DHC WG
Internet-Draft
Updates: 2131 (if approved)
Intended status: Standards Track
Expires: August 17, 2014

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

February 13, 2014

Dynamic Allocation of Shared IPv4 Addresses
draft-csf-dhc-dynamic-shared-v4allocation-00

Abstract

This memo describes the dynamic allocation of shared IPv4 addresses to clients using the DHCPv4 protocol. Address sharing allows a single IPv4 address to be allocated to multiple, active clients simultaneously, each client being differentiated by a unique set of L4 source ports. The changes necessary to existing DHCPv4 client and server behaviour are described and a new DHCPv4 option for provisioning clients with shared IPv4 addresses is included.

Due to the nature of sharing IP addresses, there are necessarily some limitations to the applicability. This memo describes those limitations and recommends suitable architectures and technologies where address sharing may be utilized.

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 RFC 2119 [RFC2119].

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/.

Copyright Notice

Table of Contents

1. Introduction

Shortages of available public IPv4 addresses mean that it is not always possible for operators to allocate a full IPv4 address to every customer. This problem may be particularly acute whilst the operator is in the migration phase from a native IPv4 network to a native IPv6 network with IPv4 provided as an overlay service. This is likely to increase the requirement on public IPv4 addresses to provide for both existing and transition networks.

Two main types of solution have emerged to ease the problem:

1. Centralised Network Address Translation (NAT44) in the core network
2. Distributing the same public IPv4 address to multiple clients using non-overlapping layer 4 port sets.

The solution described in this memo is only suitable on the second solution.

[I-D.ietf-dhc-dhcpv4-over-dhcpv6] introduces a "DHCP 4o6 Server", which is capable of servicing both DHCPv6 [RFC3315] and DHCPv4-over-DHCPv6 requests. This enables the provisioning of DHCPv4 based configuration to IPv6 connected clients over IPv6 only transport networks.

One of the benefits of the DHCPv4-over-DHCPv6 based approach is that it allows the dynamic leasing of IPv4 addresses to clients, based on existing mechanisms for address lease management available in DHCPv4 servers. This can make much more efficient use of remaining public IPv4 addresses than static pre-allocation based approaches as only IPv4 clients that are currently active need to be allocated addresses. This memo uses the defined OPTION_PORTPARAMSV4 with DHCPv4 over DHCPv6, achieving the dynamic leasing of the shared IPv4 addresses.

Due to the nature of address sharing in this manner, it is only suitable for specific architectures based on the Address plus Port Model (A+P) [RFC6346]. This model extends the unique identifier for a client from the 32-bit IPv4 address to 48-bits by including the 16-bits of the layer 4 header. Each client is allocated a unique block of layer 4 ports, and the client will generally utilize these restricted source ports by implementing a NAPT44 funtion, translating traffic from the original private IPv4 source address and unrestricted port to the allocated shared IPv4 address and unique restricted port range. [I-D.ietf-softwire-map] and [I-D.ietf-softwire-lw4over6] describe two implemented examples of the

A+P approach which may be suitable for shared, dynamic IPv4 addressing.

The use of shared addressing in other, more traditional deployment architectures must be avoided due to the fundamental incompatibilities of assigning a the same /32 IPv4 address to multiple clients attached to the same layer 2 segment.

This memo also defines OPTION_PORTPARAMSV4, a DHCPv4 option for assigning non-overlapping layer 4 port sets during the IPv4 address allocation process.

Although DHCPv4 over DHCPv6 is used as the underlying DHCPv4 transport mechanism throughout this document, OPTION_PORTPARAMSV4 may also be used in DHCPv4 over IPv6 [I-D.ietf-dhc-dhcpv4-over-ipv6] and other DHCPv4 IPv4 address allocation mechanisms. The usage of OPTION_PORTPARAMSV4 in those cases is out of scope of this document.

2. Functional Overview

Functionally, the dynamic allocation of shared IPv4 addresses by the DHCP 4o6 Server is quite similar to the normal DHCPv4 server dynamic allocation process described in [RFC2131]. The essential difference is that the DHCP 4o6 Server MAY allocate the same IPv4 address to more than one DHCP 4o6 client simultaneously, providing that each address allocation also includes a range of layer 4 source ports unique to that address (i.e. each PSID may only be allocated once per /32 address).

To enable this, the DHCP 4o6 client needs to be extended to implement OPTION_PORTPARAMSV4 (described below). This option is used to indicate to the DHCP 4o6 server the client’s support the dynamic allocation of a shared IPv4 address and also for conveying the allocated PSID back to the client.

The server must be extended to implement OPTION_PORTPARAMSV4 so that it can identify clients supporting shared, dynamic address leasing. With this option, the server can dynamically maintain shared IPv4 address leases. The server must also manage unique client leases based on the IPv4 address and PSID tuple, instead of just IPv4 address.

3. Client-Server Interaction

Section 3 of [RFC2131] describes client-server interactions necessary for leasing addresses. The following sections describe the changes necessary for the client and server to implement the dynamic allocation of a shared IPv4 address.

3.1. Allocating a Shared, Dynamic IPv4 Address

Section 3.1 of [RFC2131] describes the client-server interaction for allocating an IPv4 address. The process described below detail the changes necessary for the allocation of a shared IPv4 address.

Using DHCP 4o6, the following DHCPv4 message flow is transported within the DHCPV4-QUERY and DHCPV4-RESPONSE options, which are DHCPv6 options used for carrying DHCPv4 messages.

1. When the client constructs its DHCPv4 DHCPDISCOVER message to be transported within the DHCPv4-query message, the DHCPDISCOVER message MUST include the following options: A client Identifier (constructed as per [RFC4361] and OPTION_PORTPARAMSV4 (described below). The client MAY insert a non-zero value in the PSID-Len field within OPTION_PORTPARAMSV4 to indicate the preferred size of the restricted port range allocation to the DHCP 4o6 Server.
2. Each DHCP 4o6 Server that receives the DHCPDISCOVER message within the DHCPv4-query message responds with a DHCPOFFER message that contains an available IPv4 address in the 'yiaddr' field. The response MUST also include OPTION_PORTPARAMSV4 containing a restricted port-range. If the received OPTION_PORTPARAMSV4 field contains 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 included into the DHCPv4-response message and sent to the client.
3. The client evaluates all received DHCPOFFER messages and selects one based on the configuration parameters received, such as the size of the offered port set. The client then sends a DHCPREQUEST containing a server identifier and the corresponding OPTION_PORTPARAMSV4 received in the DHCPOFFER message.
4. The server identified in the DHCPREQUEST message (via the siaddr field) 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 referred to in any DHCP messages. The server responds with a DHCPACK message containing the configuration parameters for the requesting client. Optionally, the the server may also store the IPv6 address that the client has bound the received IPv4 paramters to.
5. The client receives the DHCPACK message with the configuration parameters. The client MUST NOT perform a final check 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 original client identifier, the leased network address and the allocated restricted source ports inlcuded in OPTION_PORTPARAMSV4.

3.2. Reusing a Previously Allocated Shared, Dynamic IPv4 address

If the client remembers the previously allocated address and restricted port range, then the process described in section 3.2 of [RFC2131] must be followed. OPTION_PORTPARAMSV4 MUST be included in the message flow, with the client's requested port set being included in the DHCPDISCOVER message.

4. Server Behavior

The DHCP 4o6 Server MUST NOT reply with the OPTION_PORTPARAMSV4 until the client has explicitly listed the option code in the Parameter Request List (Option 55) [RFC2132].

The DHCP 4o6 Server SHOULD reply with OPTION_PORTPARAMSV4 if the client includes the option 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 and port-set pool that plays the same role as address pool in a regular DHCP server. The server MUST use the combination of address and PSID as the key to maintain the state of a lease, and look for an available lease for assignment. The leasing database MUST include the information of the address and PSID.

When a server receives a DHCPDISCOVER message with OPTION_PORTPARAMSV4 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]. The difference is the server looks for the client's binding or an available lease in the server's pool of addresses and PSIDs. After selecting an available IPv4 address with a PSID, the server sends a DHCPOFFER message to the requesting client.

When the server receives a DHCPREQUEST message with OPTION_PORTPARAMSV4, the server MUST determine the client's state according to related parameters (Section 4.3.2 of [RFC2131]) and the value of OPTION_PORTPARAMSV4.

Upon reception of a DHCPRELEASE message with OPTION_PORTPARAMSV4, the server looks for the lease using the address in the message and the PSID value in the OPTION_PORTPARAMSV4, 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. The lease information for the address is applicable to the port-set as well.

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

A single DHCP 4o6 server may have clients that do not support OPTION_PORTPARAMS as well as those that do. As the rules for the allocation of shared addresses differ from the rules for full IPv4 address assigment, the DHCP 4o6 server MUST implement a mechanism to ensure that clients which do not support OPTION_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 which have requested them.

5. Client Behavior

The DHCP client applying for a port-set MUST include the OPTION_PORTPARAMSV4 code in the Parameter Request List (Option 55). The client retrieves a port set using the value contained in OPTION_PORTPARAMSV4.

When the client renews or releases the DHCP lease, it MUST put the values of offset, PSID length and PSID into the OPTION_PORTPARAMSV4, and send to the server within corresponding DHCPv4 messages.

In the DHCPDISCOVER message, the client MAY use a non-zero value for the PSID-len field within OPTION_PORTPARMAS. This is used by the client to request a specific size of port-set (i.e. the number of source ports that it will be allocated).

5.1. Client Usage of a Shared 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.

The client MUST apply the following rules for any traffic to or from the shared /32 IPv4 address:

- Only port-aware protocols or ICMP implementing [RFC5508] MUST be used
- All connections originating from the shared IPv4 address MUST use a source port taken from the allocated restricted port range.
- The client MUST NOT accept inbound connections on ports outside of the allocated restricted port range.

In order to prevent addressing conflicts which could arise from the allocation of the same IPv4 addreses, the client MUST NOT configure the received restricted IPv4 address on-link.

The mechanism by which a client implements these rules is outside 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).

6. Additional Changes to RFC 2131

In addtion to the changes mentioned elsewhere in this document, the following changes to the behaviour described in [RFC2131] are necessary in order to implement dynamic allocation of a shared IPv4 address.

Section 2.2 The client MUST NOT probe a newly received IPv4 address (e.g. with ARP) to see if it is in use by another host.
Section 3.1 Item 5. The client MUST NOT perform a final check on the assigned IPv4 address.

7. DHCPv4 Port Parameters Option

The Port Paramaters Option for DHCPv4 specifies the restricted set of layer 4 source ports that are necessary to dynamically allocate a shared address. The option uses the same fields as the MAP Port Parameters Option described in Section 4.4 of [I-D.ietf-softwire-map-dhcp], implemented as a DHCPv4 option. This is to maintain compatibility with existing implementations.

The construction and usage of OPTION_PORTPARAMSV4 is

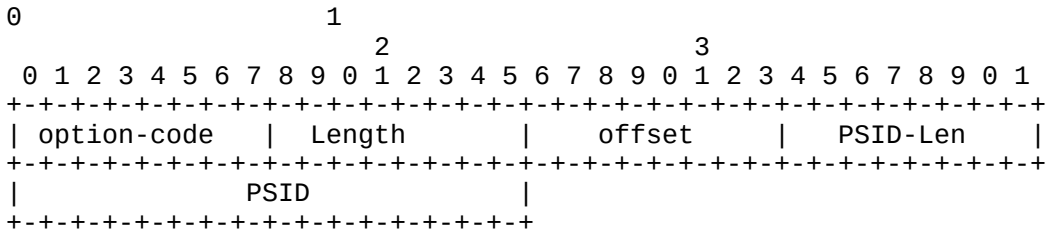

Figure 1: DHCPv4 Port Parameters Option

- option-code: OPTION_PORTPARAMSV4 (TBA)
- option-length: 3

- o offset: (PSID offset) 8 bits long field that specifies the numeric value for the MAP algorithm's excluded port range/offset bits (A-bits), as per section 5.1.1 in [I-D.ietf-softwire-map]. Allowed values are between 0 and 16, with the default value being 4 for a MAP client. This parameter is unused by a Lightweight 4over6 client and should be set to 0.
- o PSID-len: Bit length value of the number of significant bits in the PSID field (also known as 'k'). When set to 0, the PSID field is to be ignored. After the first 'a' bits, there are k bits in the port number representing valid of PSID. Subsequently, the address sharing ratio would be 2^k.
- o PSID: Explicit 16-bit (unsigned word) PSID value. The PSID value algorithmically identifies a set of ports assigned to a CE. The first k-bits on the left of this 2-octets field is the PSID value. The remaining (16-k) bits on the right are padding zeros.

[I-D.ietf-softwire-map] (Section 5.1) provides a full description of how the PSID is interpreted by the client.

When receiveing the Port Parameters option with an explicit PSID, the client MUST use this explicit PSID in configuring its DHCPv4 over DHCPv6 interface.

8. Security Consideration

8.1. Denial-of-Service

The solution is generally vulnerable to DoS when used on a 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.

8.2. 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].

9. IANA Considerations

IANA is kindly requested to allocate the following DHCPv4 option code: TBD for OPTION_PORTPARAMSV4.

10. Acknowledgements

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

11. References

11.1. Normative References

[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-dhcpv6-04 (work in progress), January 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-10 (work in progress), January 2014.

[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.

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

11.2. 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.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-dhc-dhcpv4-over-ipv6]
Cui, Y., Wu, P., Wu, J., Lemon, T., and Q. Sun, "DHCPv4 over IPv6 Transport", draft-ietf-dhc-dhcpv4-over-ipv6-08 (work in progress), October 2013.

[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-lw4over6-06 (work in progress), February 2014.

[I-D.ietf-softwire-map-dhcp]
Mrugalski, T., Troan, O., Dec, W., Bao, C., leaf.yeh.sdo@gmail.com, l., and X. Deng, "DHCPv6 Options for configuration of Softwire Address and Port Mapped Clients", draft-ietf-softwire-map-dhcp-06 (work in progress), November 2013.

[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.

[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.

[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.

[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.

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

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
2330 Central Expressway
Rennes 35000
France

Email: mohamed.boucadair@orange.com