

Network Working Group
Internet-Draft
Intended status: Informational
Expires: January 4, 2018

Q. Sun
C. Xie
China Telecom
Y. Lee
Comcast
M. Chen
FreeBit
T. Li
Tsinghua University
I. Farrer
Deutsche Telekom AG
July 3, 2017

Deployment Considerations for Lightweight 4over6
draft-ietf-softwire-lightweight-4over6-deployment-01

Abstract

Lightweight 4over6 is a mechanism for providing IPv4 services to clients connected to a single-stack IPv6 network. The architecture is similar to DS-Lite, but the network address translation function is relocated from the tunnel concentrator to the tunnel client, hence reducing the amount of state which must be maintained in the concentrator to a per-customer level. This document discusses the applicability, describes various deployment models and provides deployment considerations for Lightweight 4over6.

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 January 4, 2018.

Copyright Notice

Copyright (c) 2017 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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	3
2.	Deployment Models	4
2.1.	Top-Down Deployment Model	4
2.2.	Bottom-Up Deployment Model	5
2.3.	Campus Deployment	5
3.	Overall Deployment Considerations	5
3.1.	IP Addressing and Routing	5
3.1.1.	IPv4 Routing	5
3.1.2.	IPv6 Routing	6
3.2.	lwB4 Configuration	6
3.2.1.	DHCPv6 Based Provisioning	6
3.2.2.	DHCPv4 over DHCPv6 Based Provisioning	7
3.2.3.	PCP Based Provisioning	7
3.2.4.	NETCONF/YANG Based Provisioning	7
3.2.5.	Other lwB4 Configuration Considerations	7
3.3.	lwAFTR Discovery	8
3.4.	Impacts on Accounting	8
4.	lwAFTR Deployment Considerations	8
4.1.	Logging at the lwAFTR	9
4.2.	MTU and Fragmentation Considerations	9
4.3.	Reliability Considerations of lwAFTR	9
4.4.	Location of lwAFTRs in the Network	10
4.5.	Path Consistency Consideration	10
5.	lwB4 Deployment Considerations	11
5.1.	NAT Traversal Issues	11
5.2.	Static Port Forwarding Configuration	11
6.	DS-Lite Compatibility Consideration	12
6.1.	Case 1: Integrated Network Element with Lightweight 4over6 and DS-Lite AFTR Scenario	12
6.2.	Case 2: DS-Lite Coexistent scenario with Separated AFTR .	13
7.	Acknowledgements	14

8. References	14
Appendix A. China Telecom Experimental Results	17
A.1. Experimental Environment	18
A.2. Experimental Results	19
A.3. Conclusions	20
Appendix B. Tsinghua University Experimental Result	20
B.1. Experimental Environment	20
B.2. Experimental Results	21
B.3. Conclusion	22
Authors' Addresses	22

[1. Introduction](#)

Lightweight 4over6 [[RFC7596](#)] (lw4o6) is an extension to DS-Lite [[RFC6333](#)] which simplifies the AFTR module by relocating the NAPT function among B4 elements located at the subscriber's premises. In the lw4o6 architecture, the functional elements are referred to as the lwB4 and lwAFTR.

The lwB4 is provisioned with an IPv6 address, a public IPv4 address and a port-set. It performs port-restricted NAPT on subscriber's packets using the provisioned public IPv4 address and port-set. IPv4 packets are routed between the lwB4 and the lwAFTR encapsulated in an IPv4 in IPv6 Software. The lwAFTR maintains one binding entry per-subscriber, consisting of the lwB4's IPv6 tunnel endpoint, IPv4 address and port-set. [Section 4.4 of \[RFC6346\]](#) provides more detail of this mechanism.

This can bring a number of advantages when compared to DS-Lite:

- o Per-subscriber configuration allows for the operator to provision each subscriber according to their specific service requirements.
- o The logging requirements to meet regulatory requirements may be reduced as it is only necessary to log when a subscriber is provisioned or de-provisioned in the lwAFTR. This relaxes the need for logging on a per-session, or per port block allocation.
- o In some lw4o6 deployment topologies, the removal of per-session state means that it is possible to have very large parallelisation of lwAFTR elements. This offers excellent scaling and resilience.
- o This mechanism preserves the dynamic feature of IPv4/IPv6 address binding as in DS-Lite, so it has no coupling between IPv6 address and IPv4 address/port-set as any full stateless solution ([\[RFC6052\]](#) or [\[RFC7597\]](#)) requires.

The terminology used in this document follows the definitions and

abbreviations from [\[RFC6333\]](#) and [\[RFC7596\]](#).

2. Deployment Models

Lightweight 4over6 is suitable for operators who would like to free any correlation of the IPv6 address with IPv4 address and port-set, as the IPv4 addressing is an overlay to the IPv6 addressing architecture. Thus, IPv6 addressing is completely flexible to fit other deployment requirements, e.g., auto-configuration, service classification, user management, QoS support, etc.

Lightweight 4over6 can be deployed in a residential ISP network (depicted in Figure 1). In this scenario, a lwB4 acquires an IPv4 address and a port-set alongside IPv6 provisioning, including an address for the lwAFTR. It then establishes an IPv4-in-IPv6 software to the lwAFTR. The lwB4 function may be implemented in a CPE providing IPv4 services for a home network, or directly in a host.

The lwAFTR holds a table with the bindings between the lwB4's IPv6 addresses and their allocated IPv4 addresses + port sets. The supporting system is used to synchronise the lwB4 and lwAFTR binding information. It may also be used for logging and user management.




Figure 1 Architectural Overview

2.1. Top-Down Deployment Model

In the top-down deployment model, the supporting system holds the overall binding table for the network. It uses this to pre-provision the local binding table entries for the lwAFTR and also provision lwB4s with the correct configuration (e.g. using DHCPv6 or PCP).

With this method, one binding table entry can be present on lwAFTRs and stateless failover can be achieved.

2.2. Bottom-Up Deployment Model

In the bottom-up model, the client is first provisioned with the relevant parameters necessary for building the software. It then attempts to send traffic to the lwAFTR.

On receipt of lwB4 traffic which does not have an existing binding-table entry, one is dynamically created. The lwAFTR reports the new binding entry to the supporting system.

[[I-D.ietf-behave-syslog-nat-logging](#)] or [[I-D.ietf-behave-ipfix-nat-logging](#)] may be used for this purpose. In this way, the lwAFTR can determine the binding by its own and there is little impact on existing network architecture.

2.3. Campus Deployment

Lightweight 4over6 can also be deployed in a campus or enterprise network, (depicted in Figure 2). In this scenario, a lwB4 acts as a gateway router for a number of hosts. The lwB4 is first provisioned with an IPv4 address and port-set. It then establishes an IPv4-in-IPv6 software using the IPv6 address to deliver IPv4 services to its connected host via the lwAFTR in the network. A network management system could be used to receive statistic information of the network equipments, such as the binding table, network load, and connected device. NETCONF [[RFC6241](#)] could be used for synchronising lwB4's IPv6 address and its allocated IPv4 address + port set with the lwAFTR. The network management system may keep the binding information as well for logging and user management.

3. Overall Deployment Considerations

3.1. IP Addressing and Routing

In Lightweight 4over6, there is no inter-dependency between the IPv4 and IPv6 addressing schemes. This allows for complete flexibility in addressing architecture.

3.1.1. IPv4 Routing

The IPv4 addresses/prefixes that are allocated to customer's lwB4s are advertised to the IPv4 Internet as being reachable via the lwAFTR(s). If multiple lwAFTRs are all serving the same set of lwB4s, all will advertise the same IPv4 reachable routes.

3.1.2. IPv6 Routing

The lwAFTR provides a /128 IPv6 tunnel endpoint address which is advertised to the lwB4s. If multiple lwAFTRs are all serving the same set of lwB4s, all will advertise the same IPv6 tunnel endpoint address.

The lwB4's IPv6 addressing and routing, there are no specific topological limitations. An existing IPv6 address and routing architecture should not be affected. For example, in PPPoE scenario, a CPE could obtain a prefix via DHCPv6 prefix delegation, and the hosts behind CPE would get its own IPv6 addresses within the prefix through SLAAC or DHCPv6 statefully. This IPv6 address assignment procedure has nothing to do with restricted IPv4 address allocation.

It is worth noting that if the Top-Down provisioning model is chosen, then there must be determinism in the local address that the lwB4 uses for building its tunnel. This is so that the binding entry for the lwB4 can be pre-provisioned in the lwAFTR. [RFC7598] offers a solution for this using the 'bind-ipv6-prefix' field is used to inform the lwB4 which configured prefix to use. The suffix is then created according to [Section 6 of \[RFC7597\]](#).

3.2. lwB4 Configuration

In lw4o6, each lwB4 will get its restricted IPv4 address and a port-set after successful user authentication process and IPv6 provisioning process. This assignment can be achieved using a number of methods:

- o DHCPv6 Softwire S46 Option [[RFC7598](#)]
- o DHCPv4 over DHCPv6 [[RFC7341](#)], [[RFC7618](#)] and [[I-D.ietf-dhc-dhcp4o6-saddr-opt](#)]
- o PCP [[RFC7753](#)]
- o NETCONF/YANG [[I-D.ietf-softwire-yang](#)]

3.2.1. DHCPv6 Based Provisioning

[RFC7598] describes a set of DHCPv6 options used for provisioning lw4o6 clients. OPTION_S46_CONT_LW (96) is a DHCPv6 container option, which can hold the IPv6 of the lwAFTR (OPTION_S46_BR (90)), the lwB4's IPv4 address and IPv6 prefix hint (OPTION_S46_V4V6BIND (92)), and port set information (OPTION_S46_PORTPARAMS (93)).

In this model, the DHCPv6 server needs to be pre-provisioned with the

client configuration. Therefore, this approach is better suited to client configurations that will be long-lived.

DHCPv6 based provisioning can also be used in conjunction with a AAA server. [[I-D.ietf-softwire-map-radius](#)] describes this function.

3.2.2. DHCPv4 over DHCPv6 Based Provisioning

An operator may use DHCPv4 to provision IPv4 address to the lwB4. In a typical deployment, the DHCP server is a centralized DHCP server and lwAFTR is the DHCP relay agent to relay the dhcp messages to the server over unicast. Rarely DHCP server will collocate with the lwAFTR to provision IPv4 resources to the lwB4.

3.2.3. PCP Based Provisioning

Operator may also use PCP Port-set Option to provision IPv4 address and port-set to the lwB4. In a typical deployment, PCP server will collocate with lwAFTR, and the subscriber's binding can be determined by lwAFTR. The PCP request should be sent to the lwAFTR's tunnel end-point address. It is not common that PCP server will be centralized deployed in which the lwAFTR is the PCP proxy to relay PCP requests.

3.2.4. NETCONF/YANG Based Provisioning

Operators using NETCONF to manage customer devices can provision lw4o6 using [[I-D.ietf-softwire-yang](#)].

3.2.5. Other lwB4 Configuration Considerations

Some operators may offer different service level agreements (SLA) to users that some users may require more ports than others. In this deployment scenario, the operator can implement differentiated policies in provisioning system specified to a user's lwB4 or a group of lwB4s to allocate a certain range of port-set. The lwAFTR may also run multiple instances with different port-set sizes to build the mapping table.

It is also worth noting the compatibility between lw4o6 and Public IPv4 over IPv6 [[RFC7040](#)]. When a lw4o6 client is provisioned with a 'full' IPv4 address (i.e. with no port-set or a port-set that allows the use of all of the L4 ports), then the A+P routing model is no longer used by the lwAFTR as traffic is routed on the IPv4 address only. This function can be useful when a subscriber uses protocols which do not have L4 ports.

3.3. lwAFTR Discovery

The lwB4 needs to discover the lwAFTR's IPv6 address before it is able to set up the software tunnel and provide any IPv4 services. This address can be learned through an out-of-band channel, static configuration, or dynamic configuration. In practice, Lightweight 4over6 lwB4 can use the same DHCPv6 option [[RFC6334](#)] to discover the FQDN of the lwAFTR.

When Lightweight 4over6 is deployed in the same place with DS-Lite, either different FQDNs can be configured for Lightweight 4over6 and DS-Lite separately or different DHCPv6 options can be used for Lightweight 4over6 [[RFC7598](#)] and DS-Lite. More detailed considerations on DS-Lite compatibility will be discussed in [Section 6](#).

The lw4o6 DHCPv6 option (OPTION_S46_LW_CONT (96)) can contain OPTION_S46_BR (90) which holds the v6 address of the lwAFTR.

3.4. Impacts on Accounting

In lw4o6, the accounting impact due to the tunneling protocol is the same with DS-Lite (see [section 6.2 of \[RFC6908\]](#)). However, since in lw4o6, the IPv4 service is only available after port-set allocation, if operators regard IPv4 service as a on-demand value-added service, e.g. IPv6 connectivity is offered by default, while IPv4 connectivity will be offered until a subscriber requires, etc., IPv4 service accounting should start after port-set allocation has completed.

It should be noted that in common with all A+P mechanisms, lw4o6 can not performing per-session logging in the way that CGN based solutions do.

4. lwAFTR Deployment Considerations

As Lightweight 4over6 is an extension to DS-Lite, both technologies share similar deployment considerations. For example: the interface considerations, lawful intercept considerations, blacklisting a shared IPv4 Address, AFTR policies, and impacts on accounting processes described in [[RFC6908](#)] are also applicable here. This document only discusses additional considerations specific to Lightweight 4over6.

4.1. Logging at the lwAFTR

In lw4o6, operators only log one entry per subscriber. Each log needs to include subscriber's IPv6 address used for the software, the public IPv4 address and the allocated port-set, and the start and end times that the binding entry was valid for.

To ensure consistency of the logged information, the port set algorithm implemented in lw4o6 lwAFTR needs to be synchronized with the one implemented in the logging system. For example, if contiguous port-set algorithm is adopted in the lwAFTR, the same algorithm needs to be applied for the logging system.

Since the binding in lwAFTR does not log sessions as they are set up, operators should be aware that lw4o6 does not provide a mechanism for destination-specific logging.

4.2. MTU and Fragmentation Considerations

As Lightweight 4over6 uses a tunneling protocol, the same considerations regarding fragmentation and reassembly as for DS-Lite [[RFC6908](#)] are applicable. In order to avoid the problems that are associated with fragmentation, it is advisable to ensure that the MTU across the IPv6 domain between the lwB4 and lwAFTR is large enough to allow for the transportation of IPv4 packets plus the 40-byte overhead for IPv6 encapsulation.

4.3. Reliability Considerations of lwAFTR

Operators may deploy multiple lwAFTRs for robustness, reliability, and load balancing. In lw4o6, subscriber to IPv4 and port-set mapping needs to be pre-provisioned in the lwAFTR before an IPv4 service can be provided.

For redundancy, one or more backup lwAFTR can have the subscriber bindings already provisioned, e.g. as part of the top-down provisioning process described above. In this case, the provisioning system is responsible for ensuring that the binding tables of the lwAFTRs are consistent. In this case, as customer traffic arriving or returning through either of the lwAFTRs will be processed in the same way, an active/active redundancy model is possible.

A second option, which could be more suitable for bottom-up provisioning, is for the bindings to be replicated between the primary lwAFTR and the backup lwAFTR. When the primary lwAFTR fails, the backup lwAFTR has the necessary binding table entries to correctly forward subscriber traffic. In this mode, the internal hosts are not required to re-initiate the bindings with the external

hosts. In lw4o6, as the number of binding states have been greatly reduced compared to DS-Lite, it is reasonable to adopt Hot Standby mode when there are only two lwAFTRs (a primary and a backup lwAFTR). However, if the number of lwAFTRs is larger than two, it is not scalable to deploy using hot standby mode since each two of the lwAFTRs should to synchronize the binding states.

4.4. Location of lwAFTRs in the Network

lwAFTR(s) can be deployed in a centralized or a distributed manner.

For a centralized deployment, the lwAFTR(s) are locacate in central aggregation points in the network, such as a core site, the exit point from a MAN etc. As the lwAFTR provides the gateway between the IPv6 and IPv4 networks, it allows single stack IPv6 to be deployed in the access part of the network.

In a distributed deployment, lwAFTR function is integrated with the BRAS/SR. Since newly emerging customers might be distributed in the whole Metro area, we have to deploy lwAFTR on all BRAS/SRs. This will cost a lot in the initial phase of the IPv6 transition period. This model also has the drawback of requiring both IPv4 and IPv6 to the BRAS/Service Router devices and so is unsuitable for providers wishing to build a single stack IPv6 only core.

4.5. Path Consistency Consideration

In Lightweight 4over6, if the binding state is not synchronized among multiple lwAFTRs, the lwAFTR in which the subscriber's binding state is stored should be exactly the one to service the subscriber. Otherwise, there will be no match in the lwAFTR. This can be achieved by using a unique IPv6 tunnel endpoint address and corresponding reachable public IPv4 customer prefix for each lwAFTR.

If multiple lwAFTRs are deployed for resilience or scalability, using the top-down provisioning model, all of the lwAFTRs in this cluster will share the same IPv6 tunnel endpoint and set of reachable prefixes. In this case, any packet arriving at any of the cluster members will be processed in the same way. However, it is worth considering using ECMP with flow-hashing so that a single customer's traffic will be processed by the same lwAFTR. This will reduce the change of packet re-ordering.

In Lightweight 4over6, if the binding state is not synchronized among multiple lwAFTRs, the lwAFTR in which the subscriber's binding state is stored should be exactly the one to service the subscriber. Otherwise, there will be no match in lwAFTR. This requires the provionsion packets (either using DHCPv4-over-DHCPv6 or PCP Port-set)

should arrive at the same lwAFTR as the subsequent IP-in-IP traffic. If multiple lwAFTRs are using the same Tunnel End Point address and there are intermediate routers between lwB4 and lwAFTR, there might be a problem when intermediate routers perform ECMP based on L4 hash for the plain provisioning packets while doing L3 hash for subsequent IP-in-IP traffic. In this case, it is recommended that the provisioning packet is sent over IPv6 tunnel so that intermediate routers can only process ECMP using L3 hash.

5. lwB4 Deployment Considerations

For the lwB4, the DNS Deployment Considerations and B4 Remote Management in [[RFC6908](#)] can also be applied here. In this section, only additional considerations relevant to lw4o6 are discussed.

5.1. NAT Traversal Issues

In lw4o6, as the subscriber's traffic source port will be restricted to the port-set allocated from the provisioning system, there will be an impact on some NAT traversal mechanisms. For example, in UPnP 1.0, the external port number that can be used by a remote peer is selected by a UPnP client in end host. If the client randomly selects a port number which does not fall in that valid port-set, the UPnP process will fail.

This is likely to happen because an end-host does not have knowledge of the port-set which has been allocated to the lwB4. More detailed experimental results can be found in [[I-D.deng-aplusp-experiment-results](#)]. This problem will not exist in UPnP 2.0 because the end-host's UPnP client will negotiate the external port number with the server. Another way is to implement a mechanism (e.g. [[RFC7753](#)]) in end host to fetch the allocated port-set from the lwB4. The UPnP client can then select the port number within the port-set.

5.2. Static Port Forwarding Configuration

Currently, some externally initiated applications rely on manual configuration to reserve a port in the CPE. The restricted port-set used by the lwB4 may be problematic for manual port forwarding configuration. It is recommended that the port-set allocated from the provisioning system should be visible to the user (e.g. via the configuration interface of a HGW which implements the lwB4 function), which can be used as a hint for subscribers to add port forwarding mapping.

It should also be noted that the well-known ports are not generally allocated to a lwB4, unless the client is being allocated a full IPv4

address with no address sharing ([\[RFC7596\], Section 5.1](#)). If the user wishes to make a service running on an end-host using a well-known port externally accessible, it is necessary to configure the lwB4's port-forwarding to re-map the well-known port to a port taken from the allocated port-set.

6. DS-Lite Compatibility Consideration

Lightweight 4over6 can be either deployed as a complete solution, or in conjunction with DS-Lite. Since Lightweight 4over6 does not have any extra requirement on IPv6 addressing, it can use the same addressing scheme with DS-Lite, together with routing policy, user management policy, etc. Besides, the bottom-up model has quite similar requirement and workflow on the supporting system with DS-Lite. Therefore, it is suitable for operators to deploy incrementally in existing DS-Lite network.

6.1. Case 1: Integrated Network Element with Lightweight 4over6 and DS-Lite AFTR Scenario

In this case, DS-Lite has been deployed in the network. Later in the deployment schedule, the operator decided to implement Lightweight 4over6 lwAFTR function in the same network element (depicted in Figure3, below). Therefore, the same network element needs to support both transition mechanisms.

There are two options to distinguish the traffic from two transition mechanisms.

The first one is to distinguish using the client's source IPv4 address. The IPv4 address from Lightweight 4over6 is public address as NAT has been done in the lwB4, and IPv4 address for DS-lite is private address as NAT will be done on AFTR. When the network element receives an encapsulated packet, it would de-capsulate packet and apply the transition mechanism based on the IPv4 source address in the packet. This requires the network element to examine every packet and may introduce significant extra load to the network element. However, both the B4 element and Lightweight 4over6 lwB4 can use the same DHCPv6 option [\[RFC6334\]](#) with the same FQDN of the AFTR and lwAFTR.

The second one is to distinguish using the destination's tunnel IPv6 address. One network element can run separated instances for Lightweight 4over6 and DS-Lite with different tunnel addresses. Then B4 element and Lightweight 4over6 lwB4 can use the same DHCPv6 option [\[RFC6334\]](#) with different FQDNs pointing to corresponding tunnel addresses. This requires the supporting system should distinguish different types of users when assigning the FQDNs in DHCPv6 process.

Another option is to use a new DHCPv6 option [[RFC7598](#)] to discover the lwAFTR's FQDN.

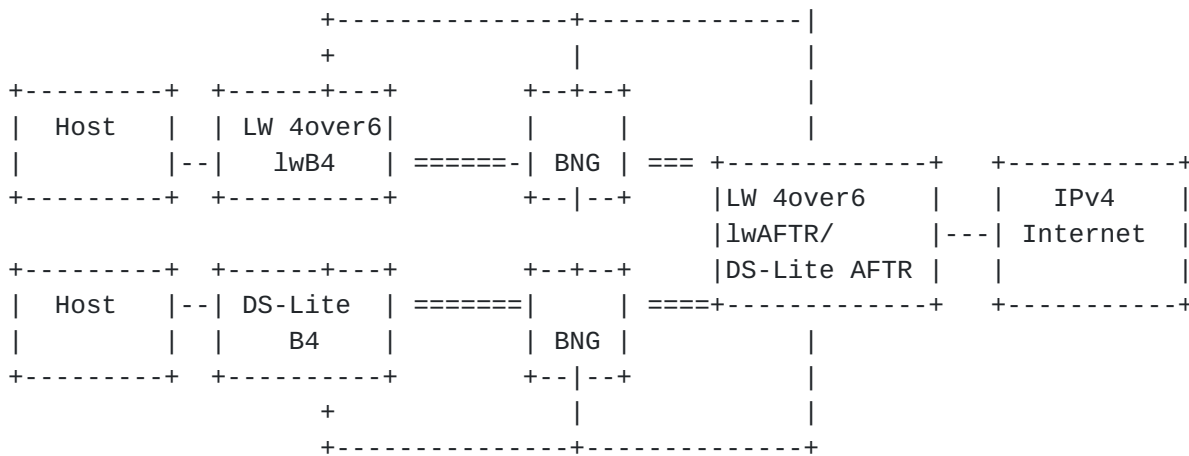


Figure 3 DS-Lite Coexistence scenario with Integrated AFTR

6.2. Case 2: DS-Lite Coexistent scenario with Separated AFTR

This is similar to Case 1. The difference is the lwAFTR and AFTR functions won't be co-located in the same network element (depicted in Figure4). This use case decouples the functions to allow more flexible deployment. For example, an operator may deploy AFTR closer to the edge and lwAFTR closer to the core. Moreover, it does not require the network element to pre-configure with the CPE's IPv6 addresses. An operator can deploy more AFTR and lwAFTR at needed. However, this requires the B4 and lwB4 to discover the corresponding network element. In this case, B4 element and Lightweight 4over6 lwB4 can still use [[RFC6334](#)] with different FQDNs pointing to corresponding tunnel end-point addresses, and the supporting system should distinguish different types of users.

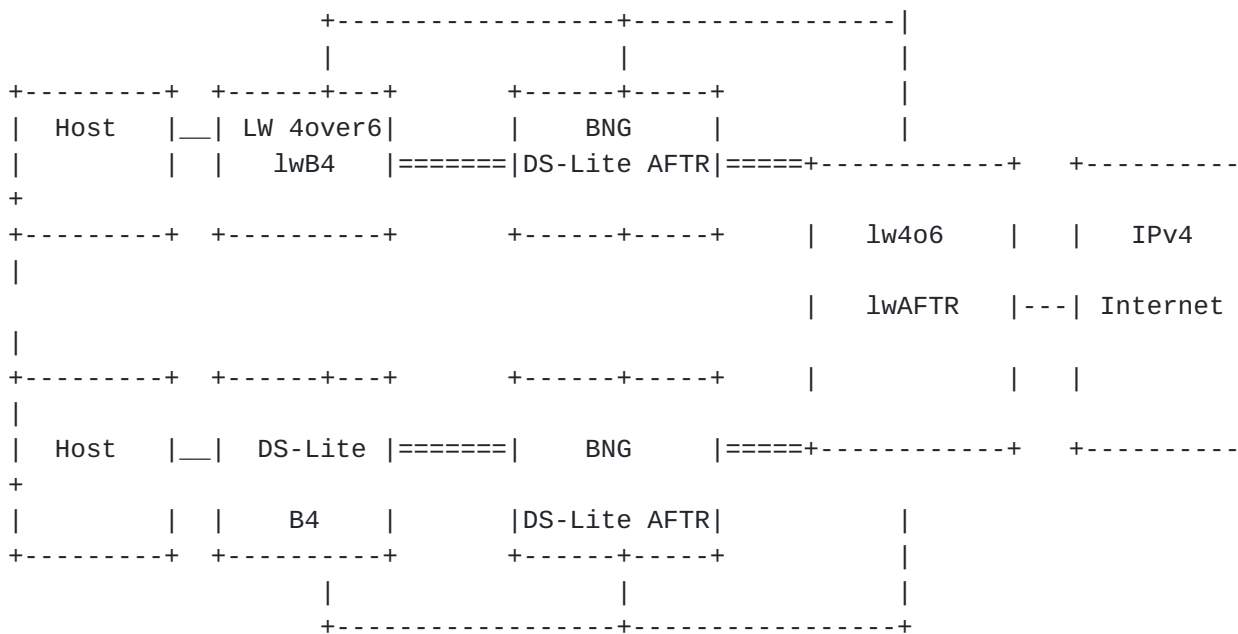


Figure 4 DS-Lite Co-existence scenario with Separated AFTR/lwAFTR

7. Acknowledgements

TBD

8. 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.cui-softwire-b4-translated-ds-lite]

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

[I-D.deng-aplusp-experiment-results]

Deng, X., Boucadair, M., and F. Telecom, "Implementing A+P in the provider's IPv6-only network", [draft-deng-aplusp-experiment-results-00](#) (work in progress), March 2011.

[I-D.ietf-behave-ipfix-nat-logging]

Sivakumar, S. and R. Penno, "IPFIX Information Elements for logging NAT Events", [draft-ietf-behave-ipfix-nat-logging-13](#) (work in progress), January 2017.

[I-D.ietf-behave-syslog-nat-logging]

Chen, Z., Zhou, C., Tsou, T., and T. Taylor, "Syslog Format for NAT Logging", [draft-ietf-behave-syslog-nat-logging-06](#) (work in progress), January 2014.

[I-D.ietf-dhc-dhcp4o6-saddr-opt]

Farrer, I., Sun, Q., Cui, Y., and L. Sun, "DHCPv4 over DHCPv6 Source Address Option", [draft-ietf-dhc-dhcp4o6-saddr-opt-00](#) (work in progress), March 2017.

[I-D.ietf-pcp-base]

Wing, D., Cheshire, S., Boucadair, M., Penno, R., and P. Selkirk, "Port Control Protocol (PCP)", [draft-ietf-pcp-base-29](#) (work in progress), November 2012.

[I-D.ietf-softwire-dslite-deployment]

Lee, Y., Maglione, R., Williams, C., Jacquenet, C., and M. Boucadair, "Deployment Considerations for Dual-Stack Lite", [draft-ietf-softwire-dslite-deployment-08](#) (work in progress), January 2013.

[I-D.ietf-softwire-map-radius]

Jiang, S., Fu, Y., Liu, B., Deacon, P., Xie, C., and T. Li, "RADIUS Attribute for Softwire Address plus Port based Mechanisms", [draft-ietf-softwire-map-radius-12](#) (work in progress), May 2017.

[I-D.ietf-softwire-yang]

Sun, Q., Wang, H., Cui, Y., Farrer, I., Zoric, S., Boucadair, M., and R. Asati, "A YANG Data Model for IPv4-in-IPv6 Softwires", [draft-ietf-softwire-yang-01](#) (work in progress), October 2016.

[I-D.lee-softwire-lw4over6-failover]

Lee, Y., Qiong, Q., and C. Liu, "Simple Failover Mechanism for Lightweight 4over6", [draft-lee-softwire-lw4over6-failover-01](#) (work in progress), July 2013.

[I-D.sun-softwire-lw4over6-dhcpv6]

Xie, C., Qiong, Q., Lee, Y., Tsou, T., and P. Wu, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Lightweight 4over6", [draft-sun-softwire-lw4over6-dhcpv6-00](#) (work in progress), July 2013.

[I-D.zhou-dime-4over6-provisioning]

Zhou, C., Taylor, T., Qiong, Q., and M. Boucadair, "Attribute-Value Pairs For Provisioning Customer Equipment Supporting IPv4-Over-IPv6 Transitional Solutions", [draft-zhou-dime-4over6-provisioning-05](#) (work in progress), September 2014.

- [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>>.
- [RFC6052] Bao, C., Huitema, C., Bagnulo, M., Boucadair, M., and X. Li, "IPv6 Addressing of IPv4/IPv6 Translators", [RFC 6052](#), DOI 10.17487/RFC6052, October 2010, <<http://www.rfc-editor.org/info/rfc6052>>.
- [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>>.
- [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>>.
- [RFC6334] Hankins, D. and T. Mrugalski, "Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Option for Dual-Stack Lite", [RFC 6334](#), DOI 10.17487/RFC6334, August 2011, <<http://www.rfc-editor.org/info/rfc6334>>.
- [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>>.
- [RFC6908] Lee, Y., Maglione, R., Williams, C., Jacquenet, C., and M. Boucadair, "Deployment Considerations for Dual-Stack Lite", [RFC 6908](#), DOI 10.17487/RFC6908, March 2013, <<http://www.rfc-editor.org/info/rfc6908>>.
- [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>>.
- [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>>.
- [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>>.
- [RFC7598] Mrugalski, T., Troan, O., Farrer, I., Perreault, S., Dec, W., Bao, C., Yeh, L., and X. Deng, "DHCPv6 Options for Configuration of Software Address and Port-Mapped Clients", [RFC 7598](#), DOI 10.17487/RFC7598, July 2015, <<http://www.rfc-editor.org/info/rfc7598>>.
- [RFC7600] Despres, R., Jiang, S., Ed., Penno, R., Lee, Y., Chen, G., and M. Chen, "IPv4 Residual Deployment via IPv6 - A Stateless Solution (4rd)", [RFC 7600](#), DOI 10.17487/RFC7600, July 2015, <<http://www.rfc-editor.org/info/rfc7600>>.
- [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>>.
- [RFC7753] Sun, Q., Boucadair, M., Sivakumar, S., Zhou, C., Tsou, T., and S. Perreault, "Port Control Protocol (PCP) Extension for Port-Set Allocation", [RFC 7753](#), DOI 10.17487/RFC7753, February 2016, <<http://www.rfc-editor.org/info/rfc7753>>.

Appendix A. China Telecom Experimental Results

We have deployed Lightweight 4over6 in our operational network of HuNan province, China. It is designed for broadband access network, and different versions of the lwB4 function have been implemented including a Linksys device, a software client for Windows XP, Windows Vista and Windows 7.

It can be integrated with existing dial-up mechanisms such as PPPoE, etc. The major objectives listed below aimed to verify the functionality and performance of Lightweight 4over6:

In our experiment, lwB4 will firstly get its IPv6 address and delegated prefix through PPPoE, and then initiate a PCP-extended request to get public IPv4 address and its valid port set. The

lwAFTR will thus create a subscriber-based state accordingly, and notify the syslog server with {IPv6 address, IPv4 address, port set, timestamp}.

A.2. Experimental Results

In our trial, we mainly focused on application and performance tests. The applications tested include web (HTTP/HTTPS), email, instant messaging, FTP, telnet, SSH, video, Video Camera, P2P, online gaming, and VoIP.

For the performance tests, we measured the number of concurrent session and throughput performance.

The experimental results are listed as follows:

Application Type	Test Result	Port Number Occupation
Web	OK IE, Firefox, Chrome	normal websites: 10~20 Ajax Flash webs: 30~40
Video	ok, web based or client based	30~40
Instant Message	OK QQ, MSN, gtalk, skype	8~20
P2P	OK utorrent,emule,xunlei	lower speed: 20~600 (per seed) higher speed: 150~300
FTP	need ALG for active mode, flashxp	2
SSH, TELNET	OK	1 for SSH, 3 for telnet
online game	OK for QQ, flash game	20~40

Figure 6 China Telecom Lightweight 4over6 Experiment Results

The performance tests for the lwAFTR are taken using a PC. Due to limitations of the PC hardware, the overall throughput is limited to around 800 Mbps. However, it can still support more than one hundred million concurrent sessions.

A.3. Conclusions

From the experiment, we reached the following conclusions:

- o Lightweight 4over6 has good scalability. As it is a lightweight solution that only maintains per-subscription state information, it can easily support a large amount of concurrent subscribers.
- o Lightweight 4over6 can be deployed rapidly. No modifications to the existing addressing and routing system in our operational network was necessary. Logging of customer address allocations was easy to implement.
- o Lightweight 4over6 can support the majority of current IPv4 applications commonly in use.

Appendix B. Tsinghua University Experimental Result

Lightweight 4over6 has been deployed in the campus of Tsinghua University, China. DHCPv4o6 [[RFC7341](#)] is used for dynamically provisioning the lwB4's IPv4 address and port set [[RFC7618](#)]. Wireless APs for Lightweight 4over6, were deployed, covering a large portion of the campus, allowing mobile devices to connect to the lwB4. We also deployed a lwB4 gateway in some of our buildings so that end device could connect directly to the lwB4. Users access the IPv4 Internet through the China Next Generation Internet (CNGI) IPv6 Network.

B.1. Experimental Environment

The network topology for this experiment is depicted in Figure 7.




Figure 7 Tsinghua University Lightweight 4over6 Experiment Topology

In this deployment model, the lwAFTR is co-located with a DHCPv4 server to assign port-restricted IPv4 addresses and port-sets to the lwB4. The lwAFTR snoops the DHCPv4 over DHCPv6 messages generated or received by the DHCPv4 server and updates its binding table accordingly.

In our experiment, the lwB4 receives its IPv6 address through static or dynamic configuration. It then sends a DHCPv4o6 request to the lwAFTR device to get the public IPv4 address and valid port-set. The lwAFTR will add the IPv6 address, IPv4 address, and port-set information of the lwB4 into its binding table.

B.2. Experimental Results

In the Tsinghua University experiment, the performance of various applications were tested including web (HTTP/HTTPS), email, instant messaging, FTP, telnet, SSH, video, Video Camera, P2P, online gaming, and VoIP. We also tested different terminal devices including PC/laptop computers, and cell phone. These devices used different operating systems, including Windows 7, MacOS, Android, and Apple IOS.

The experimental results were as follows:

Application Type	Test Applications	Test Subjects	Result
Web	IE, Chrome, Sougou	Browse websites, download files	OK
Video	Youku, pptv, qqlive (Web based, client based)	VOD, live video	OK
P2P	Bittorrent, xunlei	Download files	OK
Ping/tracert	Command line	Ping/tracert URL	OK
TELNET/SSH	Putty, secureCRT	Telnet/SSH login	OK
Email	126, QQ, hotmail (Web based, client based)	Send/receive email	OK
Cloud storage	Baidu Cloud	Upload/download files	OK
Instant messaging	Skype, QQ	Send/receive messages	OK
Online gaming	QQ game	Enter game	OK
Online payment	JD, Taobao	Complete payment	OK

Figure 8 Tsinghua University Lightweight 4over6 experimental result

B.3. Conclusion

Lightweight 4over6 supports the majority of current IPv4 applications and services. The user experience of using Lightweight 4over6 is no different from using the native IPv4 network. It can satisfy the IPv4 network service demands of IPv6 network users.

Authors' Addresses

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

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

Chongfeng Xie
China Telecom
Room 708, No.118, Xizhimennei Street
Beijing 100035
P.R.China

Phone: +86-10-58552116>
Email: xiechf@ctbri.com.cn

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

Email: yiu_lee@cable.comcast.com

Maoke Chen
FreeBit Co., Ltd.
13F E-space Tower, Maruyama-cho 3-6
Shibuya-ku, Tokyo 150-0044
Japan

Email: fibrib@gmail.com

Tianxiang Li
Tsinghua University
Beijing 100084
P.R.China

Phone: +86-10-6278-5822
Email: peter416733@gmail.com

Ian Farrer
Deutsche Telekom AG
Bonn 53227
Germany

Email: ian.farrer@telekom.de

