Internet Engineering Task Force (IETF)
Request for Comments: 7596
Category: Standards Track
ISSN: 2070-1721

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

July 2015

Lightweight 4over6: An Extension to the Dual-Stack Lite Architecture

Abstract

Dual-Stack Lite (DS-Lite) (RFC 6333) describes an architecture for transporting IPv4 packets over an IPv6 network. This document specifies an extension to DS-Lite called "Lightweight 4over6", which moves the Network Address and Port Translation (NAPT) function from the centralized DS-Lite tunnel concentrator to the tunnel client located in the Customer Premises Equipment (CPE). This removes the requirement for a Carrier Grade NAT function in the tunnel concentrator and reduces the amount of centralized state that must be held to a per-subscriber level. In order to delegate the NAPT function and make IPv4 address sharing possible, port-restricted IPv4 addresses are allocated to the CPEs.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at http://www.rfc-editor.org/info/rfc7596.

Table of Contents

1. Introduction

Dual-Stack Lite (DS-Lite) [RFC6333] defines a model for providing IPv4 access over an IPv6 network using two well-known technologies: IP in IP [RFC2473] and Network Address Translation (NAT). The DS-Lite architecture defines two major functional elements as follows:

Basic Bridging BroadBand (B4) element: A function implemented on a dual-stack-capable node (either a directly connected device or a CPE) that creates an IPv4-in-IPv6 tunnel to an AFTR.

Address Family Transition Router (AFTR) element: The combination of an IPv4-in-IPv6 tunnel endpoint and an IPv4-IPv4 NAT implemented on the same node.

As the AFTR performs the centralized NAT44 function, it dynamically assigns public IPv4 addresses and ports to a requesting host's traffic (as described in [RFC3022]). To achieve this, the AFTR must dynamically maintain per-flow state in the form of active NAPT sessions. For service providers with a large number of B4 clients, the size and associated costs for scaling the AFTR can quickly become prohibitive. Maintaining per-flow state can also place a large NAPT logging overhead on the service provider in countries where logging is a legal requirement.

This document describes a mechanism called "Lightweight 4over6" (lw4o6), which provides a solution for these problems. By relocating the NAPT functionality from the centralized AFTR to the distributed B4s, a number of benefits can be realized:

- NAPT44 functionality is already widely supported and used in today's CPE devices. lw4o6 uses this to provide private<->public NAPT44, meaning that the service provider does not need a centralized NAT44 function.

- The amount of state that must be maintained centrally in the AFTR can be reduced from per-flow to per-subscriber. This reduces the amount of resources (memory and processing power) necessary in the AFTR.

- The reduction of maintained state results in a greatly reduced logging overhead on the service provider.

Operators' IPv6 and IPv4 addressing architectures remain independent of each other. Therefore, flexible IPv4/IPv6 addressing schemes can be deployed.

Lightweight 4over6 is a solution designed specifically for complete independence between IPv6 subnet prefixes and IPv4 addresses with or without IPv4 address sharing. This is accomplished by maintaining state for each softwire (per-subscriber state) in the central lwAFTR and a hub-and-spoke forwarding architecture. "Mapping of Address and Port with Encapsulation (MAP-E)" [RFC7597] also offers these capabilities or, alternatively, allows for a reduction of the amount of centralized state using rules to express IPv4/IPv6 address mappings. This introduces an algorithmic relationship between the IPv6 subnet and IPv4 address. This relationship also allows the option of direct, meshed connectivity between users.

The tunneling mechanism remains the same for DS-Lite and Lightweight 4over6. This document describes the changes to DS-Lite that are necessary to implement Lightweight 4over6. These changes mainly concern the configuration parameters and provisioning method necessary for the functional elements.

One of the features of Lightweight 4over6 is to keep per-subscriber state in the service provider’s network. This technique is categorized as a "binding approach" [Unified-v4-in-v6] that defines a unified IPv4-in-IPv6 softwire CPE.

This document extends the mechanism defined in [RFC7040] by allowing address sharing. The solution in this document is also a variant of Address plus Port (A+P) called "Binding Table Mode" (see Section 4.4 of [RFC6346]).

This document focuses on architectural considerations, particularly on the expected behavior of the involved functional elements and their interfaces. Deployment-specific issues such as redundancy and provisioning policy are out of scope for this document.

2. Conventions

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. Terminology

This document defines the following terms:

Lightweight 4over6 (lw4o6): An IPv4-over-IPv6 hub-and-spoke mechanism that extends DS-Lite by moving the IPv4 translation (NAPT44) function from the AFTR to the B4.

Lightweight B4 (lwB4): A B4 element [RFC6333] that supports Lightweight 4over6 extensions. An lwB4 is a function implemented on a dual-stack-capable node -- either a directly connected device or a CPE -- that supports port-restricted IPv4 address allocation, implements NAPT44 functionality, and creates a tunnel to an lwAFTR.

Lightweight AFTR (lwAFTR): An AFTR element [RFC6333] that supports the Lightweight 4over6 extension. An lwAFTR is an IPv4-in-IPv6 tunnel endpoint that maintains per-subscriber address binding only and does not perform a NAPT44 function.

Restricted port set: A non-overlapping range of allowed external ports allocated to the lwB4 to use for NAPT44. Source ports of IPv4 packets sent by the B4 must belong to the assigned port set. The port set is used for all port-aware IP protocols (TCP, UDP, the Stream Control Transmission Protocol (SCTP), etc.).

Port-restricted IPv4 address: A public IPv4 address with a restricted port set. In Lightweight 4over6, multiple B4s may share the same IPv4 address; however, their port sets must be non-overlapping.

Throughout the remainder of this document, the terms "B4" and "AFTR" should be understood to refer specifically to a DS-Lite implementation. The terms "lwB4" and "lwAFTR" refer to a Lightweight 4over6 implementation.

4. Lightweight 4over6 Architecture

The Lightweight 4over6 architecture is functionally similar to DS-Lite. lwB4s and an lwAFTR are connected through an IPv6-enabled network. Both approaches use an IPv4-in-IPv6 encapsulation scheme to deliver IPv4 connectivity. The following figure shows the data plane with the main functional change between DS-Lite and lw4o6:

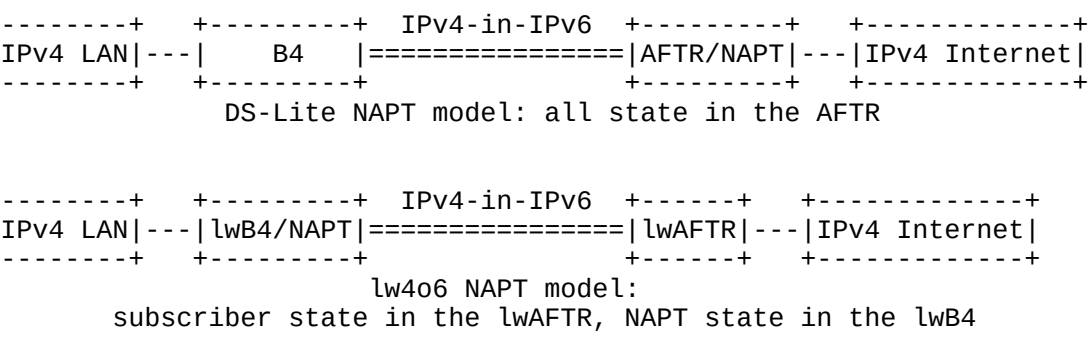

Figure 1: Comparison of DS-Lite and Lightweight 4over6 Data Plane

There are three main components in the Lightweight 4over6 architecture:

- The lwB4, which performs the NAPT function and IPv4/IPv6 encapsulation/decapsulation.
- The lwAFTR, which performs the IPv4/IPv6 encapsulation/ decapsulation.
- The provisioning system, which tells the lwB4 which IPv4 address and port set to use.

The lwB4 differs from a regular B4 in that it now performs the NAPT functionality. This means that it needs to be provisioned with the public IPv4 address and port set it is allowed to use. This information is provided through a provisioning mechanism such as DHCP, the Port Control Protocol (PCP) [RFC6887], or the Broadband Forum's TR-69 specification [TR069].

The lwAFTR needs to know the binding between the IPv6 address of each subscriber as well as the IPv4 address and port set allocated to each subscriber. This information is used to perform ingress filtering upstream and encapsulation downstream. Note that this is per-subscriber state, as opposed to per-flow state in the regular AFTR case.

The consequence of this architecture is that the information maintained by the provisioning mechanism and the one maintained by the lwAFTR MUST be synchronized (see Figure 2). The precise mechanism whereby this synchronization occurs is out of scope for this document.

The solution specified in this document allows the assignment of either a full or a shared IPv4 address to requesting CPEs. [RFC7040] provides a mechanism for assigning a full IPv4 address only.

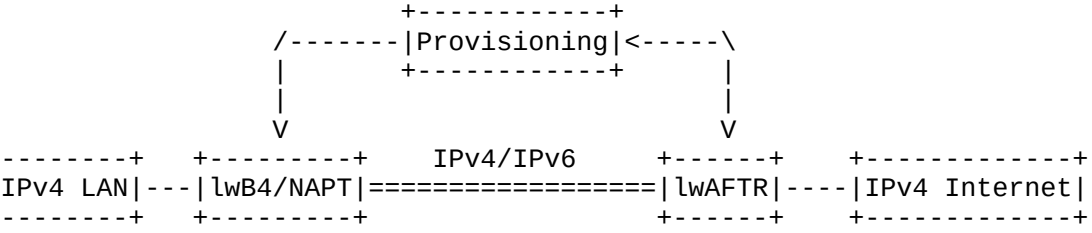

Figure 2: Lightweight 4over6 Provisioning Synchronization

5. Lightweight B4 Behavior

5.1. Lightweight B4 Provisioning with DHCPv6

With DS-Lite, the B4 element only needs to be configured with a single DS-Lite-specific parameter so that it can set up the softwire (the IPv6 address of the AFTR). Its IPv4 address can be taken from the well-known range 192.0.0.0/29.

In lw4o6, a number of lw4o6-specific configuration parameters must be provisioned to the lwB4. These are:

- IPv6 address for the lwAFTR
- IPv4 external (public) address for NAPT44
- Restricted port set to use for NAPT44
- IPv6 binding prefix

The lwB4 MUST implement DHCPv6-based configuration using OPTION_S46_CONT_LW as described in Section 5.3 of [RFC7598]. This means that the lifetime of the softwire and the derived configuration information (e.g., IPv4 shared address, IPv4 address) are bound to the lifetime of the DHCPv6 lease. If stateful IPv4 configuration or additional IPv4 configuration information is required, DHCP 4o6 [RFC7341] MUST be used.

Although it would be possible to extend lw4o6 to have more than one active lw4o6 tunnel configured simultaneously, this document is only concerned with the use of a single tunnel.

The IPv6 binding prefix field is provisioned so that the Customer Edge (CE) can identify the correct prefix to use as the tunnel source. On receipt of the necessary configuration parameters listed above, the lwB4 performs a longest-prefix match between the IPv6 binding prefix and its currently active IPv6 prefixes. The result forms the subnet to be used for sourcing the lw4o6 tunnel. The full /128 address is then constructed in the same manner as [RFC7597].

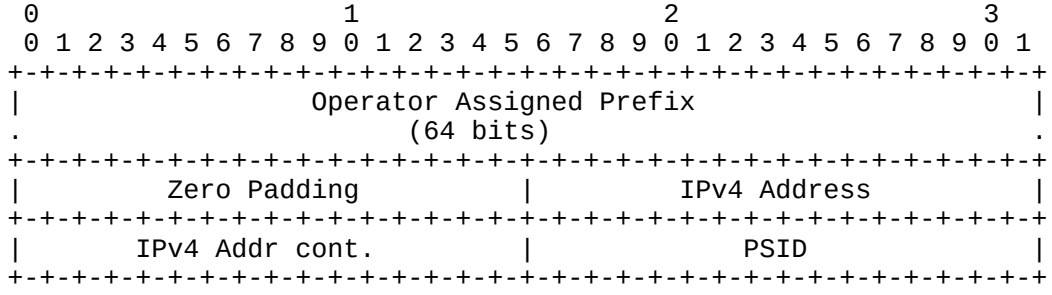

Figure 3: Construction of the lw4o6 /128 Prefix

Operator Assigned Prefix:
IPv6 prefix allocated to the client. If the prefix length is less than 64, it is right-padded with zeros to 64 bits.

Padding: Padding (all zeros).

IPv4 Address: Public IPv4 address allocated to the client.

PSID: Port Set ID. Allocated to the client; left-padded with zeros to 16 bits. If no PSID is provisioned, all zeros.

In the event that the lwB4's IPv6 encapsulation source address is changed for any reason (such as the DHCPv6 lease expiring), the lwB4's dynamic provisioning process MUST be re-initiated. When the lwB4's public IPv4 address or Port Set ID is changed for any reason, the lwB4 MUST flush its NAPT table.

An lwB4 MUST support dynamic port-restricted IPv4 address provisioning. The port-set algorithm for provisioning this is described in Section 5.1 of [RFC7597]. For lw4o6, the number of a-bits SHOULD be 0, thus allocating a single contiguous port set to each lwB4.

Provisioning of the lwB4 using DHCPv6 as described here allocates a single PSID to the client. In the event that the client is concurrently using all of the provisioned L4 ports, it may be unable to initiate any additional outbound connections. DHCPv6-based provisioning does not provide a mechanism for the client to request more L4 port numbers. Other provisioning mechanisms (e.g., PCP-based provisioning [PCP-PORT_SET]) provide this function. Issues relevant to IP address sharing are discussed in more detail in [RFC6269].

Unless an lwB4 is being allocated a full IPv4 address, it is RECOMMENDED that PSIDs containing the system ports (0-1023) not be allocated to lwB4s. The reserved ports are more likely to be reserved by middleware, and therefore we recommend that they not be issued to clients other than as a deliberate assignment. Section 5.2.2 of [RFC6269] provides analysis of allocating system ports to clients with IPv4 address sharing.

In the event that the lwB4 receives an ICMPv6 error message (Type 1, Code 5) originating from the lwAFTR, the lwB4 interprets this to mean that no matching entry in the lwAFTR's binding table has been found, so the IPv4 payload is not being forwarded by the lwAFTR. The lwB4 MAY then re-initiate the dynamic port-restricted provisioning process. The lwB4's re-initiation policy SHOULD be configurable.

On receipt of such an ICMP error message, the lwB4 MUST validate the source address to be the same as the lwAFTR address that is configured. In the event that these addresses do not match, the lwB4 MUST discard the ICMP error message.

In order to prevent forged ICMP messages (using the spoofed lwAFTR address as the source) from being sent to lwB4s, the operator can implement network ingress filtering as described in [RFC2827].

The DNS considerations described in Sections 5.5 and 6.4 of [RFC6333] apply to Lightweight 4over6; lw4o6 implementations MUST comply with all requirements stated there.

5.2. Lightweight B4 Data-Plane Behavior

Several sections of [RFC6333] provide background information on the B4's data-plane functionality and MUST be implemented by the lwB4, as they are common to both solutions. The relevant sections are:

5.2 Encapsulation	Covering encapsulation and decapsulation of tunneled traffic
5.3 Fragmentation and Reassembly	Covering MTU and fragmentation considerations (referencing [RFC2473])
7.1 Tunneling	Covering tunneling and Traffic Class mapping between IPv4 and IPv6 (referencing [RFC2473]). Also see [RFC2983]

The lwB4 element performs IPv4 address translation (NAPT44) as well as encapsulation and decapsulation. It runs standard NAPT44 [RFC3022] using the allocated port-restricted address as its external IPv4 address and range of source ports.

The working flow of the lwB4 is illustrated in Figure 4.

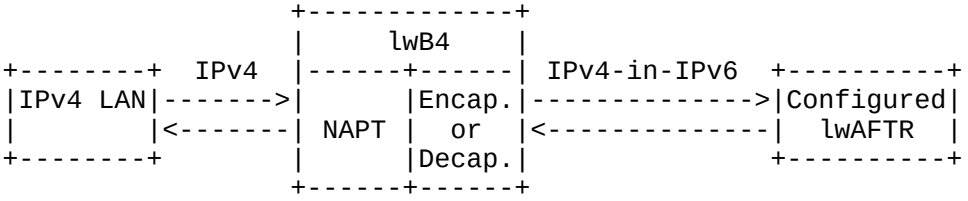

Figure 4: Working Flow of the lwB4

Hosts connected to the customer's network behind the lwB4 source IPv4 packets with an [RFC1918] address. When the lwB4 receives such an IPv4 packet, it performs a NAPT44 function on the source address and port by using the public IPv4 address and a port number from the allocated port set. Then, it encapsulates the packet with an IPv6 header. The destination IPv6 address is the lwAFTR's IPv6 address, and the source IPv6 address is the lwB4's IPv6 tunnel endpoint address. Finally, the lwB4 forwards the encapsulated packet to the configured lwAFTR.

When the lwB4 receives an IPv4-in-IPv6 packet from the lwAFTR, it decapsulates the IPv4 packet from the IPv6 packet. Then, it performs NAPT44 translation on the destination address and port, based on the available information in its local NAPT44 table.

If the IPv6 source address does not match the configured lwAFTR address, then the packet MUST be discarded. If the decapsulated IPv4 packet does not match the lwB4's configuration (i.e., invalid destination IPv4 address or port), then the packet MUST be dropped. An ICMPv4 error message (Type 3, Code 13 -- Destination Unreachable, Communication Administratively Prohibited) MAY be sent back to the lwAFTR. The ICMP policy SHOULD be configurable.

The lwB4 is responsible for performing Application Layer Gateway (ALG) functions (e.g., SIP, FTP) and other NAPT traversal mechanisms (e.g., Universal Plug and Play (UPnP) IGD (Internet Gateway Device), the NAT Port Mapping Protocol (NAT-PMP), manual binding configuration, PCP) for the internal hosts, if necessary. This requirement is typical for NAPT44 gateways available today.

It is possible that an lwB4 is co-located in a host. In this case, the functions of NAPT44 and encapsulation/decapsulation are implemented inside the host.

5.2.1. Fragmentation Behavior

For TCP and UDP traffic, the NAPT44 implemented in the lwB4 MUST conform to the behavior and best current practices documented in [RFC4787], [RFC5508], and [RFC5382]. If the lwB4 supports the Datagram Congestion Control Protocol (DCCP), then the requirements in [RFC5597] MUST be implemented.

The NAPT44 in the lwB4 MUST implement ICMP message handling behavior conforming to the best current practice documented in [RFC5508]. If the lwB4 receives an ICMP error (for errors detected inside the IPv6 tunnel), the node relays the ICMP error message to the original source (the lwAFTR). This behavior SHOULD be implemented conforming to Section 8 of [RFC2473].

If IPv4 hosts behind different lwB4s sharing the same IPv4 address send fragments to the same IPv4 destination host outside the Lightweight 4over6 domain, those hosts may use the same IPv4 fragmentation identifier, resulting in incorrect reassembly of the fragments at the destination host. Given that the IPv4 fragmentation identifier is a 16-bit field, it could be used similarly to port ranges: An lwB4 could rewrite the IPv4 fragmentation identifier to be within its allocated port set, if the resulting fragment identifier space is large enough related to the rate at which fragments are

sent. However, splitting the identifier space in this fashion would increase the probability of reassembly collision for all connections through the lwB4. See also Section 5.3.1 of [RFC6864].

6. Lightweight AFTR Behavior

6.1. Binding Table Maintenance

The lwAFTR maintains an address binding table containing the binding between the lwB4's IPv6 address, the allocated IPv4 address, and the restricted port set. Unlike the DS-Lite extended binding table, which is a 5-tuple NAPT table and is defined in Section 6.6 of [RFC6333], each entry in the Lightweight 4over6 binding table contains the following 3-tuples:

- IPv6 address for a single lwB4
- Public IPv4 address
- Restricted port set

The entry has two functions: the IPv6 encapsulation of inbound IPv4 packets destined to the lwB4 and the validation of outbound IPv4-in-IPv6 packets received from the lwB4 for decapsulation.

The lwAFTR does not perform NAPT and so does not need session entries.

The lwAFTR MUST synchronize the binding information with the port-restricted address provisioning process. If the lwAFTR does not participate in the port-restricted address provisioning process, the binding MUST be synchronized through other methods (e.g., out-of-band static update).

If the lwAFTR participates in the port-restricted provisioning process, then its binding table MUST be created as part of this process.

For all provisioning processes, the lifetime of binding table entries MUST be synchronized with the lifetime of address allocations.

6.2. lwAFTR Data-Plane Behavior

Several sections of [RFC6333] provide background information on the AFTR's data-plane functionality and MUST be implemented by the lwAFTR, as they are common to both solutions. The relevant sections are:

6.2 Encapsulation	Covering encapsulation and decapsulation of tunneled traffic
6.3 Fragmentation and Reassembly	Fragmentation and reassembly considerations (referencing [RFC2473])
7.1 Tunneling	Covering tunneling and Traffic Class mapping between IPv4 and IPv6 (referencing [RFC2473]). Also see [RFC2983]

When the lwAFTR receives an IPv4-in-IPv6 packet from an lwB4, it decapsulates the IPv6 header and verifies the source addresses and port in the binding table. If both the source IPv4 and IPv6 addresses match a single entry in the binding table and the source port is in the allowed port set for that entry, the lwAFTR forwards the packet to the IPv4 destination.

If no match is found (e.g., no matching IPv4 address entry, port out of range), the lwAFTR MUST discard or implement a policy (such as redirection) on the packet. An ICMPv6 Type 1, Code 5 (Destination Unreachable, source address failed ingress/egress policy) error message MAY be sent back to the requesting lwB4. The ICMP policy SHOULD be configurable.

When the lwAFTR receives an inbound IPv4 packet, it uses the IPv4 destination address and port to look up the destination lwB4's IPv6 address in its binding table. If a match is found, the lwAFTR encapsulates the IPv4 packet. The source is the lwAFTR's IPv6 address, and the destination is the lwB4's IPv6 address from the matched entry. Then, the lwAFTR forwards the packet to the lwB4 natively over the IPv6 network.

If no match is found, the lwAFTR MUST discard the packet. An ICMPv4 Type 3, Code 1 (Destination Unreachable, Host Unreachable) error message MAY be sent back. The ICMP policy SHOULD be configurable.

The lwAFTR MUST support hairpinning of traffic between two lwB4s, by performing decapsulation and re-encapsulation of packets from one lwB4 that need to be sent to another lwB4 associated with the same AFTR. The hairpinning policy MUST be configurable.

7. Additional IPv4 Address and Port-Set Provisioning Mechanisms

In addition to the DHCPv6-based mechanism described in Section 5.1, several other IPv4 provisioning protocols have been suggested. These protocols MAY be implemented. These alternatives include:

- DHCPv4 over DHCPv6: [RFC7341] describes implementing DHCPv4 messages over an IPv6-only service provider's network. This enables leasing of IPv4 addresses and makes DHCPv4 options available to the DHCPv4-over-DHCPv6 client. An lwB4 MAY implement [RFC7341] and [Dyn-Shared-v4Alloc] to retrieve a shared IPv4 address with a set of ports.

- PCP [RFC6887]: an lwB4 MAY use [PCP-PORT_SET] to retrieve a restricted IPv4 address and a set of ports.

In a Lightweight 4over6 domain, the binding information MUST be synchronized across the lwB4s, the lwAFTRs, and the provisioning server.

To prevent interworking complexity, it is RECOMMENDED that an operator use a single provisioning mechanism / protocol for their implementation. In the event that more than one provisioning mechanism / protocol needs to be used (for example, during a migration to a new provisioning mechanism), the operator SHOULD ensure that each provisioning mechanism has a discrete set of resources (e.g., IPv4 address/PSID pools, as well as lwAFTR tunnel addresses and binding tables).

8. ICMP Processing

For both the lwAFTR and the lwB4, ICMPv6 MUST be handled as described in [RFC2473].

ICMPv4 does not work in an address-sharing environment without special handling [RFC6269]. Due to the port-set style of address sharing, Lightweight 4over6 requires specific ICMP message handling not required by DS-Lite.

8.1. ICMPv4 Processing by the lwAFTR

For inbound ICMP messages, the following behavior SHOULD be implemented by the lwAFTR to provide ICMP error handling and basic remote IPv4 service diagnostics for a port-restricted CPE:

1. Check the ICMP Type field.

2. If the ICMP Type field is set to 0 or 8 (echo reply or request), then the lwAFTR MUST take the value of the ICMP Identifier field as the source port and use this value to look up the binding table for an encapsulation destination. If a match is found, the lwAFTR forwards the ICMP packet to the IPv6 address stored in the entry; otherwise, it MUST discard the packet.

3. If the ICMP Type field is set to any other value, then the lwAFTR MUST use the method described in REQ-3 of [RFC5508] to locate the source port within the transport-layer header in the ICMP packet's data field. The destination IPv4 address and source port extracted from the ICMP packet are then used to make a lookup in the binding table. If a match is found, it MUST forward the ICMP reply packet to the IPv6 address stored in the entry; otherwise, it MUST discard the packet.

Otherwise, the lwAFTR MUST discard all inbound ICMPv4 messages.

The ICMP policy SHOULD be configurable.

8.2. ICMPv4 Processing by the lwB4

The lwB4 MUST implement the requirements defined in [RFC5508] for ICMP forwarding. For ICMP echo request packets originating from the private IPv4 network, the lwB4 SHOULD implement the method described in [RFC6346] and use an available port from its port set as the ICMP identifier.

9. Security Considerations

As the port space for a subscriber shrinks due to address sharing, the randomness for the port numbers of the subscriber is decreased significantly. This means that it is much easier for an attacker to guess the port number used, which could result in attacks ranging from throughput reduction to broken connections or data corruption.

The port set for a subscriber can be a set of contiguous ports or non-contiguous ports. Contiguous port sets do not reduce this threat. However, with non-contiguous port sets (which may be generated in a pseudorandom way [RFC6431]), the randomness of the

port number is improved, provided that the attacker is outside the Lightweight 4over6 domain and hence does not know the port-set generation algorithm.

The lwAFTR MUST rate-limit ICMPv6 error messages (see Section 5.1) to defend against DoS attacks generated by an abuse user.

More considerations about IP address sharing are discussed in Section 13 of [RFC6269], which is applicable to this solution.

This document describes a number of different protocols that may be used for the provisioning of lw4o6. In each case, the security considerations relevant to the provisioning protocol are also relevant to the provisioning of lw4o6 using that protocol. lw4o6 does not add any other security considerations specific to these provisioning protocols.

10. References

10.1. Normative References

[RFC1918] Rekhter, Y., Moskowitz, B., Karrenberg, D., de Groot, G., and E. Lear, "Address Allocation for Private Internets", BCP 5, RFC 1918, DOI 10.17487/RFC1918, February 1996, <http://www.rfc-editor.org/info/rfc1918>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC2473] Conta, A. and S. Deering, "Generic Packet Tunneling in IPv6 Specification", RFC 2473, DOI 10.17487/RFC2473, December 1998, <http://www.rfc-editor.org/info/rfc2473>.

[RFC4787] Audet, F., Ed., and C. Jennings, "Network Address Translation (NAT) Behavioral Requirements for Unicast UDP", BCP 127, RFC 4787, DOI 10.17487/RFC4787, January 2007, <http://www.rfc-editor.org/info/rfc4787>.

[RFC5382] Guha, S., Ed., Biswas, K., Ford, B., Sivakumar, S., and P. Srisuresh, "NAT Behavioral Requirements for TCP", BCP 142, RFC 5382, DOI 10.17487/RFC5382, October 2008, <http://www.rfc-editor.org/info/rfc5382>.

[RFC5508] Srisuresh, P., Ford, B., Sivakumar, S., and S. Guha, "NAT Behavioral Requirements for ICMP", BCP 148, RFC 5508, DOI 10.17487/RFC5508, April 2009, <http://www.rfc-editor.org/info/rfc5508>.

[RFC5597] Denis-Courmont, R., "Network Address Translation (NAT) Behavioral Requirements for the Datagram Congestion Control Protocol", BCP 150, RFC 5597, DOI 10.17487/RFC5597, September 2009, <http://www.rfc-editor.org/info/rfc5597>.

[RFC6333] Durand, A., Droms, R., Woodyatt, J., and Y. Lee, "Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion", RFC 6333, DOI 10.17487/RFC6333, August 2011, <http://www.rfc-editor.org/info/rfc6333>.

[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>.

10.2. Informative References

[B4-Trans-DSLite]
Cui, Y., Sun, Q., Boucadair, M., Tsou, T., Lee, Y., and I. Farrer, "Lightweight 4over6: An Extension to the DS-Lite Architecture", Work in Progress, draft-cui-softwire-b4-translated-ds-lite-11, February 2013.

[DSLite-LW-Ext]
Deng, X., Boucadair, M., and C. Zhou, "NAT offload extension to Dual-Stack lite", Work in Progress, draft-zhou-softwire-b4-nat-04, October 2011.

[Dyn-Shared-v4Alloc]
Cui, Y., Sun, Q., Farrer, I., Lee, Y., Sun, Q., and M. Boucadair, "Dynamic Allocation of Shared IPv4 Addresses", Work in Progress, draft-ietf-dhc-dynamic-shared-v4allocation-09, May 2015.

[PCP-PORT_SET]
Sun, Q., Boucadair, M., Sivakumar, S., Zhou, C., Tsou, T., and S. Perreault, "Port Control Protocol (PCP) Extension for Port Set Allocation", Work in Progress, draft-ietf-pcp-port-set-09, May 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>.

[RFC2983] Black, D., "Differentiated Services and Tunnels", RFC 2983, DOI 10.17487/RFC2983, October 2000, <http://www.rfc-editor.org/info/rfc2983>.

[RFC3022] Srisuresh, P. and K. Egevang, "Traditional IP Network Address Translator (Traditional NAT)", RFC 3022, DOI 10.17487/RFC3022, January 2001, <http://www.rfc-editor.org/info/rfc3022>.

[RFC6269] Ford, M., Ed., Boucadair, M., Durand, A., Levis, P., and P. Roberts, "Issues with IP Address Sharing", RFC 6269, DOI 10.17487/RFC6269, June 2011, <http://www.rfc-editor.org/info/rfc6269>.

[RFC6346] Bush, R., Ed., "The Address plus Port (A+P) Approach to the IPv4 Address Shortage", RFC 6346, DOI 10.17487/RFC6346, August 2011, <http://www.rfc-editor.org/info/rfc6346>.

[RFC6431] Boucadair, M., Levis, P., Bajko, G., Savolainen, T., and T. Tsou, "Huawei Port Range Configuration Options for PPP IP Control Protocol (IPCP)", RFC 6431, DOI 10.17487/RFC6431, November 2011, <http://www.rfc-editor.org/info/rfc6431>.

[RFC6864] Touch, J., "Updated Specification of the IPv4 ID Field", RFC 6864, DOI 10.17487/RFC6864, February 2013, <http://www.rfc-editor.org/info/rfc6864>.

[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>.

[RFC7040] Cui, Y., Wu, J., Wu, P., Vautrin, O., and Y. Lee, "Public IPv4-over-IPv6 Access Network", RFC 7040, DOI 10.17487/RFC7040, November 2013, <http://www.rfc-editor.org/info/rfc7040>.

[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>.

[RFC7597] Troan, O., Ed., Dec, W., Li, X., Bao, C., Matsushima, S., Murakami, T., and T. Taylor, Ed., "Mapping of Address and Port with Encapsulation (MAP-E)", RFC 7597, DOI 10.17487/RFC7597, July 2015, <http://www.rfc-editor.org/info/rfc7597>.

[Stateless-DS-Lite] Penno, R., Durand, A., Clauberg, A., and L. Hoffmann, "Stateless DS-Lite", Work in Progress, draft-penno-softwire-sdnat-02, March 2012.

[TR069] Broadband Forum TR-069, "CPE WAN Management Protocol", Amendment 5, CWMP Version: 1.4, November 2013, <https://www.broadband-forum.org>.

[Unified-v4-in-v6] Boucadair, M., Farrer, I., Perreault, S., Ed., and S. Sivakumar, Ed., "Unified IPv4-in-IPv6 Softwire CPE", Work in Progress, draft-ietf-softwire-unified-cpe-01, May 2013.

Acknowledgements

The authors would like to thank Ole Troan, Ralph Droms, and Suresh Krishnan for their comments and feedback.

This document is a merge of three documents: [B4-Trans-DSLite], [DSLite-LW-Ext], and [Stateless-DS-Lite].

Contributors

The following individuals contributed to this effort:

Jianping Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
China
Phone: +86-10-62785983
Email: jianping@cernet.edu.cn

Peng Wu
Tsinghua University
Department of Computer Science, Tsinghua University
Beijing 100084
China
Phone: +86-10-62785822
Email: pengwu.thu@gmail.com

Qi Sun
Tsinghua University
Beijing 100084
China
Phone: +86-10-62785822
Email: sunqi@csnet1.cs.tsinghua.edu.cn

Chongfeng Xie
China Telecom
Room 708, No. 118, Xizhimennei Street
Beijing 100035
China
Phone: +86-10-58552116
Email: xiechf@ctbri.com.cn

Xiaohong Deng
The University of New South Wales
Sydney NSW 2052
Australia
Email: dxhbupt@gmail.com

Cathy Zhou
Huawei Technologies
Section B, Huawei Industrial Base, Bantian Longgang
Shenzhen 518129
China
Email: cathyzhou@huawei.com

Alain Durand
Juniper Networks
1194 North Mathilda Avenue
Sunnyvale, CA 94089-1206
United States
Email: adurand@juniper.net

Reinaldo Penno
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134
United States
Email: repenno@cisco.com

Axel Clauberg
Deutsche Telekom AG
CTO-ATI
Landgrabenweg 151
Bonn 53227
Germany
Email: axel.clauberg@telekom.de

Lionel Hoffmann
Bouygues Telecom
TECHNOPOLE
13/15 Avenue du Marechal Juin
Meudon 92360
France
Email: lhoffman@bouyguestelecom.fr

Maoke Chen (a.k.a. Noriyuki Arai)
BBIX, Inc.
Tokyo Shiodome Building, Higashi-Shimbashi 1-9-1
Minato-ku, Tokyo 105-7310
Japan
Email: maoke@bbix.net

Authors' Addresses

Yong Cui
Tsinghua University
Beijing 100084
China

Phone: +86-10-62603059
Email: yong@csnet1.cs.tsinghua.edu.cn

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

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

Mohamed Boucadair
France Telecom
Rennes 35000
France

Email: mohamed.boucadair@orange.com

Tina Tsou
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050
United States

Phone: +1-408-330-4424
Email: tena@huawei.com

Yiu L. Lee
Comcast
One Comcast Center
Philadelphia, PA 19103
United States

Email: yiu_lee@cable.comcast.com

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

Email: ian.farrer@telekom.de