DHC Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 15, 2014

Q. Sun
Tsinghua University
Y. Lee
Comcast
Q. Sun
China Telecom
G. Bajko
Nokia
M. Boucadair
France Telecom

October 12, 2013

# Dynamic Host Configuration Protocol (DHCP) Option for Port Set Assignment
## draft-sun-dhc-port-set-option-02

Abstract

Because of the exhaustion of the IPv4 address space, several techniques have been proposed to share the same IPv4 address among several uses. As an alternative to introducing a level of NAT in the provider's core network, this document provides a mechanism to assign non-overlapping layer 4 port sets to users assigned with the same IPv4 address: Port Set DHCPv4 Option.

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 April 15, 2014.

Table of Contents

## 1. Introduction

Currently some large ISPs still have a large enough IPv4 address pool to be able to allocate public IPv4 addresses for their subscribers. However, due to the exhaustion of the global IPv4 address space, these ISPs expect the situation is unsustainable and they will not be able anymore to assign to every requesting host a public IPv4 address.

Two solutions have been proposed so far: (1) Deploy Network Address Translation (NAT) or (2) Allocate the same public IPv4 address with non-overlapped layer 4 port sets directly to multiple connected devices (which can be CPEs or end hosts). This document focuses on the second solution.

This document describes a new DHCPv4 option which allows the DHCPv4 server to assign a set of ports to a user device during the IPv4 address provisioning process. By assigning the same IPv4 address with non-overlapped port sets to multiple clients, the clients is enabled to share the IPv4 address and continue to deliver IPv4 services to subscribers.

When using this DHCPv4 option, the underlaying forwarding carrier should be other than IPv4 to avoid affecting the current IPv4-only architecture, for example IPv6 [I-D.ietf-dhc-dhcpv4-over-ipv6], [I-D.ietf-dhc-dhcpv4-over-dhcpv6], etc. The server has to manage to forward DHCP responses to the right client.

The Port Set Option described in this document can be used in various deployment scenarios, some of which are described in [RFC6346]

## 2. 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].

## 3. DHCPv4 Port Set Option

### 3.1. Port Set Option Format

The format of Port Set Option is shown in Figure 1.

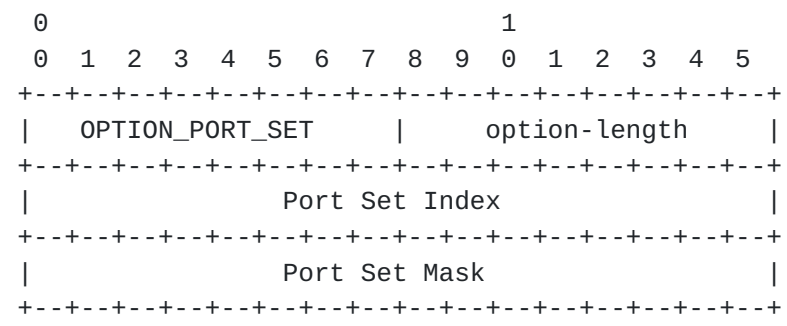


Figure 1 Port Set Option Format

- option-code: OPTION_PORT_SET (TBD)

- option-length: An 8-bit field indicating the length of the option excluding the 'Option Code' and the 'Option Length' fields. In this option, the option-length is 4 octets.

- Port Set Index: Port Set Index identifies a set of ports assigned to a device. The first k bits on the left of the 2-octet field is the Port Set Index value, with the rest of the field right padding zeros.

- Port Set Mask: Port Set Mask indicates the position of the bits used to build the mask. The first k bits on the left is padding ones while the remained (16-k) bits of the 2-octet field on the right is padding zeros.

In the context of Port Set Option, the port number should consist of port set prefix and port number suffix. The port set prefix can be got from Port Set Index and Port Set Mask, while port number suffix can change continuously. The format of port number is shown in Figure 2.

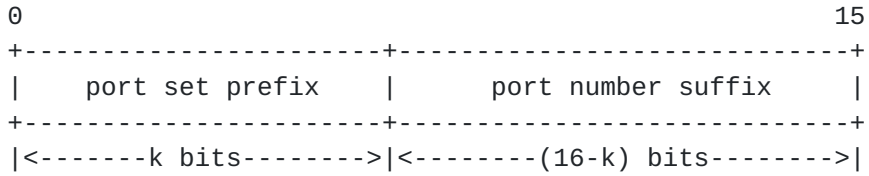


Figure 2 Bit Representation of a port number

In order to exclude the system ports ([I-D.ietf-tsvwg-iana-ports]) or ports saved by SPs, the former port-sets that contains well-known ports SHOULD NOT be assigned.

For example: If k is 10 (the left 10 bits of Port Set Mask is '1'), the first 16 port sets is located in well-known port space, which

should not be allocated. Or,

For example: If k is 4 (the left 4 bits of Port Set Mask is '1'), the first port set (0 - 4095) contains the well-know port space. It should be perceived as well.

### 3.2. Port Set Option Example

The Port Set Option is used to specify one contiguous port set pertaining to the given IP address.

Concretely, this option is used to notify a remote DHCP client about the port set prefix to be applied when selecting a port value as a source port. The Port Set Option is used to infer a set of allowed contiguous port values. Two port numbers are said to belong to the same Port Set if and only if, they have the same port set prefix.

The following Port Set Index and Port Set Mask are conveyed using DHCP to assign a contiguous port set with excluding well-know ports (with Port Set Index not zero):

Port Set Index: 0001 0100 0000 0000 (5120)

Port Set Mask: 1111 1100 0000 0000 (64512)

The device will get a contiguous port set: 5120 - 6143

## 4. Server Behavior

The server will not reply with the option until the client has explicitly listed the option code in the Parameter Request List (Option 55).

Server MUST reply with Port Set Option if the client requested OPTION_PORT_SET in its Parameter Request List. In order to achieve the dynamic management of IPv4 address and port set in the address sharing environment, the server MUST run an address & port-set pool that plays the same role as address pool in regular DHCP server. The address and port-set pool MUST follow the Port-Mask-format port-set. The server MUST use a combination of address & port-set as a key to maintain the state of a lease, and look for an available lease for assignment. The leasing database MUST also include the information of the address and port-set.

When a server receives a DHCPDISCOVER message with OPTION_PORT_SET in the Parameter Request List from a client, the server chooses an IPv4 address and a port-set for the requesting client. The logic of

choosing is similar to that in Section 4.3.1 of [RFC2131], while the difference is the server looks for the client's binding or an available lease in the server's pool of addresses & port-sets. After selecting an available combination of an address and a port-set, the server puts the address into the 'yiaddr' field and the port-set (in the Port-Mast-format) into the Port Set Option.

If the server receives a DHCPDISCOVER message containing a Port Set Option, this means the client is requesting a specific port set. The Port Set Mask field in the option indicates the size of port set that the client requests. The server MAY reply with a Port Set Option whose Port Set Mask is as requested, if the server has such one port set. Or the server can ignore the request and just assign a port set from the pool.

When a server receives a DHCPREQUEST message with Port Set Option, the server MUST determine the client's state according to related parameters (Section 4.3.2 of [RFC2131]) and the value of Port Set Option.

Upon reception of a DHCPRELEASE message with Port Set Option, the server looks for the lease using the address and the value in the Port Set Option, and marks it as unallocated.

The port-set assignment MUST be coupled with the address assignment process. Therefore server MUST assign the address and port set in the same DHCP messages. And the lease information for the address is applicable to the port-set as well.

## 5. Client Behavior

The DHCP client applying for a port-set MUST include the OPTION_PORT_SET code in the Parameter Request List (Option 55). The client retrieves a Port Set Option and use the Port Set Index and Port Set Mask to perform the port mask algorithm to get the contiguous port set.

When the client renews or releases the DHCP lease, it MUST put its Port Set Index and Port Set Mask into the Port Set Option, and send to the server within corresponding DHCPv4 messages.

The client MAY include a Port Set Option in the DHCPDISCOVER message, in which the Port Set Mask field indicates the requested size of a port set from the client.

## 6. Security Consideration

### 6.1. Denial-of-Service

The solution is generally vulnerable to DoS when used in shared medium or when access network authentication is not a prerequisite to IP address assignment. The solution SHOULD only be used on point-to-point links, tunnels, and/or in environments where authentication at link layer is performed before IP address assignment, and not shared medium.

### 6.2. Port Randomization

Preserving port randomization [RFC6056] may be more or less difficult depending on the address sharing ratio (i.e., the size of the port space assigned to a CPE). The host can only randomize the ports inside a fixed port range [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 [I-D.vixie-dnsext-dns0x20] to protect against DNS attacks, [RFC5961] to improve the robustness of TCP against Blind In-Window Attacks, use IPv6).

A proposal to preserve the entropy when selecting port is discussed in [I-D.bajko-pripaddrassign]

## 7. IANA Consideration

IANA is kindly requested to allocate DHCP option code to the OPTION_PORT_SET. The code should be added to the DHCP option code space.

## 8. Contributors List

Many thanks for valuable comments and great efforts from the following contributors:

Peng Wu
Tsinghua University

peng-wu@foxmail.com

Teemu Savolainen
Nokia

teemu.savolainen@nokia.com

Ted Lemon
Nominum, Inc.

mellon@nominum.com

Tina Tsou
Huawei Technologies

tena@huawei.com

Pierre Levis
France Telecom

Email: pierre.levis@orange.com

Cong Liu
Tsinghua University

Email: gnocuil@gmail.com

## 9. References

### 9.1. 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.

[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.

[RFC6269] Ford, M., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, June 2011.

[RFC6346] Bush, R., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, August 2011.

## 9.2. Informative References

[I-D.bajko-pripaddrassign]
Bajko, G., Savolainen, T., Boucadair, M., and P. Levis, "Port Restricted IP Address Assignment", draft-bajko-pripaddrassign-04 (work in progress), April 2012.

[I-D.ietf-dhc-dhcpv4-over-dhcpv6]
Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4 over DHCPv6 Transport", draft-ietf-dhc-dhcpv4-over-dhcpv6-01 (work in progress), July 2013.

[I-D.ietf-dhc-dhcpv4-over-ipv6]
Cui, Y., Wu, P., Wu, J., and T. Lemon, "DHCPv4 over IPv6 Transport", draft-ietf-dhc-dhcpv4-over-ipv6-07 (work in progress), September 2013.

[I-D.ietf-tsvwg-iana-ports]
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", draft-ietf-tsvwg-iana-ports-10 (work in progress), February 2011.

[I-D.vixie-dnsext-dns0x20]
Vixie, P. and D. Dagon, "Use of Bit 0x20 in DNS Labels to Improve Transaction Identity", draft-vixie-dnsext-dns0x20-00 (work in progress), March 2008.

[RFC6842] Swamy, N., Halwasia, G., and P. Jhingran, "Client Identifier Option in DHCP Server Replies", RFC 6842,

January 2013.

Authors' Addresses

Qi Sun
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

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

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

Phone:
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

Gabor Bajko
Nokia

Phone:
Email: gabor.Bajko@nokia.com

Mohamed Boucadair
France Telecom
2330 Central Expressway
Rennes 35000
France

Phone:
Email: mohamed.boucadair@orange.com