the DHCPv4-query and DHCPv4-response messages as in DHCPv4 over DHCPv6 [\[RFC7341\]](#).

1. When the client constructs the DHCPv4 DHCPDISCOVER message to be transported within the DHCPv4-query message, the DHCPDISCOVER message MUST include the client identifier option (constructed as per [\[RFC4361\]](#)) and the Parameter Request List (PRL) option with the code of `OPTION_V4_PORTPARAMS`. The client MAY insert an `OPTION_V4_PORTPARAMS` with preferred values in related fields as a suggestion to the DHCP 4o6 Server.

2. DHCP 4o6 Servers that receive the DHCPDISCOVER message and support shared IPv4 addresses respond with a DHCPOFFER message with the shared IPv4 address in the 'yiaddr' field and MUST add an OPTION_V4_PORTPARAMS option containing an available restricted port set. If the DHCPDISCOVER included an OPTION_V4_PORTPARAMS option containing a non-zero PSID-Len field, the DHCP 4o6 Server MAY allocate a port set of the requested size to the client (depending on policy). The DHCPOFFER message is then encapsulated in the DHCPv4-response message and sent to the client.
3. The client evaluates all received DHCPOFFER messages and selects one (e.g., based on the configuration parameters received, such as the size of the offered port set). The client then sends a DHCPREQUEST encapsulated in the DHCPv4-query message containing the corresponding OPTION_V4_PORTPARAMS received in the DHCPOFFER message.
4. The server identified in the DHCPREQUEST message creates a binding for the client. The binding includes the client identifier, the IPv4 address and the PSID. These parameters are used by both the server and the client to identify a lease in any DHCP message. The server MUST respond with a DHCPACK message containing OPTION_V4_PORTPARAMS for the requesting client.
5. On receipt of the DHCPACK message with the configuration parameters, the client MUST NOT perform an in-use probe on the address, such as ARPing for a duplicate allocated address.
6. If the client chooses to relinquish its lease by sending a DHCPRELEASE message, the client MUST include the leased network address and OPTION_V4_PORTPARAMS (with the allocated PSID) to identify the lease to be released.

In the case that the client has stored the previously allocated address and restricted port set, the logic described in [Section 3.2 of \[RFC2131\]](#) MUST be followed on the condition that the client's source IPv6 address for DHCP 4o6 does not change. Note, this corresponds to the INIT-REBOOT state defined in [\[RFC2131\]](#). The client MUST include the OPTION_V4_PORTPARAMS with the requested port set information in the message flow, which starts with a DHCPREQUEST message. If the client's DHCP 4o6 IPv6 source address is changed for any reason, the client MUST re-initiate the DHCP 4o6 shared-address provisioning process by sending a DHCPDISCOVER message.

7. Client Behavior

A DHCP 4o6 client sending a DHCPDISCOVER message for a shared IPv4 address MUST include the OPTION_V4_PORTPARAMS option code in the Parameter Request List option. If a client has been successfully allocated an IPv4 address and PSID previously, the client's DHCPDISCOVER message MAY include the 'requested IP address' option

along with an `OPTION_V4_PORTPARAMS` to request that a specific IPv4 address and PSID be re-assigned. Alternatively, the client MAY omit the 'requested IP address' option, but include an `OPTION_V4_PORTPARAMS` with a non-zero value in only the PSID-Len field, as a hint to the server for the preferred size of the port set.

A client that requests `OPTION_V4_PORTPARAMS`, but receives `DHCP OFFER` and `DHCP ACK` messages without `OPTION_V4_PORTPARAMS` SHOULD proceed as defined in [\[RFC7341\]](#) and configure a full IPv4 address with no address sharing (see [Section 8.1](#) for the server's behavior).

When receiving a `DHCP ACK` message containing `OPTION_V4_PORTPARAMS`, the client MUST use the received explicit PSID for configuring the interface for which the DHCP 4o6 request was made.

The client MUST NOT probe a newly received IPv4 address (e.g., using ARP) to see if it is in use by another host.

When the client renews or releases its DHCP lease, it MUST put the values of offset, PSID length and PSID into `OPTION_V4_PORTPARAMS`, and send it to the server within corresponding DHCPv4 messages that are conveyed through DHCPv4-query message.

In the event that the client's DHCP 4o6 IPv6 source address is changed for any reason, the client MUST re-initiate the DHCP 4o6 shared-address provisioning process by sending a `DHCP DISCOVER` message.

7.1. Restrictions to Client Usage of a Shared IPv4 Address

As a single IPv4 address is being shared between a number of different clients, the allocated shared address is only suitable for certain uses. The client MUST implement a function to ensure that only the allocated layer 4 ports of the shared IPv4 address are used for sourcing new connections, or accepting inbound connections.

The client MUST apply the following rules for all traffic destined to or originating from the shared IPv4 address:

- o The client MUST use only port-aware protocols (e.g., TCP, UDP, DCCP etc.) or ICMP implementing [\[RFC5508\]](#).
- o All connections originating from the shared IPv4 address MUST use a source port taken from the allocated restricted port set.
- o The client MUST NOT accept inbound connections on ports outside of the allocated restricted port set.

In order to prevent addressing conflicts which could arise from the allocation of the same IPv4 address, the client MUST NOT use the received restricted IPv4 address to perform ARP operations.

The mechanism by which a client implements the above rules is out of the scope of this document.

In the event that the DHCPv4 over DHCPv6 configuration mechanism fails for any reason, the client MUST NOT configure an IPv4 link-local address [[RFC3927](#)] (taken from the 169.254.0.0/16 range).

8. Server Behavior

The DHCP 4o6 Server MUST NOT reply with OPTION_V4_PORTPARAMS unless the client has explicitly listed the option code in the Parameter Request List (Option 55) [[RFC2132](#)].

The DHCP 4o6 Server SHOULD reply with OPTION_V4_PORTPARAMS if the client includes OPTION_V4_PORTPARAMS in its Parameter Request List. In order to achieve the dynamic management of shared IPv4 addresses, the server is required to implement an address and port-set pool that provides the same function as the address pool in a regular DHCP server. Also, the server uses the combination of address and PSID as the key for maintaining the state of a lease, and for searching for an available lease for assignment. The leasing database is required to include the IPv4 address, PSID and client identifier of the requesting client.

When a server receives a DHCPDISCOVER message with OPTION_V4_PORTPARAMS in the Parameter Request List option, the server determines an IPv4 address with a PSID for the requesting client. If an IPv4 address with a PSID is available, the server SHOULD follow the logic below to select which specific address and PSID to provision to the client. The logic is similar to that in [Section 4.3.1 of \[RFC2131\]](#).

- o The client's current address with the PSID as recorded in the client's current lease binding, ELSE
- o The client's previous address with PSID as recorded in the client's (expired or released) binding, if that address with PSID is in the server's pool of available addresses and PSIDs, and not already allocated, ELSE
- o The address requested in the 'Requested IP Address' option along with the PSID parameters requested in the OPTION_V4_PORTPARAMS, if that pair of address and PSID is valid and not already allocated, ELSE
- o A new address with a PSID allocated from the server's pool of available addresses and PSIDs.

Upon receipt of a DHCPRELEASE message with OPTION_V4_PORTPARAMS, the server searches for the lease using the address in the 'ciaddr' field and the PSID information in the OPTION_V4_PORTPARAMS, and marks the lease as unallocated if a record (matching that PSID) is maintained by the server for that client.

The port-set assignment MUST be coupled with the address assignment process. Therefore the server MUST assign the address and port set in the same DHCP message.

When defining the pools of IPv4 addresses and PSIDs which are available to lease to clients, the server MUST implement a mechanism to reserve some port ranges (e.g., 0-1023) from allocation to clients. The reservation policy SHOULD be configurable.

8.1. Leasing Shared and Non-Shared IPv4 Addresses from a Single DHCP 4o6 Server

A single DHCP 4o6 server may serve clients that do not support OPTION_V4_PORTPARAMS as well as those that do. As the rules for the allocation of shared addresses differ from the rules for full IPv4 address assignment, the DHCP 4o6 server MUST implement a mechanism to ensure that clients not supporting OPTION_V4_PORTPARAMS do not receive shared addresses. For example, two separate IPv4 addressing pools could be used, one of which allocates IPv4 addresses and PSIDs only to clients that have requested them.

If the server is only configured with address pools for shared address allocation, it MUST discard requests that do not contain OPTION_V4_PORTPARAMS in the Parameter Request List option.

A server configured with non-shared address pools can be instructed to honor received requests that contain OPTION_V4_PORTPARAMS in the Parameter Request List option (that is ignore OPTION_V4_PORTPARAMS and serve the requesting clients with non-shared IPv4 addresses).

9. DHCPv4 Port Parameters Option

The meaning of 'offset', 'PSID-len', and 'PSID' fields of the DHCPv4 Port Parameters Option is identical to that of 'offset', 'PSID-len', and 'PSID' fields of the S46 Port Parameters Option (Section 4.5 of [\[I-D.ietf-softwire-map-dhcp\]](#)). The use of the same encoding in both options is meant to ensure compatibility with existing port set implementations.

The format of OPTION_V4_PORTPARAMS is shown in Figure 1.

The security considerations described in [RFC2131] and [RFC7341] are also potentially applicable to this solution. Unauthorised DHCP 4o6 servers in the network could be used to stage an amplification attack or to supply invalid configuration leading to service disruption. The risks of these types of attacks can be reduced through the use of unicast DHCP 4o6 message flows (enabled by supplying DHCP 4o6 server unicast addresses within the OPTION DHCP4_0 DHCP6_SERVER option).

A malicious user could attempt a DoS attack by requesting a large number of IPv4 address (or fractional address) and port sets allocations, exhausting the available addresses and port sets for other clients. This can be mitigated through DHCPv4 address allocation policy, limiting the number of simultaneously active IPv4 leases for clients whose request originate from each customer site.

Additional security considerations are discussed in Section 11 of [\[I-D.ietf-softwire-map\]](#) and [Section 9 of \[I-D.ietf-softwire-lw4over6\]](#).

10.1. Port Randomization

Preserving port randomization [\[RFC6056\]](#) may be more difficult because the host can only randomize the ports inside a fixed port range (see [Section 13.4 of \[RFC6269\]](#)).

More discussion to improve the robustness of TCP against Blind In-Window Attacks can be found at [\[RFC5961\]](#). Other means than the (IPv4) source port randomization to provide protection against attacks should be used (e.g., use [\[RFC5961\]](#) to improve the robustness of TCP against Blind In-Window Attacks, use IPv6).

11. IANA Considerations

IANA is requested to assign the following new DHCPv4 Option Code in the registry maintained in: <http://www.iana.org/assignments/bootp-dhcp-parameters/>:

Option Name	Value	Data length	Meaning
OPTION_V4_PORTPARAMS	TBA	4	This option is used to configure a set of ports bound to a shared IPv4 address.

12. Acknowledgements

This document is merged from [\[I-D.sun-dhc-port-set-option\]](#) and [\[I-D.farrer-dhc-shared-address-lease\]](#).

The authors would like to thank Peng Wu, Gabor Bajko, Teemu Savolainen, Ted Lemon, Tina Tsou, Pierre Levis, Cong Liu, and Marcin Siodelski for their contributions.

Many thanks to Brian Haberman for the review.

13. References

13.1. Normative References

- [I-D.ietf-softwire-lw4over6]
Cui, Y., Qiong, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", [draft-ietf-softwire-lw4over6-13](#) (work in progress), November 2014.
- [I-D.ietf-softwire-map]
Troan, O., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, "Mapping of Address and Port with Encapsulation (MAP)", [draft-ietf-softwire-map-13](#) (work in progress), March 2015.
- [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.
- [RFC2132] Alexander, S. and R. Droms, "DHCP Options and BOOTP Vendor Extensions", [RFC 2132](#), March 1997.
- [RFC4361] Lemon, T. and B. Sommerfeld, "Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)", [RFC 4361](#), February 2006.
- [RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", [RFC 5961](#), August 2010.
- [RFC6056] Larsen, M. and F. Gont, "Recommendations for Transport-Protocol Port Randomization", [BCP 156](#), [RFC 6056](#), January 2011.
- [RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", [RFC 7341](#), August 2014.

13.2. Informative References

- [I-D.farrer-dhc-shared-address-lease]
Farrer, I., "Dynamic Allocation of Shared IPv4 Addresses using DHCPv4 over DHCPv6", [draft-farrer-dhc-shared-address-lease-00](#) (work in progress), June 2013.

- [I-D.ietf-softwire-map-dhcp] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for configuration of Softwire Address and Port Mapped Clients", [draft-ietf-softwire-map-dhcp-12](#) (work in progress), March 2015.
- [I-D.sun-dhc-port-set-option] Qiong, Q., Lee, Y., Sun, Q., Bajko, G., and M. Boucadair, "Dynamic Host Configuration Protocol (DHCP) Option for Port Set Assignment", [draft-sun-dhc-port-set-option-02](#) (work in progress), October 2013.
- [RFC3927] Cheshire, S., Aboba, B., and E. Guttman, "Dynamic Configuration of IPv4 Link-Local Addresses", [RFC 3927](#), May 2005.
- [RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", [BCP 148](#), [RFC 5508](#), April 2009.
- [RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", [RFC 6269](#), June 2011.
- [RFC6335] Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S. Cheshire, "Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry", [BCP 165](#), [RFC 6335](#), August 2011.
- [RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", [RFC 6346](#), August 2011.
- [RFC6888] Perreault, S., Yamagata, I., Miyakawa, S., Nakagawa, A., and H. Ashida, "Common Requirements for Carrier-Grade NATs (CGNs)", [BCP 127](#), [RFC 6888](#), April 2013.

Authors' Addresses

Yong Cui
Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6260-3059
Email: yong@csnet1.cs.tsinghua.edu.cn

Qi Sun
Tsinghua University
Beijing 100084
P.R. China

Phone: +86-10-6278-5822
Email: sunqi@csnet1.cs.tsinghua.edu.cn

Ian Farrer
Deutsche Telekom AG
CTO-ATI, Landgrabenweg 151
Bonn, NRW 53227
Germany

Email: ian.farrer@telekom.de

Yiu L. Lee
Comcast
One Comcast Center
Philadelphia PA 19103
USA

Email: yiu_lee@cable.comcast.com

Qiong Sun
China Telecom
Room 708, No.118, Xizhimennei Street
Beijing 100035
P.R. China

Phone: +86-10-58552936
Email: sunqiong@ctbri.com.cn

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

