



■ **Deployment Guide**

DS-Lite

AX Series

TABLE OF CONTENTS

1	Dual-Stack Lite Overview	3
1.1.1	Architecture	3
1.1.2	Home Gateway.....	5
1.1.3	Address Family Transition Router.....	6
2	Configuration.....	7
2.1	IPv6 Configuration.....	9
2.1.1	DS-Lite Configuration	11
2.2	Configuration Options	15
2.2.1	DS-Lite ALG Support	15
2.2.2	DS-Lite Port Reservation	15
2.2.3	IPv4 Address Filtering	16
2.2.4	Additional DS-Lite Options	16
2.2.5	Fragmentation Options.....	17
2.2.6	TCP Maximum Segment Size Clamping.....	18
2.2.7	Logging.....	19
2.2.8	Network Integration	19
3	Conclusion	21

1 DUAL-STACK LITE OVERVIEW

As the available IPv4 address space continues to erode, many service providers are investigating solutions that allow rapid migration to IPv6. However, an abrupt transition to IPv6 is simply not practical due to many factors including access network upgrade costs, network availability, customer SLA, lack of content over IPv6, and legacy IPv4 clients that have minimal or no support for IPv6. Therefore service providers must be prepared to support both IPv6 and IPv4-based customer applications well into the future while planning and executing their IPv6 migration strategy.

Dual-Stack Lite (DS-Lite), defined in RFC 6333, allows a service provider to share existing IPv4 address space and support both IPv6 and IPv4 clients utilizing an IPv6 infrastructure. This allows for preservation of the IPv4 address space by reclaiming addresses from the access network as it migrates to IPv6, and sharing the existing IPv4 addresses among its customer base. Unlike other migration strategies, DS-Lite combines both tunneling and network address translation technologies, and decouples the service provider's access network from the public internet. These features can simplify the migration to IPv6 by allowing incremental IPv6 deployment within the service provider's network while continuing to support legacy IPv4 clients.

1.1.1 ARCHITECTURE

The DS-Lite architecture consists of the following elements:

- DS-Lite home gateway containing the Basic Bridging Broadband (B4) element
- Address Family Transition router (AFTR)
- IPv4-in-IPv6 tunnels, referred to as "softwires" (defined in RFC 4925), connecting the B4 element to the AFTR.

Please refer to [Figure 1](#).

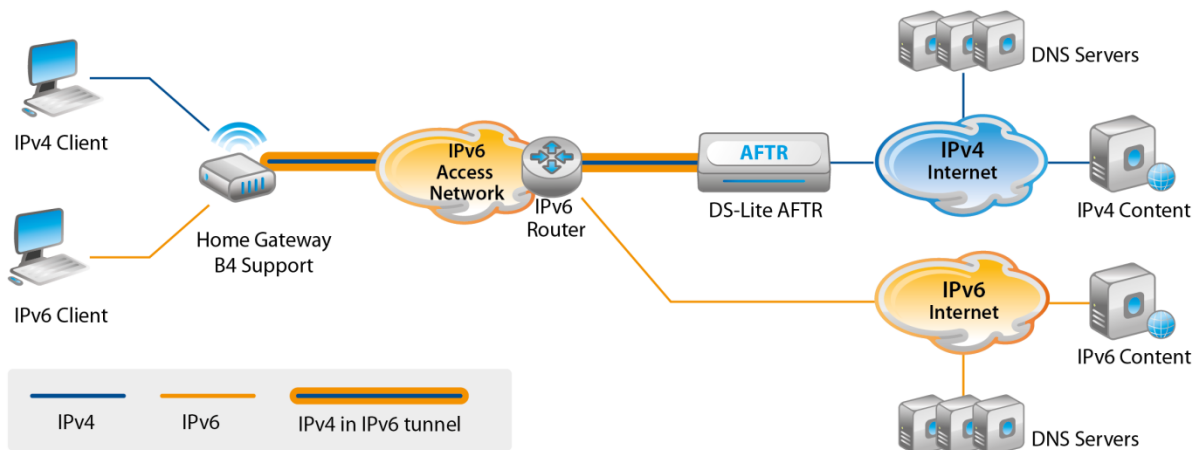


Figure 1: DS-Lite architecture

DS-Lite is deployed across an all-IPv6 infrastructure, which natively supports IPv6 clients, and tunnels IPv4 packets using softwires to the AFTR. The AFTR inspects the IPv6 packet, exposing the IPv4 source and destination information, and then performs NAT44, translating the source address using pooled IPv4 resources. The home gateway is not required to perform NAT, but rather utilizes tunneling to transport IPv4 packets to the AFTR, which performs NAT using Carrier Grade NAT (CGN/CGNAT). This architecture model provides the following advantages:

- DS-Lite allows incremental migration to IPv6.
- DS-Lite provides native IPv6 client access to IPv6 content as well as IPv4 legacy client support for IPv4 internet services.
- The DS-lite model is implemented across an all-IPv6 environment, reclaiming IPv4 address space and allowing the sharing of IPv4 addresses with CGN.
- Softwires are created automatically within the AFTR upon receipt of a valid IPv6 packet designated for DS-Lite processing. No specific tunnel configuration is required.
- Softwires can terminate anywhere in a service provider's network, allowing for easy load distribution and scaling.
- The binding table within the AFTR includes the softwire ID (IPv6 address of the B4), eliminating the need for unique IPv4 client addresses.
- DS-Lite does not require, nor does it implement, IP protocol translation (IPv4-to-IPv6 translation), minimizing application breakage due to embedded IP literals.

1.1.2 HOME GATEWAY

The DS-Lite home gateway is implemented across an all-IPv6 infrastructure; therefore, only a single IPv6 address is configured on the IPv6 interface. This address will serve as the softwire ID, uniquely identifying the CPE within the AFTR. The home gateway acts as an IPv4 DHCP server delivering RFC 1918 addresses to IPv4 clients, and supports DHCPv6 or address auto-configuration for IPv6 clients. The home gateway does not perform NAT; rather, it delivers native IPv6 packets directly to the IPv6 access network and IPv4 packets to the B4 element within the gateway.

The B4 element is responsible for softwire initiation, constructing the tunnel to the AFTR. The softwire is a multipoint-to-point tunnel, terminating on a designated AFTR. When a client accesses an IPv4 service, it sources the packet from its private IP address (RFC 1918) and sends it to the home gateway. The B4 element intercepts the packet, encapsulates it within IPv6, and delivers it to its designated AFTR through the softwire. This is in direct contrast to the home gateway NAT operations required by other strategies such as NAT444. The AFTR inspects the packet, exposing the IPv4 source and destination information, and translates the source address using the AFTR's CGN capability. The IPv6 address of the AFTR endpoint can be configured manually on the home gateway or discovered through DHCPv6 options, as defined in RFC 6334.

DS-Lite supports an all-IPv6 access infrastructure; since the home gateway is an IPv6 end-point, all DNS resolution must occur over IPv6. This operation should be confined within the service provider's IPv6 network, such that the DNS requests over IPv6 are never sent to the AFTR. The home gateway can easily pass the IPv6 DNS server address to its downstream IPv6 clients but not to the IPv4 clients. To resolve this issue, the B4 should implement a DNS proxy function, to allow IPv4 clients access to DNS services. In the case where DNS proxy is not supported, or other circumstances negate its use, the DNS requests originating from the IPv4 clients may flow across the softwire to the AFTR, where they will be forwarded to the DNS server, as defined on the client. This has the effect of creating a binding within the CGN device, possibly taking resources from traffic bearing services. The AX Series AFTR supports a binding lifetime of 3 seconds for DNS requests.

Configuration of the DS-Lite home gateway is vendor-specific and out of the scope of this document; however, [Figure 2](#) shows a sample configuration.

INTERNET CONNECTION TYPE

Choose the mode to be used by the router to connect to the Internet.

My Internet Connection is :

AFTR ADDRESS INTERNET CONNECTION TYPE :

Enter the AFTR address information provided by your Internet Service Provider(ISP).

DS-Lite Configuration DS-Lite DHCPv6 Option Manual Configuration

AFTR IPv6 Address :

B4 IPv4 Address : (optional)

WAN IPv6 Address : **2001::2:1/112**

IPv6 WAN Default Gateway : **2001::2:2**

Figure 2: Sample DS-Lite configuration on a home router

1.1.3 ADDRESS FAMILY TRANSITION ROUTER

The AFTR terminates the softwire (multipoint-to-point IPv6 tunnel) and decapsulates the packet, exposing the IPv4 source and destination information. The CGN function then maps the source address (RFC 1918 address of the IPv4 client) to a public address from the NAT pool, and creates a binding including the softwire ID (IPv6 address of the home gateway). The softwire ID is unique, allowing for overlapping RFC 1918 addresses across the service provider's home gateways.

DS-Lite performance and feature capability is highly dependent upon the CGN implementation, architecture, and performance. A10 Networks' DS-Lite implementation supports the latest, industry-leading CGN feature set, including stateful high availability (HA), logging reduction techniques, end-point-independent mapping and filtering (EIM/EIF), address-pool paired behavior (sticky NAT), override actions, hairpinning, user/session quotas, and NAT session/STUN timeouts. These features are described in more detail in the [Carrier Grade NAT \(CGN\) Deployment Guide](#).

2 CONFIGURATION

This section presents a typical DS-Lite topology, including detailed configuration for the DS-Lite AFTR and CGN.

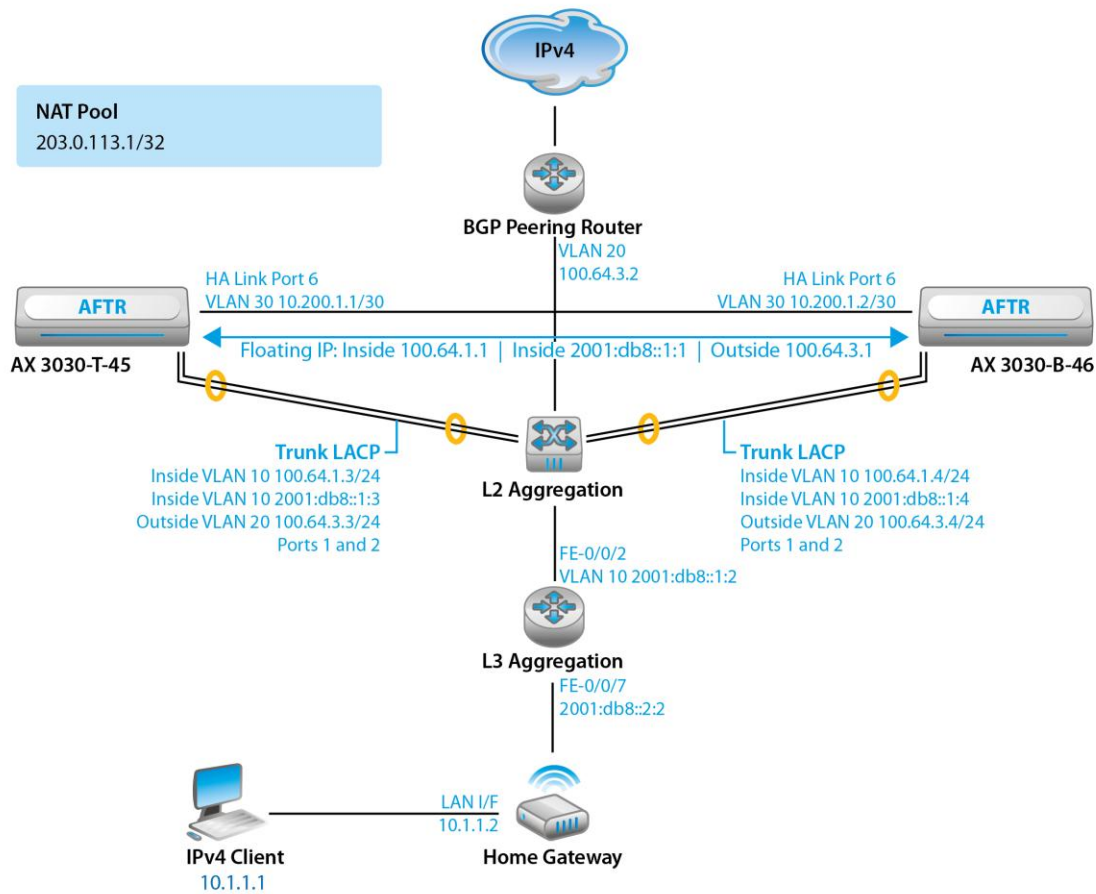


Figure 3: DS-Lite lab topology

The configuration example in [Figure 3](#) illustrates a DS-Lite deployment that consists of a single IPv4 client, directly connected to the DS-Lite home gateway. The IPv4 client connects to the IPv4 Internet through dynamic CGN mapping, utilizing the 203.0.113.1/32 NAT pool. The aggregation router and Layer 2 switch are connected to two AX devices (providing AFTR function) running in HA mode, and are connected using LACP link aggregation. A dedicated HA link is utilized for clarity and simplicity, although the HA protocol also can use LACP connections. Finally, the BGP peering router is connected to the aggregation router to provide the uplink to the Internet.

This example uses dynamic routing protocols to redistribute the NAT pool and the floating IP address. IPv6 static routing is used between the DS-Lite home gateway and the aggregation router; OSPFv3 is used in the access network between the aggregation layers and the AX devices.

Internet routing is provided by BGP and is configured between the BGP peering router and the AX device. The BGP peering router announces a default route towards the AX device, and the AX device announces the NAT pool, and modifies the next-hop to the outside floating IP address (100.64.3.1).

In the access network, OSPF3 is configured between the AX inside interfaces and the aggregation router. The AX redistributes floating IP address 2001:db8::1:1 (the AFTR's IPv6 tunnel end-point) to the aggregation router, and receives the home gateway subnet (2001:db8::2:0/112) through OSPFv3. Finally, the DS-Lite home gateway is configured with local IPv6 address 2001:db8::2:1, default-gateway 2001:db8::2:2, and AFTR address 2001:db8::1:1

To further clarify, below is a packet walkthrough of the topology:

1. IPv4 client sends a DNS request for www.example.com using IPv4.
2. Home gateway proxies the request from the client and sends it to the service provider's IPv6 DNS servers. DNS returns 192.0.43.10.
3. Client sends a TCP SYN to 192.0.43.10:80 source address 10.1.1.1:2000
4. Home gateway B4 element intercepts the IPv4 packet, encapsulates it into IPv6, and sends the encapsulated packet to AFTR. IPv6 destination is the floating address of the AFTR HA pair (2001:db8::1:1).
5. The AFTR receives the IPv6 packet and strips off the IPv6 header, exposing the original IPv4 packet. The CGN component translates the source address and port 10.1.1.1:2000 to an address selected from the NAT pool (203.0.113.1:2000). The CGN device creates a binding that includes the IPv6 softwire ID (2001:db8::2:1), originating IPv4 client address (10.1.1.1:2000), and translated source address (203.0.113.1:2000).
6. The AFTR sends the packet to the destination (192.0.43.10:80) with source address 203.0.113.1:2000.
7. Site www.example.com responds with a TCP SYN-ACK with destination address 203.0.113.1:2000 and source address 192.0.43.10.
8. AFTR receives the packet, consults the NAT binding table, and finds an entry corresponding to 203.0.113.1:2000. The AFTR translates the destination address to 10.1.1.1:2000 and encapsulates it into IPv6 using the softwire ID in the binding table. In this case, the softwire ID and destination IPv6 address are 2001:db8::2:1.
9. B4 element strips off the IPv6 header, exposing the IPv4 destination and source address. The home gateway then routes the packet to the destination (10.1.1.1:2000).

2.1 IPV6 CONFIGURATION

1. Configure an IPv6 IP address to the inside Virtual Ethernet (VE) interface for VLAN 10:

```
AX3030-T-47-Active(config)#interface ve 10
AX3030-T-47-Active(config-if:ve10)#ipv6 address 2001:db8::1:3/112
```

Note: This guide builds upon the “NAT44” configuration described in the [Carrier Grade NAT \(CGN\) Deployment Guide](#).

2. Add the IPv6 floating IP address to the HA configuration:

```
AX3030-T-47-Active(config)#floating-ip 2001:db8::1:1 ha-group 1
```

3. Configure OSPF3:

```
AX3030-T-47-Active(config)#inter ve 10
```

```
AX3030-T-47-Active(config-if:ve10)#ipv6 router ospf area 0 tag 1 instance-id 0
AX3030-T-47-Active(config-if:ve10)#exit
AX3030-T-47-Active(config)#router ospf
AX3030-T-47-Active(config-router)#router-id 10.10.10.10
```

Note: OSPFv3 supports multiple instances per interface. Most routers default to instance ID 0. If the instance ID does not match, the adjacency cannot form.

4. Verify IPv6 address settings:

```
AX3030-T-47-Active#show ipv6 interfaces
```

Port	IPv6 Address	Scope	Type
ve10	2001:db8::1:3/112	global	unicast
	fe80::21f:a0ff:fe02:7f86/64	link-local	unicast

```
AX3030-T-47-Active#show ipv6 neighbor
```

```
Total IPv6 neighbor entries: 4
```

IPv6 Address	MAC Address	Type	Age	State	Interface	Vlan
fe80::21f:a0ff:fe02:8076	001F.A002.8076	Dynamic	178	Reachable	ve10	10
2001:db8::1:4	001F.A002.8076	Dynamic	11	Reachable	ve10	10
2001:db8::1:2	0026.88FA.2FC2	Dynamic	207	Reachable	ve10	10
fe80::226:88ff:fefa:2fc2	0026.88FA.2FC2	Dynamic	173	Reachable	ve10	10

AX3030-T-47-Active#show ipv6 traffic

Traffic Type	Received	Sent	Errors
Router Solicit	0	0	0
Router Adverts	0	0	0
Neigh Solicit	8	265	0
Neigh Adverts	0	4	0
Echo Request	0	0	0
Echo Replies	0	0	0
Other ICMPv6 Errs	10	0	0

AX3030-T-47-Active#show ipv6 route

IPv6 Routing Table

Codes: K - kernel route, C - connected, S - static, R - RIP, O - OSPF,
I - IS-IS, B - BGP

Timers: Uptime

C 2001:db8::1:0/112 via ::, ve 10, 00:36:49

O 2001:db8::2:0/112 [110/101] via fe80::226:88ff:fefa:2fc2, ve 10, 00:36:02

AX3030-T-47-Active#show ipv6 ospf interface

ve 10 is up, line protocol is up

Interface ID 78

IPv6 Prefixes

fe80::21f:a0ff:fe02:7f86/64 (Link-Local Address)

2001:db8::1:3/112

OSPFv3 Process (1), Area 0.0.0.0, Instance ID 0

Router ID 10.10.10.10, Network Type BROADCAST, Cost: 100

Transmit Delay is 1 sec, State DROther, Priority 1

```
Designated Router (ID) 100.100.100.101
  Interface Address fe80::226:88ff:fe02:2fc2
Backup Designated Router (ID) 10.10.10.13
  Interface Address fe80::21f:a0ff:fe02:8076
Timer interval configured, Hello 10, Dead 40, Wait 40, Retransmit 5
Hello due in 00:00:02
Neighbor Count is 2, Adjacent neighbor count is 2
```

```
AX3030-T-47-Active#show ipv6 ospf neighbor
OSPFv3 Process (1)
Neighbor ID      Pri   State           Dead Time   Interface   Instance ID
10.10.10.13      1    Full/Backup     00:00:29   ve 10      0
100.100.100.101 128  Full/DR         00:00:40   ve 10      0
```

2.1.1 DS-LITE CONFIGURATION

The following steps are required to configure DS-Lite:

- Configure a NAT pool (or group of pools) containing the IPv4 addresses to use for traffic translation.
 - Configure a Limit ID (LID) and add the pool or pool group.
 - Configure a class list that matches the DS-Lite home gateway IPv6 addresses, and map it to the LID.
 - Bind the class list to the DS-Lite feature. The class list applies to packets from the inside NAT interface to the outside NAT interface. There can be at most 1 class list used for this purpose.
 - Enable IPv6 inside NAT on the interface connected to the IPv6 clients.
 - Enable IPv4 outside NAT on the interface connected to the IPv4 Internet.
1. To configure the NAT pool for DS-Lite translation, include the pool name, addresses and HA group. Note that the **Isn** option is required for this pool group. Enter the following command at the global configuration level.

```
AX3030-T-47-Active(config)# ip nat pool ds_lite_1 203.0.113.1 203.0.113.1
netmask /32 ha-group-id 1 lsn
```

Alternatively, NAT pools can be combined into pool groups. This simplifies future changes to the configuration and allows non-contiguous address bundling. Use the following command to create a pool group. Declare a pool group name (example) and list the NAT pools to be included in the group.

```
AX3030-T-47-Active(config)#ip nat pool-group example ds_lite_1
ha-group-id 1
```

2. Create the DS-Lite LID. The LID associates the DS-Lite pool or pool group with specific configuration options, including user quota, IPv4 address filtering, and rule lists. The operator can specify up to 1024 LIDs. Begin the configuration by assigning a LID number. This enters the LSN-LID configuration level.

```
AX3030-T-47-Active(config)#lsn-lid 2
```

Specify the NAT pool or pool groups to assign to this LID.

```
AX3030-T-47-Active(config-lsn-lid)#source-nat-pool ds_lite_1
```

Specify optional parameters for this NAT pool. Refer to the “Configuration Options” section in the [Carrier Grade NAT \(CGN\) Deployment Guide](#) for more details.

```
AX3030-T-47-Active(config)#user-quota icmp 50
AX3030-T-47-Active(config)#user-quota udp 250 reserve 0
AX3030-T-47-Active(config)#user-quota tcp 250 reserve 0
```

3. Create the class list, specifying the internal subnets and hosts that are associated with a specific LID. In this example, the class list “ds_lite_cn” classifies all DS-Lite home gateways assigned with an address from the subnet 2001:db8::2:0/112, and is tied to the configuration in LID 2. This address also is the B4 tunnel end-point that represents the softwire ID.

```
AX3030-T-47-Active(config)#class-list ds_lite_cn
AX3030-T-47-Active(config-class list)#2001:db8::2:0/112 lsn-lid 2
```

Note: The wildcard IPv6 address (::/0) and host addresses may be configured within the class list.

4. Bind the class list to the DS-Lite process.

```
AX3030-T-47-Active(config)#ds-lite inside source class-list ds_lite_cn
```

5. Declare the interfaces for NAT operation. Inside NAT is configured for client-side interfaces, while outside NAT is configured for interfaces connecting to the public Internet.

```
AX3030-T-47-Active (config) #interface ve 20
AX3030-T-47-Active (config-if:ve20) #ip nat outside
AX3030-T-47-Active (config-if:ve20) #inter ve 10
AX3030-T-47-Active (config-if:ve10) #ipv6 nat inside
```

6. Verify proper DS-Lite operation.

```
AX3030-T-47-Active#show session ds-lite
Traffic Type                               Total
-----
TCP Established                             2
TCP Half Open                              0
UDP                                          1
Non TCP/UDP IP sessions                    0
Other                                       0
Reverse NAT TCP                             0
Reverse NAT UDP                             0
Curr Free Conn                             33325051
Conn Count                                  1061
Conn Freed                                  1058
TCP SYN Half Open                           0
Conn SMP Alloc                              0
Conn SMP Free                               0
Conn SMP Aged                               0
Conn Type 0 Available                       66322432
Conn Type 1 Available                       33325051
Conn Type 2 Available                       16662496
Conn Type 3 Available                       8290304
Conn SMP Type 0 Available                   66322432
Conn SMP Type 1 Available                   33161216
Conn SMP Type 2 Available                   16580608
Conn SMP Type 3 Available                   8298479
```

```
Prot IP Type  IP Address and Port                               Age  Hash  Flags
-----
Udp  Fwd Src  [2001:db8::2:1]10.1.1.1:123                                60   1     NS
     Fwd Dst  [2001:db8::1:1]17.171.4.36:123
     Rev Src  17.171.4.36:123
     Rev Dst  203.0.113.1:19735
Tcp  Fwd Src  [2001:db8::2:1]10.1.1.1:49501                              300  2     NS
     Fwd Dst  [2001:db8::1:1]74.125.134.102:80
     Rev Src  74.125.134.102:80
     Rev Dst  203.0.113.1:49501
Tcp  Fwd Src  [2001:db8::2:1]10.1.1.1:49499                              300  5     NS
     Fwd Dst  [2001:db8::1:1]74.125.134.138:80
     Rev Src  74.125.134.138:80
     Rev Dst  203.0.113.1:49499
```

```
AX3030-T-47-Active#show ds-lite statistics
DS-Lite Statistics:
-----
Total TCP Ports Allocated           321
Total TCP Ports Freed                321
Total UDP Ports Allocated            741
Total UDP Ports Freed                741
Total ICMP Ports Allocated           2
Total ICMP Ports Freed                2
Data Session Created                 1063
Data Session Freed                   1063
User-Quota Created                   37
User-Quota Freed                     37
User-Quota Creation Failed           0
TCP NAT Port Unavailable              0
UDP NAT Port Unavailable              0
ICMP NAT Port Unavailable             0
New User NAT Resource Unavailable     0
TCP User-Quota Exceeded               0
UDP User-Quota Exceeded               0
ICMP User-Quota Exceeded              0
Extended User-Quota Matched           0
Extended User-Quota Exceeded          0
Data Session User-Quota Exceeded      0
TCP Full-cone Session Created         0
TCP Full-cone Session Freed           0
UDP Full-cone Session Created         0
UDP Full-cone Session Freed           0
Full-cone Session Creation Failed     0
Hairpin Session Created               0
Self-Hairpinning Drop                 0
Endpoint-Independent Mapping Matched  0
Endpoint-Independent Filtering Matched 0
Endpoint-Dependent Filtering Drop     0
NAT Pool Mismatch Drop                0
TCP Port Overloaded                   0
UDP Port Overloaded                   0
TCP Port Overloading Session Created  0
UDP Port Overloading Session Created  0
TCP Port Overloading Session Freed    0
UDP Port Overloading Session Freed    0
Truncated Packet                      0
LSN LID Drop                           0
LSN LID Pass-through                  0
No Class-List Match                   0
Permit Class-List Drop                 0
```

2.2 CONFIGURATION OPTIONS

2.2.1 DS-LITE ALG SUPPORT

The AX series DS-Lite implementation provides ALG support for the following protocols:

- FTP
- PPTP
- RTSP
- SIP
- TFTP

FTP is supported by default. To enable additional ALG support, use the following command:

ds-lite alg {ftp | pptp | rtsp | sip | tftp} enable

For example, to enable ALG support for SIP, enter the following command:

```
AX3030-T-47-Active (config) #ds-lite alg sip enable
```

For SIP contact NAT mappings, the corresponding full-cone session's Session Traversal Utilities for NAT (STUN) timeout is set to the "Expires" value in the SIP Registration packet's payload. For SIP RTP/RTCP NAT mappings, the corresponding full-cone session's STUN timeout is configurable. The RTP/RTSP STUN timeout can be 2-10 minutes. The default is 5 minutes. For example, to change the RTP/RTCP STUN timeout for full-cone sessions used for SIP NAT mappings to 10 minutes, use the following command at the global configuration level of the CLI:

```
AX3030-T-47-Active (config) #ip nat lsn alg sip rtp-stun-timeout 10
```

2.2.2 DS-LITE PORT RESERVATION

The AX Series' implementation of DS-Lite supports static mappings for a range of protocol ports. To enable static mappings, enter the following command at the global configuration level:

```
ds-lite port-reservation
  inside ipv6-tunnel-source ipv6-tunnel-destination
  ipv4-inside-addr inside-start-port inside-end-port
  nat nat-ipaddr nat-start-portnum nat-end-portnum
```

For example, the following command maps all ports from client address 10.1.1.1 on the home gateway with softwire ID 2001:db8::2:1 to the NAT pool address 203.0.113.1:

```
AX3030-T-47-Active(config)#ds-lite port-reservation inside 2001:db8::2:1 2001:db8::1:1
10.1.1.1 1 65535 nat 203.0.113.1 1 65535
```

To verify port reservations, use the following CLI command:

```
AX3030-T-47-Active#show ds-lite port-reservations
```

```
LSN Port Reservations
```

Tunnel Src IPv6 Address	Tunnel Dst IPv6 Address	Inside Address	Start	End	NAT Address	Start	End
2001:db8::2:1	2001:db8::1:1	10.1.1.1	1	65535	203.0.113.1	1	65535

1

2.2.3 IPV4 ADDRESS FILTERING

By default, the AFTR will accept any IPv4 address within an established software and send it to the CGN process for translation. The AX Series allows for explicit filtering of IPv4 addresses within the tunnel, which is accomplished by configuring the addresses to be allowed within a class-list (implicit deny), and enabling the filter under the DS-Lite configured LID. For example, if the AFTR should deny all IPv4 addresses with the exception of 10.1.1.1 issue the following commands from the CLI:

```
AX3030-T-47-Active(config)#class-list IPV4_Permit
AX3030-T-47-Active(config-class list)#10.1.1.1 /32
exit
AX3030-T-47-Active(config)#lsn-lid 2
AX3030-T-47-Active(config-lsn lid)#ds-lite inside-src-permit-list IPV4_Permit
```

2.2.4 ADDITIONAL DS-LITE OPTIONS

The AX DS-Lite implementation shares many features in common with CGN. These features include:

- NAT session timeouts
- Endpoint-Independent Mapping (EIM)
- Endpoint-Independent Filtering (EIF)
- Hairpinning
- NAT IP selection method
- Port-batching
- Port-overloading

- STUN-timeout
- SYN-timeout
- User and Session quotas

These configuration options can be found under the *ltn-lid*, *ip nat*, or *ip nat ltn* configuration hierarchy. Please refer to “Advanced Configuration Options” of the [Carrier Grade NAT \(CGN\) Deployment Guide](#) for more information and configuration details.

2.2.5 FRAGMENTATION OPTIONS

DS-Lite is designed to use a tunneling architecture, which introduces the possibility of MTU issues. Tunneling requires additional header overhead, which could exceed the physical interface MTU, resulting in poor performance for both web and FTP applications. Where possible, the service provider should increase the MTU size to account for the additional tunnel overhead contributed by the IPv6 header. For those cases where it is not possible to support an appropriately sized MTU from the B4 to the AFTR, the operator can configure the AX series to support fragmentation of inbound and outbound packets.

By default, fragmentation for inbound IPv6 tunnel packets is enabled but fragmentation of IPv4 packets within the tunnel is disabled. In the outbound direction, IPv4 fragmentation is enabled by default and cannot be disabled. If an IPv4 packet has the Don't Fragment bit set, the AX device does not fragment the packet, and instead returns an ICMP Type 3 Code 4 (Fragmentation Needed and DF Set) message.

These behaviors can be modified in the CLI, with the following options:

- DS-Lite fragmentation inbound – The following options are available:
 - drop – Drops inbound fragmented packets
 - ipv4 – Uses IPv4 fragmentation for oversize packets
 - ipv6 – Uses IPv6 fragmentation for oversize packets; **the default behavior**
 - df-set – If the DF bit is set, there are additional options:
 - drop – Drops silently
 - ipv4 – Overrides using IPv4 fragmentation for oversize packets
 - ipv6 – Overrides using IPv6 fragmentation for oversize packets
 - send-icmp – Sends ICMP Type 3 Code 4 (Fragmentation Needed and DF Set); **the default behavior**
- DS-Lite fragmentation outbound. The following options are available:
 - Drop – Drops fragmented packets
 - ipv4 – Uses IPv4 fragmentation for oversize packets; **the default behavior**
 - send-icmpv6 – Sends ICMPv6 Type 2 Code 0 (Packet Too Big)
 - df-set – If the DF bit is set, there are additional options:
 - drop – Drops silently
 - ipv4 – Overrides using IPv4 Fragmentation

- send-icmp – Sends ICMP Type 3 Code 4 (Fragmentation Needed and DF Set); **the default behavior**
- send-icmpv6 – Sends ICMP Type 2 Code 0 (Packet Too Big)

For example, to configure DS-Lite fragmentation behavior to drop all fragmented packets in the outbound direction, enter the following command at the global configuration level:

```
AX3030-T-47-Active(config)# ds-lite fragmentation outbound drop
```

2.2.6 TCP MAXIMUM SEGMENT SIZE CLAMPING

In addition to fragmentation, the AX supports TCP Maximum Segment Size (MSS) adjustment for DS-Lite. The MSS specifies the maximum length of the TCP segment allowed in any single IP packet, excluding the IP and TCP headers. The MSS value is set by the client, allowing for the size of the additional IPv4 headers. However, IPv4 clients behind DS-Lite home gateways are not DS-Lite aware, and cannot account for the tunnel overhead introduced by the IPv6 header.

To minimize fragmentation, the DS-Lite AFTR must ensure that TCP segments placed into the IPv6 tunnel allow for the additional overhead of both the IPv4 and IPv6 headers. The process checks the MSS value during TCP connection initiation and makes adjustments if required. The following configuration options are available:

- None – The AX device does not change the MSS value.
- Fixed value – The AX device changes the MSS to the length you specify.
- Subtract – The AX device reduces the MSS if it is longer than the specified number of bytes. This option sets the MSS based on the following calculations:
 - If MSS minus S is greater than N , subtract S from the MSS.
 - If MSS minus S is less than or equal to N , set the MSS to N .
 - Using these values, the default MSS clamping calculations are as follows:
 - If MSS minus 40 is greater than 416, subtract 40 from the MSS.
 - If MSS minus 40 is less than or equal to 416, set the MSS to 416.

To configure the MSS Clamping feature, use the following command at the global configuration level:

```
ds-lite tcp mss-clamp {none | fixed  $n$  | subtract  $s$  [min  $n$ ]}
```

Note: For more details please consult the *IPv4-to-IPv6 Transition Solutions Guide*.

2.2.7 LOGGING

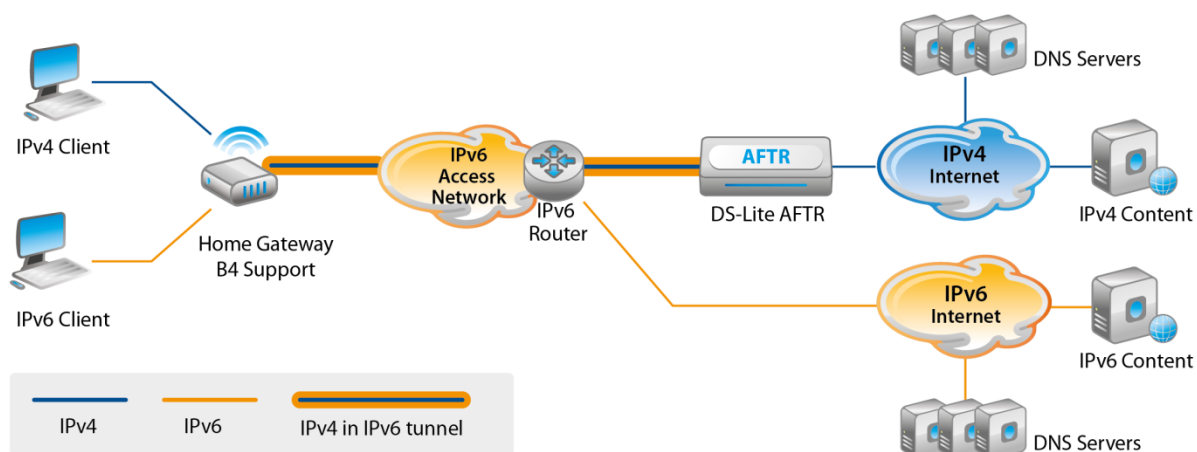
The AX DS-Lite implementation supports all CGN log-file size and volume reduction logging features including Compact (hex), Binary, and RFC 5424 formats; and Log Batching. Standard Syslog is supported, as well as Syslog over TCP, Netflow v9/v10¹ and RADIUS logging. Please refer to the “Logging Configuration” section of the [Carrier Grade NAT \(CGN\) Deployment Guide](#) for more operational details and configuration guidelines.

The following example illustrates DS-Lite session creation and port-mapping logs in ASCII format:

```
Nov  1 11:24:45 10.240.240.241 Nov  1 14:28:26 AX3030-T-47 NAT-UDP-N:
[2001:db8::2:1]10.1.1.1:57238<-->[2001:db8::1:1]66.18.32.2: 3, 66.18.32.2:53<--
>203.0.113.1:57238
Nov  1 11:24:45 10.240.240.241 Nov  1 14:28:26 AX3030-T-47 NAT-UDP-C:
[2001:db8::2:1]10.1.1.1:57238 -> 203.0.113.1:57238
Nov  1 11:24:47 10.240.240.241 Nov  1 14:28:28 AX3030-T-47 NAT-UDP-D:
[2001:db8::2:1]10.1.1.1:57238<-->[2001:db8::1:1]66.18.32.2: 3, 66.18.32.2:53<--
>203.0.113.1:57238
Nov  1 11:24:47 10.240.240.241 Nov  1 14:28:28 AX3030-T-47 NAT-UDP-F:
[2001:db8::2:1]10.1.1.1:57238 -> 203.0.113.1:57238
```

2.2.8 NETWORK INTEGRATION

A10 Networks provides flexible options for integration of the AX series' DS-Lite features into the service provider's network. In the example shown in



, BGP is used on the public Internet-facing interfaces and OSPFv3 is configured for the access-network

¹ NetFlow v10 (IPFIX) is available as of ACOS 2.6.6-P4.

interfaces. For proper operation, the aggregation router must have a route to the AFTR IPv6 address, and the AFTR must have a route to the access network subnets containing the home gateways. For HA environments, the AFTR address is the floating IP address assigned to the HA group, and it must be redistributed to the aggregation router.

For the public Internet-facing interfaces, the AX must redistribute the NAT pool addresses to its upstream peering router. By default, the next hop of a NAT pool address is the physical interface address of the active AX device. For proper HA operation, the next-hop address must be modified to represent the floating IP address associated with the HA group.

The following commands configure redistribution of the floating IP and NAT pool addresses. Refer to [Figure 3](#).

1. Configure the floating IP addresses:

```
AX3030-T-47-Active (config) #floating-ip 100.64.3.1 ha-group 1
AX3030-T-47-Active (config) #floating-ip 2001:db8::1:1 ha-group 1
```

2. Redistribute the IPv6 floating IP address (AFTR tunnel end-point) to downstream aggregation routers:

```
AX3030-T-47-Active (config) #router ipv6 ospf 1
AX3030-T-47-Active (config-router) #redistribute floating-ip
```

3. Create a route map to modify the next-hop of NAT pool addresses sent to upstream peering routers:

```
AX3030-T-47-Active (config) #route-map outside_fip4 permit 1
AX3030-T-47-Active (config-route-map) #set ip next-hop 100.64.3.1
```

4. Redistribute NAT pool addresses:

```
AX3030-T-47-Active (config) #router bgp 65000
AX3030-T-47-Active (config-router) #redistribute ip-nat route-map outside_fip4
```

3 CONCLUSION

This DS-Lite Deployment Guide is part of a series IPv6 deployment guides. This deployment guide builds upon the first deployment guide, which covers Carrier Grade NAT(CGN/CGNAT)/Large Scale NAT(LSN). The examples in this guide were created on an AX device running ACOS version 2.6.6-P3.