Network Working Group
Internet-Draft
Intended status: Informational
Expires: December 31, 2016

C. Liu
Q. Sun
J. Wu
Tsinghua University
I. Farrer
Deutsche Telekom AG

June 29, 2016

Dynamic IPv4 Provisioning for Lightweight 4over6
draft-liu-softwire-lw4over6-dynamic-provisioning-02

Abstract

Lightweight 4over6 [RFC7596] is an IPv4 over IPv6 hub-and-spoke mechanism that provides overlay IPv4 services in an IPv6-only access network. It uses a deterministic, DHCPv6 based method for the provisioning of IPv4 addresses and port sets to customer CE devices. This document describes how existing specifications can be used for the dynamic IPv4 provisioning of Lightweight 4over6, based on DHCPv4 over DHCPv6 [RFC7341].

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 December 31, 2016.

Table of Contents

1. Introduction 2
2. Terminology 3
3. Dynamic Provisioning Model 4
3.1. Flow 1: lwB4's IPv6 Addressing and DHCPv6 Configuration . 4
3.2. Flow 2: DHCP 4o6 Function 5
3.3. Flow 3: lwAFTR Binding Table Maintainence 5
3.3.1. Flow 3a: Binding Table Maintenance for Co-located lwAFTR/DHCP 4o6 Functions 5
3.3.2. Flow 3b: Binding Table Maintenance for Distributed lwAFTR/DHCP 4o6 Functions 6
4. Security Considerations 6
4.1. Data Retention Requirements 6
5. IANA Considerations 7
6. References 7
6.1. Normative References 7
6.2. Informative References 8
Authors' Addresses 8

1. Introduction

Lightweight 4over6 [RFC7596] (lw4o6) provides IPv4 access over an IPv6 network with a hub-and-spoke softwire architecture. In Lightweight 4over6, each Lightweight B4 (lwB4) is assigned a full, or shared (port-restricted) IPv4 address to be used for IPv4 communication. Provisioning the lwB4 with its IPv4 address, port set and other parameters necessary for building the softwire is a core function of the lw4o6 control plane.

[RFC7596] describes the use of DHCPv6 for deterministic IPv4 provisioning. The IPv4 address, port set ID (PSID) and address of the lwAFTR are carried in DHCPv6 options defined in [RFC7598].

However, the deterministic provisioning of the IPv4 parameters imposes restrictions on the deployment:

- The IPv4 address' life time is bound to the client's IPv6 tunnel endpoint life time

- The tunnel must be initiated from a fixed and predictable /64 prefix in the home network topology

- The IPv4 address and PSID need to be embedded into the IID of the clients' /128 IPv6 address

- IPv4 address resources are permanently reserved for a client whether it is active or not. This results in less efficient public IPv4 address usage

This document describes how lw4o6 uses DHCPv4 over DHCPv6 to achieve dynamic IPv4 address provisioning. The main advantages of using a dynamic provisioning model over a deterministic provisioning model are as follows:

- No inherent restrictions on the IPv6 source address within the customer internal network that the client uses for sourcing its tunneled traffic

- The lifetimes of IPv6 and IPv4 addresses are decoupled, allowing for more flexibility in the service provider's addressing policy

- Inactive clients' addresses can be released/reclaimed for allocation to active clients, so more efficient address usage is possible

Since DHCPv4 over IPv4 cannot be used natively in a pure IPv6 network, DHCPv4 over DHCPv6 (DHCP 4o6) [RFC7341] allows DHCPv4 messages to be trasported over a pure IPv6 network by encapsulating DHCPv4 messages into specific DHCPv6 options and messages.

Note that the dynamic provisioning decouples the IPv6 and IPv4 addresses, the binding info required by lwAFTR turns to be an ayschronous combiantion of (restricted) IPv4 address and IPv6 address. [I-D.fsc-softwire-dhcp4o6-saddr-opt] defines a DHCP 4o6 based mechanism for the lwB4 to inform the server of its binding between dynamically allocated IPv4 address and Port Set ID and the IPv6 address that it will use for accessing IPv4-over-IPv6 services

The architecture which is described in this document can be implemented with or without the sharing of IPv4 addresses between multiple clients. If IPv4 address sharing is required, then [RFC7618] describes the necessary extensions to the DHCPv4 server and client provisioning for the allocation and lease management of shared IPv4 addresses.

2. Terminology

Terminology defined in [RFC7341] and [RFC7596] is used extensively throughout this document.

Determinstic provisioning: Lightweight B4 provisioning with DHCPv6 as described in section 5.1 of [RFC7596]. The IPv4 address, restricted port set and the address of lwAFTR are carried in DHCPv6 options defined in [RFC7598].

Dynamic provisioning: Lightweight B4 provisioning with DHCPv4 over DHCPv6 as described in this document. The IPv4 address and rescricted port set are allocated through DHCP 4o6 transport as defined in [RFC7341]. The allocation of lwAFTR's IPv6 address is descirbed in [I-D.fsc-softwire-dhcp4o6-saddr-opt].

3. Dynamic Provisioning Model

As shown in Figure 1, the dynamic provisioning model consists of four functional elements: lwB4, lwAFTR, DHCPv6 Server and DHCP 4o6 Server. Note that these elements are not necessarily separate devices, one or more functional elements could be located on a single device. One existing example of this is the co-location of the DHCP 4o6 Server and lwAFTR as a single gateway device. The differences in the message flow from this co-location are also described below.

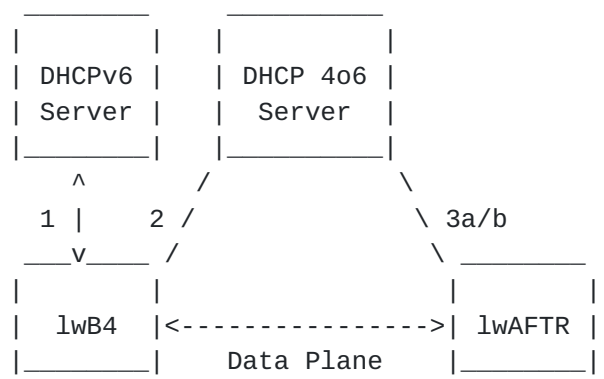

The numbers corresponding to each of the provisioning flows are described in more detail below.

Figure 1: Dynamic lw4o6 Provisioning Model

3.1. Flow 1: lwB4's IPv6 Addressing and DHCPv6 Configuration

Before attempting the DHCP 4o6 configuration process to obtain IPv4 configuration, the lwB4 requires an IPv6 address of a suitable scope to allow communication with the lwAFTR (e.g. a link-local address cannot be used). This IPv6 address can be configured using any applicable method (e.g. SLAAC, DHCPv6, etc.).

To enable DHCPv4 over DHCPv6 transport, the lwB4 needs to perform DHCPv6 to retrieve the DHCP 4o6 server's IPv6 address. The option

code OPTION_DHCP4_O_DHCP6_SERVER (88) is included in the client's ORO. The DHCPv6 server responds the DHCP 4o6 server's IPv6 address by carrying the addresses in DHCP 4o6 Server Address option as defined in [RFC7341].

3.2. Flow 2: DHCP 4o6 Function

Once the lwB4 has acquired the IPv6 address of the DHCP 4o6 server, stateful configuration using DHCP 4o6 is performed to obtain an IPv4 address and (optionally) a port set. The lwB4 sends a DHCPv4 DISCOVER message in a DHCPv4-query Message to the DHCP 4o6 server(s) to activate the DHCP 4o6 transport. To obtain a shared IPv4 address, the lwB4 also has to include Parameter Request List option with the option code OPTION_V4_PORTPARAMS (159) as described in [RFC7618].

To obtain the IPv6 address of lwAFTR and inform the DHCP4o6 server of the lwB4's IPv6 tunnle source address, the message flow described in section 5 of [I-D.fsc-softwire-dhcp4o6-saddr-opt] is followed by the lwB4.

Once successfully completed, the client has been provisioned with the IPv6 address of the lwAFTR, an IPv4 address and (optionally) a range of source ports. The server has the /128 IPv6 address that the client will use its tunnel source associated with the IPv4 lease.

3.3. Flow 3: lwAFTR Binding Table Maintainence

As stated in [RFC7596], the lwAFTR MUST synchronize the binding information with the port-restricted address provisioning process. In the dynamic provisioning model described in this document, once the lwB4's provisioning process is completed and the DHCP 4o6 server holds the client's /128 IPv6 tunnel endpoint address, this binding information can be syncronized with the lwAFTR. The method for this synchronization is dependent on whether the DHCP 4o6 and lwAFTR functions are co-located on the same physical host.

3.3.1. Flow 3a: Binding Table Maintenance for Co-located lwAFTR/DHCP 4o6 Functions

Here, the lwAFTR maintains its binding table as per section 6.1 of [RFC7596] and is synchronized with DHCP 4o6 process. The following DHCP 4o6 messages trigger binding table modification:

DHCPACK: Generated by the DHCP 4o6 server, triggers lwAFTR to add a new entry or modify an existing entry.

DHCPRELEASE: Generated by lwB4, triggers lwAFTR to delete an existing entry.

When the DHCP 4o6 server generates a DHCPACK message, information about the new lease including the client's IPv6 tunnel endpoint address and IPv4 address/PSID tuple is sent to the lwAFTR process. The lwAFTR performs a check that this new binding does not match an existing binding (both the v6 and v4 information must be checked independently to ensure no conflicts). If no conflicts are found the lwAFTR creates a new binding table entry for the client.

If there an existing entry is found, the lwAFTR updates the IPv6 address and lifetime fields of the entry.

When the DHCP 4o6 server receives a DHCPRELEASE message, the lwAFTR looks up the binding table using the lwB4's IPv6 address, IPv4 address and PSID as index. The lwAFTR deletes the entry either by removing it from the binding table or by marking the lifetime field with an invalid value (e.g. 0).

3.3.2. Flow 3b: Binding Table Maintenance for Distributed lwAFTR/DHCP 4o6 Functions

With this architecture, NETCONF [RFC6241] is used for syncronising client DHCP 4o6 provisioning and the lwAFTR binding table. A YANG model for lw4o6 is defined in [I-D.sun-softwire-yang]. In this deployment model, the DHCP 4o6 server and lwAFTR also implements a NETCONF server. When an IPv4 leasing event occurs (DHCPACK/ DHCPRELEASE messages, or an active lease expiring), the DHCP 4o6 server informs the operator's centralised configuration database of the change.

The operator's configuration database will then use NETCONF to update the lwAFTR of the relevant change by adding or removing the binding table entry which matches he DHCP 4o6 server's IPv4 address lease.

4. Security Considerations

Security considerations in [RFC7596] and [RFC7341] are also relevant here.

The DHCP message triggered binding table maintenance may be used by an attacker to send fake DHCP messages to lwAFTR. The operator network should deploy [RFC2827] to prevent this kind of attack.

4.1. Data Retention Requirements

In some countries, regulations require a service providers to retain the necessary information to link IP allocatoin information to a specific customer at a specific point in time.

With a deterministic provisioning model, any individual client will always receive a pre-determined set of IPv4 provisioning requirements. In this scenario, the logging requirement may be met by retaining information on how the DHCPv6 server has been pre-provisioned, with timestamp information on when changes to the pre-provisioning have come into effect.

The dynamic provisioning model that is described in this document brings an additional logging requirement to the service provider: The retention logs holding allocated IPv4 address and ports, the associated IPv6 tunnel endpoint and timestamps marking the start and end of the lease. This is a higher logging overheard than deterministic provisioning, but is in line with the amount of logging that service providers currently have.

5. IANA Considerations

This document does not include an IANA request.

6. References

6.1. Normative References

[I-D.fsc-softwire-dhcp4o6-saddr-opt]
Farrer, I., Sun, Q., and Y. Cui, "DHCPv4 over DHCPv6 Source Address Option", draft-fsc-softwire-dhcp4o6-saddr-opt-04 (work in progress), November 2015.

[RFC2827] Ferguson, P. and D. Senie, "Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing", BCP 38, RFC 2827, DOI 10.17487/RFC2827, May 2000, <http://www.rfc-editor.org/info/rfc2827>.

[RFC7341] Sun, Q., Cui, Y., Siodelski, M., Krishnan, S., and I. Farrer, "DHCPv4-over-DHCPv6 (DHCP 4o6) Transport", RFC 7341, DOI 10.17487/RFC7341, August 2014, <http://www.rfc-editor.org/info/rfc7341>.

[RFC7596] Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture", RFC 7596, DOI 10.17487/RFC7596, July 2015, <http://www.rfc-editor.org/info/rfc7596>.

[RFC7618] Cui, Y., Sun, Q., Farrer, I., Lee, Y., Sun, Q., and M. Boucadair, "Dynamic Allocation of Shared IPv4 Addresses", RFC 7618, DOI 10.17487/RFC7618, August 2015, <http://www.rfc-editor.org/info/rfc7618>.

6.2. Informative References

[I-D.sun-softwire-yang]
Sun, Q., Wang, H., Cui, Y., Farrer, I., Boucadair, M., and R. Asati, "YANG Data Model for IPv4-in-IPv6 Softwire", draft-sun-softwire-yang-04 (work in progress), October 2015.

[RFC3315] Droms, R., Ed., Bound, J., Volz, B., Lemon, T., Perkins, C., and M. Carney, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6)", RFC 3315, DOI 10.17487/RFC3315, July 2003, <http://www.rfc-editor.org/info/rfc3315>.

[RFC6241] Enns, R., Ed., Bjorklund, M., Ed., Schoenwaelder, J., Ed., and A. Bierman, Ed., "Network Configuration Protocol (NETCONF)", RFC 6241, DOI 10.17487/RFC6241, June 2011, <http://www.rfc-editor.org/info/rfc6241>.

[RFC6887] Wing, D., Ed., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", RFC 6887, DOI 10.17487/RFC6887, April 2013, <http://www.rfc-editor.org/info/rfc6887>.

[RFC7598] 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", RFC 7598, DOI 10.17487/RFC7598, July 2015, <http://www.rfc-editor.org/info/rfc7598>.

Authors' Addresses

Cong Liu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5822
Email: cong-liu13@mails.tsinghua.edu.cn

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

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5983
Email: jianping@cernet.edu.cn

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

Email: ian.farrer@telekom.de