

Technology Overview

Configuring Dual-Stack Lite for IPv6 Access

Release
13.1



Published: 2013-02-08

Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, California 94089
USA
408-745-2000
www.juniper.net

This product includes the Envoy SNMP Engine, developed by Epilogue Technology, an Integrated Systems Company. Copyright © 1986-1997, Epilogue Technology Corporation. All rights reserved. This program and its documentation were developed at private expense, and no part of them is in the public domain.

This product includes memory allocation software developed by Mark Moraes, copyright © 1988, 1989, 1993, University of Toronto.

This product includes FreeBSD software developed by the University of California, Berkeley, and its contributors. All of the documentation and software included in the 4.4BSD and 4.4BSD-Lite Releases is copyrighted by the Regents of the University of California. Copyright © 1979, 1980, 1983, 1986, 1988, 1989, 1991, 1992, 1993, 1994. The Regents of the University of California. All rights reserved.

GateD software copyright © 1995, the Regents of the University. All rights reserved. Gate Daemon was originated and developed through release 3.0 by Cornell University and its collaborators. Gated is based on Kirton's EGP, UC Berkeley's routing daemon (routed), and DCN's HELLO routing protocol. Development of Gated has been supported in part by the National Science Foundation. Portions of the GateD software copyright © 1988, Regents of the University of California. All rights reserved. Portions of the GateD software copyright © 1991, D. L. S. Associates.

This product includes software developed by Maker Communications, Inc., copyright © 1996, 1997, Maker Communications, Inc.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.

Technology Overview Configuring Dual-Stack Lite for IPv6 Access

Release 13.1

NCE0025

Copyright © 2013, Juniper Networks, Inc.

All rights reserved.

The information in this document is current as of the date on the title page.

YEAR 2000 NOTICE

Juniper Networks hardware and software products are Year 2000 compliant. Junos OS has no known time-related limitations through the year 2038. However, the NTP application is known to have some difficulty in the year 2036.

END USER LICENSE AGREEMENT

The Juniper Networks product that is the subject of this technical documentation consists of (or is intended for use with) Juniper Networks software. Use of such software is subject to the terms and conditions of the End User License Agreement ("EULA") posted at <http://www.juniper.net/support/eula.html>. By downloading, installing or using such software, you agree to the terms and conditions of that EULA.

Table of Contents

Introduction	1
Overview of Dual-Stack Lite	1
DS-Lite Implementation	2
Transition of IPv4 Traffic to IPv6 Addresses Using Dual-Stack Lite	3
Example: Configuring Dual-Stack Lite for IPv6 Access	3
Redundancy and Load Balancing Using IPv6 Anycast Addresses	13
Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs	14
Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs	27

Introduction

This document describes dual-stack lite (DS-Lite), a technology that enables Internet service providers (ISPs) to move to an IPv6 network while simultaneously handling IPv4 address depletion. It also provides step-by-step configuration examples for configuring DS-Lite for IPv6 access, configuring redundancy and load-balancing using a single DS-Lite Address Family Transition Router (AFTR), and configuring redundancy with two or more DS-Lite AFTRs using a single IPv6 anycast address.

Overview of Dual-Stack Lite

Because IPv4 addresses are becoming depleted, broadband service providers (DSL, cable, and mobile) need new addresses to supply new customers. Providing IPv6 addresses alone is often not workable because most of the systems that make up the public Internet are still enabled to support only IPv4, and many customer systems do not yet fully support IPv6.

Dual-stack lite (DS-Lite) provides one solution to this problem for Internet service providers (ISPs). DS-Lite allows an ISP to migrate to an IPv6 access network without changing end-user software. The device that accesses the Internet remains the same.

The DS-Lite architecture uses IPv6-only links between the provider and the customer while maintaining the IPv4 (or dual-stack) hosts in the customer network.

When a customer's device sends an IPv4 packet to an external destination, the IPv4 packet is encapsulated in an IPv6 packet for transport into the provider network. These IPv4-in-IPv6 tunnels are called *softwires*. Tunneling IPv4 over IPv6 is simpler than translation and eliminates performance and redundancy concerns.

The softwires terminate in a softwire concentrator in the service provider network, which decapsulates the IPv4 packets and performs Network Address Translation (NAT). The packets undergo source-NAT processing to hide the original source address.

IPv6 packets originated by hosts in the subscriber's home network are transported natively over the access network.

The IPv4 packets originated by the end hosts have private (and possibly overlapping) IP addresses. Therefore, NAT must be applied to these packets. If end hosts have overlapping addresses, Network Address Port Translation (NAPT) is needed.

When using NAPT, the system uses an algorithm that takes the IPv6 packet's source address, private IPv4 address, and port to map the IPv4 packet to a unique combination of an IPv4 public address and port. Because each customer's IPv6 address is unique, the combination of the IPv6 source address with the IPv4 source address and port creates an unambiguous mapping.

The system takes the following actions when it receives a responding IPv4 packet from outside the subscriber network:

- Matches the IPv4 destination address and port for the packet to a specific customer based on the IPv6 address in the mapping table
- Maps the packet's IPv4 destination address and port to the IPv4 destination address and port inside the subscriber network
- Encapsulates the IPv4 packet in an IPv6 packet using the mapped IPv6 address as the IPv6 destination address
- Forwards the packet to the customer

For more information, see the following documents:

- Internet draft draft-ietf-softwire-dual-stack-lite-06, *Dual-Stack Lite Broadband Deployments Following IPv4 Exhaustion*, August 2010.
- RFC 2473, *Generic Packet Tunneling in IPv6 Specification*, December 1998.
- RFC 2663, *IP Network Address Translator (NAT) Terminology and Considerations*, August 1999.
- RFC 4787, *Network Address Translation (NAT) Behavioral Requirements for Unicast UDP, BCP 127*, January 2007.
- RFC 4925, *Softwire Problem Statement*, July 2007.
- RFC 5382, *NAT Behavioral Requirements for TCP, BCP 142*, October 2008.
- RFC 5508, *NAT Behavioral Requirements for ICMP, BCP 148*, April 2009.
- <http://www.potaroo.net/tools/ipv4/index.html>
- <http://www.iana.org/assignments/ipv4-address-space/ipv4-address-space.xhtml>

DS-Lite Implementation

In Juniper Networks® Junos® operating systems (Junos OS) Release 10.4 and later, Juniper Networks has implemented an Address Family Transition Router (AFTR) in its Services Physical Interface Cards (PICs) and Services Dense Port Concentrators (DPCs). An AFTR consists of the combination of an IPv4-in-IPv6 tunnel end-point and an IPv4-IPv4 NAT implemented on the same device.

A Basic Bridging Broadband Elements (B4 or a softwire initiator) is a function implemented on a dual-stack capable node, either as a directly connected device or a home gateway that creates a tunnel to an AFTR. IPv6 packets destined for the softwire concentrator's address are sent to a Services PIC, where the system creates a softwire according to the configuration. The system then extracts the IPv4 packets, performs NAT rule lookup and address translation, and sends the translated IPv4 packets to the Internet. The system performs these functions in a single pass through the Services PIC.

In the reverse path, the system sends IPv4 packets to the Services PIC, where the reverse translation is performed. The resulting packet is encapsulated in an IPv6 packet corresponding to the proper softwire and sent to the B4.

The system automatically creates softwires as IPv6 packets are received. IPv4 flows created by the encapsulated packets are associated with the specific software that initially carried them. When the last IPv4 flow associated with a software is completed, the software itself goes away. Thus, there is no need to create or manage tunnel interfaces, which simplifies the configuration.

The number of established softwires does not affect throughput, and scalability is independent of the number of interfaces.

Transition of IPv4 Traffic to IPv6 Addresses Using Dual-Stack Lite

ISPs can use DS-Lite to migrate over to an IPv6 access network without changing end-user software. Customers can still access the Internet using their current hardware. DS-Lite accomplishes this by encapsulating the IPv4 packets originated by the existing end hosts into IPv6 packets.

DS-Lite enables an IPv4 host to communicate with a NAT endpoint over an IPv6 network using softwires. DS-Lite configurations create the IPv6 softwires, which terminate on the Services PIC. You can also configure the Services PIC to apply other services such as NAT on the packets that exit from the software. The aim of this implementation is to enable packets to travel over softwires to a carrier-grade NAT endpoint where they undergo source-NAT processing to hide the original source address.

Currently, a NAT rule configuration is required with a software configuration. NAT processing from IPv4 to IPv6 address pools and vice versa is not currently supported. The application-level gateway (ALG) currently supports HTTP, FTP, RSTP, and ICMP.

DS-Lite is supported on Juniper Networks M Series Multiservice Edge Routers and T Series Core Routers with Multiservices 100, 400, and 500 PICs and MX Series 3D Universal Edge Routers with Multiservices DPCs.

Related Documentation

- [Redundancy and Load Balancing Using IPv6 Anycast Addresses on page 13](#)
- [Stateful NAT64 Overview](#)
- [Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs on page 27](#)
- [Example: Configuring Dual-Stack Lite for IPv6 Access on page 3](#)
- [Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs on page 14](#)

Example: Configuring Dual-Stack Lite for IPv6 Access

This example shows how to configure DS-Lite for IPv6 access.

- [Requirements on page 4](#)
- [Overview on page 4](#)
- [Configuration on page 5](#)
- [Verification on page 10](#)

Requirements


This example uses the following hardware and software components:

- Juniper Networks MX Series 3D Universal Edge Routers with Multiservices Dense Port Concentrators (DPCs)
- Juniper Networks® Junos® operating system (Junos OS) 10.4 or later running on the Address Family Transition Routers (AFTRs)

Overview

In [Figure 1 on page 4](#), the AFTR is running on an MX Series router with two Gigabit Ethernet interfaces and a Multiservices DPC. The interface toward the Basic Bridging BroadBand Element (B4) is **ge-3/1/5**, and the interface toward the Internet is **ge-3/1/0**.

Figure 1: Logical Topology



In [Figure 1 on page 4](#):

- The source IPv4 address connected to the home router is 10.0.0.1.
- The source address (or B4 interface address) of the IPv4-in-IPv6 software is 2001:0:0:1::1.
- The address of the NAT pool between the AFTR and the Internet is 129.0.0.1.
- The address of the IPv4 host connected to the Internet is 128.0.0.1.
- The address of the software on the AFTR is 2001:0:0:2::1/48.

Configuration

Configuring DS-Lite involves the following tasks:

- [Enabling the Layer-3 Service Package on page 5](#)
- [Configuring Network Address and Port Translation on page 6](#)
- [Configuring the Softwire Concentrator on page 6](#)
- [Configuring the Service Set with Softwire and NAT Rules on page 7](#)
- [Configuring Interfaces and Associating Service Sets with the Interfaces on page 8](#)

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
AFTR set chassis fpc 0 pic 0 adaptive-services service-package layer-3
set interfaces sp-0/0/0 unit 0 family inet
set interfaces sp-0/0/0 unit 0 family inet6
set interfaces ge-3/1/0 description AFTR-Internet
set interfaces ge-3/1/0 unit 0 family inet address 128.0.0.2/24
set interfaces ge-3/1/5 description AFTR-B4
set interfaces ge-3/1/5 unit 0 family inet
set interfaces ge-3/1/5 unit 0 family inet6
set interfaces ge-3/1/5 unit 0 family inet6 service input service-set sset
set interfaces ge-3/1/5 unit 0 family inet6 service output service-set sset
set interfaces ge-3/1/5 unit 0 family inet6 address 2001:0:0:2::1/48
set services service-set sset syslog host local services any
set services service-set sset softwire-rules r1
set services service-set sset tcp-mss 1024
set services service-set sset nat-rules r1
set services service-set sset interface-service service-interface sp-0/0/0.0
set services softwire softwire-concentrator ds-lite ds1 softwire-address 1001::1
set services softwire softwire-concentrator ds-lite ds1 mtu-v6 1460
set softwire softwire-concentrator ds-lite ds1 copy-dscp
set softwire softwire-concentrator ds-lite ds1 flow-limit 10
set services softwire rule r1 match-direction input
set services softwire rule r1 term t1 then ds-lite ds1
set services nat pool p1 address 129.0.0.1/32
set services nat pool p1 port automatic
set services nat rule r1 match-direction input
set services nat rule r1 term t1 from source-address 10.0.0.0/16
set services nat rule r1 term t1 then translated source-pool p1
set services nat rule r1 term t1 then translated translation-type napt-44
set services nat rule r1 term t1 then syslog
```

Enabling the Layer-3 Service Package

Step-by-Step Procedure The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the CLI User Guide.

- Configure the Layer 3 service package.

This example assumes that the PIC is in FPC 0, slot 0.

```
[edit chassis]
```

```
user@AFTR# set fpc 0 pic 0 adaptive-services service-package layer-3
```

The service package with its associated **sp-** interface is for manipulating traffic before it is delivered to its destination. For details about configuring service packages, see the *Junos OS Services Interfaces Configuration Guide*.

Configuring Network Address and Port Translation

Step-by-Step Procedure

To configure NAT and Port Address Translation (PAT) rules:

1. Configure an IPv4 address and port for the NAT pool to specify the IPv4-to-IPv6 translation for packets traveling between the AFTR router and the Internet.

```
[edit services nat]
```

```
user@AFTR# set pool p1 address 129.0.0.1/32
```

```
user@AFTR# set pool p1 port automatic
```

2. Configure a NAT rule to translate the private IPv4 address from the home network to NAT **pool p1**.

NAT rules specify the traffic to be matched and the action to be taken when traffic matches the rule. In this example, only one rule is required to accomplish the address translation. The rule selects all traffic coming from the source address **10.0.0.0**.

```
[edit services nat]
```

```
user@AFTR# set rule r1 match-direction input
```

```
user@AFTR# set rule r1 term t1 from source-address 10.0.0.0/16
```

```
user@AFTR# set rule r1 term t1 then translated source-pool p1
```

```
user@AFTR# set rule r1 term t1 then translated translation-type napt-44
```

```
user@AFTR# set rule r1 term t1 then syslog
```

Configuring the Software Concentrator

Step-by-Step Procedure

1. Create a software concentrator object of type **ds-lite** and associate it with the IPv6 address of the software.

Specify a name for the software concentrator to facilitate references in logs, in the CLI, and in other operations and management activities.

```
[edit services software]
```

```
user@AFTR# set software-concentrator ds-lite ds1 software-address 1001::1
```

2. Configure the maximum transmission unit (ranging from 1280 to 9192 bytes) for the software for encapsulating IPv4 packets to IPv6.

This is the maximum packet size that can be sent on a tunnel from the AFTR to B4 without fragmentation. If the final length of the packet is greater than the MTU, the IPv6 packet would be fragmented.



NOTE: Including the **mtu-v6** statement is mandatory, and you cannot commit the example configuration unless this statement is configured.

```
[edit services software]
user@AFTR# set software-concentrator ds-lite ds1 mtu-v6 1460
```

3. (Optional) Configure the software to copy DSCP information from the IPv6 header into the decapsulated IPv4 header.

```
[edit services software]
user@AFTR# set software-concentrator ds-lite ds1 copy-dscp
```

4. (Optional) Configure a flow limit for the maximum number of IPv4 flows or sessions (ranging from 1280 to 9192 bytes) per software from the IPv4 host to the B4.

```
[edit services software]
user@AFTR# set software-concentrator ds-lite ds1 flow-limit 10
```

5. Create a software rule.

The rule in this example specifies that any traffic destined for the software concentrator **ds1** creates a new software. You can also configure more elaborate match conditions to perform as part of software initiator actions.

```
[edit services software]
user@AFTR# set rule r1 match-direction input
user@AFTR# set rule r1 term t1 then ds-lite ds1
```

Configuring the Service Set with Software and NAT Rules

Step-by-Step Procedure

To configure the service set on service interface **sp-0/0/0** to contain the software and NAT rules:

1. Configure a service set using the same NAT and software rules configured in the previous two procedures.

```
[edit services]
user@AFTR# set service-set sset software-rules r1
user@AFTR# set service-set sset nat-rules r1
user@AFTR# set service-set sset interface-service service-interface sp-0/0/0.0
```

2. Configure the service interface.

In this example, the interface is **sp-0/0/0**.

```
[edit interfaces]
user@AFTR# set sp-0/0/0 unit 0 family inet
user@AFTR# set sp-0/0/0 unit 0 family inet6
```

3. (Optional) Configure a TCP maximum segment size value on a service-set basis to ensure that TCP traffic works through links with different MTUs.

```
[edit services]
user@AFTR# set service-set sset tcp-mss 1024
```

4. Associate the software and NAT rules and the service interface with the service set.

```
[edit services]
user@AFTR# set service-set sset interface-service service-interface sp-0/0/0.0
user@AFTR# set service-set sset software-rules r1
user@AFTR# set service-set sset nat-rules r1
```

5. Configure system log parameters for the service set.

```
[edit services]
user@AFTR# set service-set sset syslog host local services any
```

Configuring Interfaces and Associating Service Sets with the Interfaces

Step-by-Step Procedure

1. Configure the **ge-3/1/5** interface between the home router running the B4 and the router in the ISP network running the AFTR.

```
[edit interfaces]
user@AFTR# set ge-3/1/5 description AFTR-B4
user@AFTR# set ge-3/1/5 unit 0 family inet
user@AFTR# set ge-3/1/5 unit 0 family inet6
```



NOTE: Even if IPv6 packets are received on the media interface, configuring family inet on this interface is very important for the DS-Lite configuration to work properly.

2. Associate the appropriate service set for the NAT and DS-Lite services.
Service sets can be configured in either interface style or nexthop VPN routing and forwarding (VRF) style. This example depicts the interface style configuration.

```
[edit interfaces]
user@AFTR# set ge-3/1/5 unit 0 family inet6 service input service-set sset
user@AFTR# set ge-3/1/5 unit 0 family inet6 service output service-set sset
```

3. Include the IPv6 software address of the AFTR router.

```
[edit interfaces]
user@AFTR# set ge-3/1/5 unit 0 family inet6 address 2001:0:0:2::1/48
```

4. Configure the **ge-3/1/0** interface between the AFTR and the Internet, and specify the IPv4 address connected to the Internet.

```
[edit interfaces]
user@AFTR# set ge-3/1/0 description AFTR-Internet
user@AFTR# set ge-3/1/0 unit 0 family inet address 128.0.0.2/24
```

Results In configuration mode, confirm your configuration by entering the **show chassis**, **show services nat**, **show services software**, **show services service-set**, and **show interfaces** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@AFTR# show chassis
fpc 0 {
  pic 0 {
    adaptive-services {
      service-package layer-3;
    }
  }
}

user@AFTR# show services nat
pool p1 {
```

```
address 129.0.0.1/32;
port {
    automatic;
}
}
rule r1 {
    match-direction input;
    term t1 {
        from {
            source-address {
                10.0.0.0/16;
            }
        }
        then {
            translated {
                source-pool p1;
                translation-type {
                    napt-44;
                }
            }
        }
        syslog;
    }
}
}
```

user@AFTR# show services software

```
software-concentrator {
    ds-lite ds1 {
        software-address 1001::1;
        mtu-v6 1460;
    }
    rule r1 {
        match-direction input;
        term t1 {
            then {
                ds-lite ds1;
            }
        }
    }
}
```

user@AFTR# show services service-set sset

```
syslog {
    host local {
        services any;
    }
    software-rules r1;
    nat-rules r1;
    interface-service {
        service-interface sp-0/0/0.0;
    }
}
```

user@AFTR# show interfaces

```
sp-0/0/0 {
    unit 0 {
        family inet;
```

```

        family inet6;
    }
}
ge-3/1/0 {
    description AFTR-Internet;
    unit 0 {
        family inet {
            address 128.0.0.2/24;
        }
    }
}
ge-3/1/5 {
    description AFTR-B4;
    unit 0 {
        family inet;
        family inet6 {
            service {
                input {
                    service-set sset;
                }
                output {
                    service-set sset;
                }
            }
            address 2001:0:0:2::1/48;
        }
    }
}
}

```

If you are done configuring the device, enter **commit** from configuration mode.

Verification

Confirm that the configuration is working properly.

- [Verifying Softwires on page 10](#)
- [Verifying NAT Flows on page 11](#)
- [Verifying Traceroute on page 13](#)

Verifying Softwires

Purpose Verify the creation of the softwires.

- Action** 1. Issue the **show services software** command to view information about the softwires created.

```

user@AFTR> show services software
Interface: sp-0/0/0, Service set: sset
Software
2001::3      ->      1001::1      Direction      Flow count
                        I                        3

Interface: sp-1/3/0, Service set: dslite-svc-set1
Software
2001::2      ->      1001::1      Direction      Flow count
2001::4      ->      1001::1      I                        3

```

-
2. Issue the **show services software statistics ds-lite** command to view details of global software statistics.

```
user@AFTR> show services software statistics ds-lite
DS-Lite Statistics:
```

```
Service PIC Name:                               :sp-0/0/0
```

```
Statistics
-----
```

```
Softwires Created                               :6
Softwires Deleted                               :6
Softwires Flows Created                         :6
Softwires Flows Deleted                        :6
Slow Path Packets Processed                     :6
Fast Path Packets Processed                     :21
Fast Path Packets Encapsulated                  :20
Rule Match Succeeded                           :6
Rule Match Failed                             :0
IPv6 Packets Fragmented                        :0
IPv4 Client Fragments                          :0
ICMPv4 Error Packets sent                      :0
ICMPv6 Packets sent                            :0
```

```
Transient Errors
-----
```

```
Flow Creation Failed - Retry                    :0
Slow Path Failed - Retry                       :0
```

```
Errors
-----
```

```
Software Creation Failed                       :0
Flow Creation Failed                           :0
Slow Path Failed                              :0
Packet not IPv4-in-IPv6                       :0
IPv6 Fragmentation Error                      :0
Slow Path Failed - IPv6 Next Header Offset    :0
Decapsulated Packet not IPv4                  :0
Fast Path Failed - IPv6 Next Header Offset    :0
No Software ID                                :0
No Flow Extension                             :0
Flow Limit Exceeded                           :0
```

Verifying NAT Flows

Purpose Verify pre-NAT and post-NAT flows.

- Action**
1. On the host router, issue the **show services stateful-firewall flows** command to verify the creation of the softwires, pre-NAT flows, and post-NAT flows within the configuration.

```
user@AFTR> show services stateful-firewall flows
```

```
Interface: sp-0/0/0, Service set: sset
```

Flow	State	Dir	Frm
count			

```

TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I          107621
      NAT source      20.20.1.2:1025  ->    44.44.44.1:1024
      Software        2001::3      ->    1001::1
TCP          200.200.200.2:80  ->    44.44.44.1:1024  Forward  O          208420
      NAT dest        44.44.44.1:1024  ->    20.20.1.2:1025
      Software        2001::3      ->    1001::1
DS-LITE      2001::3      ->    1001::1      Forward  I          322166

```

In this example:

- In the output direction (O), the protocol (TCP) line shows the Internet-to-IPv4 host address translated to the address of the AFTR.
 - In the output direction, the NAT-translated IPv4 address is translated to the IPv4 address of the home host (NAT dest).
 - In the output direction, the IPv6 address of the B4 is translated to the IPv6 address of the AFTR (Software).
 - In the input direction (I), the protocol (TCP) line shows the address of the home host sending the packet to the address of the Internet-to-IPv4 host.
 - In the input direction, the IPv6 address of the B4 is translated to the IPv6 address of the AFTR (NAT source).
2. Issue the **show services stateful-firewall conversations** command to verify the conversations (collections of related flows).

```
user@AFTR> show services stateful-firewall conversations
```

```
Interface: sp-0/0/0, Service set: sset
```

```
Conversation: ALG protocol: tcp
```

```

      Number of initiators: 1, Number of responders: 1
Flow count      State  Dir  Frm
TCP          20.20.1.2:1025  ->  200.200.200.2:80    Forward  I          189280
      NAT source      20.20.1.2:1025  ->    44.44.44.1:1024
      Software        2001::3      ->    1001::1
TCP          200.200.200.2:80  ->    44.44.44.1:1024  Forward  O          363675
      NAT dest        44.44.44.1:1024  ->    20.20.1.2:1025
      Software        2001::3      ->    1001::1

```

3. Issue the **show services nat pool detail** command to display global NAT statistics related to pool usage.

You normally use this command in conjunction with the **show services stateful-firewall flows** command, which displays the source and output of the translation.

```
user@AFTR> show services nat pool detail
```

```
Interface: sp-0/0/0, Service set: sset
```

```

      NAT pool: p1, Translation type: dynamic
      Address range: 129.0.0.1-129.0.0.1
      Port range: 512-65535, Ports in use: 16, Out of port errors: 0, Max ports
      used: 17

```

Verifying Traceroute

Purpose Examine the traceroute from the IPv4 host on the home network to the IPv4 node on the Internet.

Action Examine the traceroute.

The following output of a traceroute from the client, to the home host, to the IPv4 host on the Internet is based on [“Configuring the Software Concentrator” on page 6](#).

```
user@AFTR> show services stateful-firewall flows
```

```
Interface: sp-0/0/0, Service set: sset
Flow
ICMP      10.0.0.1      -> 128.0.0.1      State Dir  Frm count
          NAT source  10.0.0.1      -> 129.0.0.1      Watch I    4
          Softwire    2001:0:0:1::1 -> 2002:0:0:2::1
ICMP      128.0.0.1    -> 129.0.0.1      State Dir  Frm count
          NAT dest   129.0.0.1     -> 10.0.0.1      Watch 0    1
          Softwire    2001:0:0:1::1 -> 2002:0:0:2::1
DS-LITE    2001:0:0:1::1 -> 2002:0:0:2::1 Forward I    322166
```



NOTE: If a traceroute starts from the home host and goes to an IPv4 host on the Internet, the software concentrator does not return an ICMP error and, therefore, is not properly identified as an intermediate hop. However, the traceroute still functions.

Meaning The ICMP source and destination addresses in the output indicate that the traffic is flowing from the IPv4 host on the home network (10.0.0.1) to the IPv4 node on the Internet (128.0.0.1).

- Related Documentation**
- [Overview of Dual-Stack Lite on page 1](#)
 - [Redundancy and Load Balancing Using IPv6 Anycast Addresses on page 13](#)
 - [Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs on page 27](#)
 - [Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs on page 14](#)
 - [Example: Configuring Stateful NAT64 for Handling IPv4 Address Depletion](#)

Redundancy and Load Balancing Using IPv6 Anycast Addresses

An IPv6 anycast address is like any unicast address that is assigned to more than one interface (typically belonging to different nodes) where a packet sent to an anycast address is routed using the preferred route from the routing table. Anycasting is commonly used to load-balance between geographically dispersed servers.

For DS-Lite, redundancy and load balancing can be accomplished by:

- Using multiple Services Physical Interface Cards (PICs) on the same Address Family Transition Router (AFTR) and a single anycast address for the software address.
- Configuring the same anycast address on multiple AFTRs.

The advantage of configuring anycast addresses for AFTRs is that Internet Service Providers (ISPs) get service continuity and load balancing. In addition, the Basic Bridging BroadBand Element (B4) or the software initiator only needs to know a single AFTR IPv6 address.

**Related
Documentation**

- [Overview of Dual-Stack Lite on page 1](#)
- [Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs on page 27](#)
- [Example: Configuring Dual-Stack Lite for IPv6 Access on page 3](#)
- [Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs on page 14](#)

Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs

This example shows how to configure redundancy and load balancing using a single DS-Lite Address Family Transition Router (AFTR).

- [Requirements on page 14](#)
- [Overview on page 15](#)
- [Configuration on page 16](#)
- [Verification on page 23](#)

Requirements


This example uses the following hardware and software components:

- Juniper Networks MX Series 3D Universal Edge Routers with Multiservices Dense Port Concentrators (DPCs)
- Juniper Networks® Junos® operating system (Junos OS) 10.4 or later running on the AFTR

Overview

You can provide redundancy and load balancing using multiple Services PICs on the same AFTR and a single anycast address where the two Services PICs actively load-balance traffic. In [Figure 2 on page 15](#), three Basic Bridging BroadBand Elements (B4s or softwire initiators) are connected to the AFTR's softwire (ID 1001::1) using different tunnels. The AFTR has two services for load balancing and redundancy. When HTTP clients connect to the server, traffic is load-balanced between the Services PICs. In addition, when one of the Services PICs is down, traffic from all three B4s is channelized through the other Services PIC.

Figure 2: Sample Topology for DS-Lite Anycast Configuration Using Multiple Services PICs



- The IPv4 client or host in the home network is configured with an IPv4 interface to the ISP and a static route to the IPv4 server on the Internet.
- The multiple B4s or softwire initiators are configured with an IPv4 interface, an IPv6 interface, and an IPv4-in-IPv6 tunnel to an anycast address.
- The pure IPv6 node in the IPv6 cloud is configured with interfaces to the IPv6 interfaces.

- The address range of the NAT pool between the AFTR and the Internet is 33.33.33.1 through 33.33.33.32 corresponding to NAT rule **dslite-nat-rule1**, and 44.44.44.1 through 44.44.44.32 corresponding to NAT rule **dslite-nat-rule2**.
- NAT rule **dslite-nat-rule1** corresponds to Services PIC **sp-0/1/0**, and NAT rule **dslite-nat-rule2** corresponds to Services PIC **sp-1/3/0**.
- The AFTR is configured with anycast address 2001::1/16 for the interface toward the three B4s. Address 200.200.200.1/24 is configured for the interface from the AFTR toward the Internet. The two Services PICs are **sp-0/1/0** and **sp-1/3/0**.
- The IPv4 node on the Internet is configured with an IPv4 interface and routes for reverse traffic.

Configuration

CLI Quick Configuration To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
AFTR
set chassis fpc 1 pic 1 adaptive-services service-package layer-3
set services nat pool dslite-pool1 address-range low 33.33.33.1 high 33.33.33.32
set services nat pool dslite-pool1 port automatic
set services nat pool dslite-pool2 address-range low 44.44.44.1 high 44.44.44.32
set services nat pool dslite-pool2 port automatic
set services nat rule dslite-nat-rule1 match-direction input
set services nat rule dslite-nat-rule1 term t1 from source-address 20.20.0.0/16
set services nat rule dslite-nat-rule1 term t1 then translated source-pool dslite-pool1
set services nat rule dslite-nat-rule1 term t1 then translated translation-type napt-44
set services nat rule dslite-nat-rule2 match-direction input
set services nat rule dslite-nat-rule2 term t1 from source-address 20.20.0.0/16
set services nat rule dslite-nat-rule2 term t1 then translated source-pool dslite-pool2
set services nat rule dslite-nat-rule2 term t1 then translated translation-type napt-44
set services softwire softwire-concentrator ds-lite ds1 softwire-address 1001::1
set services softwire softwire-concentrator ds-lite ds1 mtu-v6 9192
set services softwire rule dslite-rule match-direction input
set services softwire rule dslite-rule term t1 then ds-lite ds1
set services service-set dslite-svc-set1 syslog host local services any
set services service-set dslite-svc-set1 softwire-rules dslite-rule
set services service-set dslite-svc-set1 stateful-firewall-rules sfw-r1
set services service-set dslite-svc-set1 nat-rules dslite-nat-rule1
set services service-set dslite-svc-set1 next-hop-service inside-service-interface sp-0/1/0.1
set services service-set dslite-svc-set1 next-hop-service outside-service-interface
  sp-0/1/0.2
set services service-set dslite-svc-set2 syslog host local services any
set services service-set dslite-svc-set2 softwire-rules dslite-rule
set services service-set dslite-svc-set2 stateful-firewall-rules sfw-r1
set services service-set dslite-svc-set2 nat-rules dslite-nat-rule2
set services service-set dslite-svc-set2 next-hop-service inside-service-interface sp-1/3/0.1
set services service-set dslite-svc-set2 next-hop-service outside-service-interface
  sp-1/3/0.2
set services stateful-firewall rule sfw-r1 match-direction input
set services stateful-firewall rule sfw-r1 term t1 from applications junos-http
set services stateful-firewall rule sfw-r1 term t1 from applications junos-ftp
```

```

set services stateful-firewall rule sfw-r1 term t1 from applications junos-rtsp
set services stateful-firewall rule sfw-r1 term t1 from applications junos-icmp-all
set services stateful-firewall rule sfw-r1 term t1 then accept
set services stateful-firewall rule sfw-r1 term t1 then syslog
set interfaces ge-0/0/2 unit 0 family inet
set interfaces ge-0/0/2 unit 0 family inet6 address 2001::1/16
set interfaces ge-0/0/3 unit 0 family inet address 200.200.200.1/24
set interfaces sp-0/1/0 services-options syslog host local services any
set interfaces sp-0/1/0 unit 0 family inet
set interfaces sp-0/1/0 unit 0 family inet6
set interfaces sp-0/1/0 unit 1 family inet6
set interfaces sp-0/1/0 unit 1 service-domain inside
set interfaces sp-0/1/0 unit 2 family inet6
set interfaces sp-0/1/0 unit 2 service-domain outside
set interfaces sp-1/3/0 services-options syslog host local services any
set interfaces sp-1/3/0 unit 0 family inet
set interfaces sp-1/3/0 unit 0 family inet6
set interfaces sp-1/3/0 unit 1 family inet6
set interfaces sp-1/3/0 unit 1 service-domain inside
set interfaces sp-1/3/0 unit 2 family inet6
set interfaces sp-1/3/0 unit 2 service-domain outside
set routing-options forwarding-table export load-balancing-policy
set policy-options policy-statement load-balancing-policy then load-balance per-packet
set routing-options rib inet6.0 static route 1001::1/128 next-hop sp-1/3/0.1
set routing-options rib inet6.0 static route 1001::1/128 next-hop sp-0/1/0.1
set forwarding-options hash-key family inet6 layer-3 destination-address
set forwarding-options hash-key family inet6 layer-3 source-address

```

Configuring the AFTR

Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the CLI User Guide.

1. Configure the Layer 3 service package.
This example assumes that the PIC is in FPC 1, slot 1.

[edit chassis]

```
user@AFTR# set fpc 1 pic 1 adaptive-services service-package layer-3
```

The service package with its associated **sp-** interface is for manipulating traffic before it is delivered to its destination. For details about configuring service packages, see the *Junos OS Services Interfaces Configuration Guide*.

2. Configure two different NAT pools and NAT for the two Services PICs.

[edit services nat]

```

user@AFTR# set pool dslite-pool1 address-range low 33.33.33.1 high 33.33.33.32
user@AFTR# set pool dslite-pool1 port automatic
user@AFTR# set pool dslite-pool2 address-range low 44.44.44.1 high 44.44.44.32
user@AFTR# set pool dslite-pool2 port automatic
user@AFTR# set rule dslite-nat-rule1 match-direction input
user@AFTR# set rule dslite-nat-rule1 term t1 from source-address 20.20.0.0/16
user@AFTR# set rule dslite-nat-rule1 term t1 then translated source-pool dslite-pool1
user@AFTR# set rule dslite-nat-rule1 term t1 then translated translation-type
napt-44
user@AFTR# set rule dslite-nat-rule2 match-direction input

```

```

user@AFTR# set rule dslite-nat-rule2 term t1 from source-address 20.20.0.0/16
user@AFTR# set rule dslite-nat-rule2 term t1 then translated source-pool
dslite-pool2
user@AFTR# set rule dslite-nat-rule2 term t1 then translated translation-type
napt-44

```

3. Configure the software concentrator and create the software rule.

```

[edit services software]
user@AFTR# set software-concentrator ds-lite ds1 software-address 1001::1
user@AFTR# set software-concentrator ds-lite ds1 mtu-v6 9192
user@AFTR# set rule dslite-rule match-direction input
user@AFTR# set rule dslite-rule term t1 then ds-lite ds1

```

4. Configure next-hop-style service sets **dslite-svc-set1** and **dslite-svc-set2** for Services PICs **sp-0/1/0** and **sp-1/3/0**, respectively.

```

[edit services]
user@AFTR# set service-set dslite-svc-set1 syslog host local services any
user@AFTR# set service-set dslite-svc-set1 software-rules dslite-rule
user@AFTR# set service-set dslite-svc-set1 stateful-firewall-rules sfw-r1
user@AFTR# set service-set dslite-svc-set1 nat-rules dslite-nat-rule1
user@AFTR# set service-set dslite-svc-set1 next-hop-service inside-service-interface
sp-0/1/0.1
user@AFTR# set service-set dslite-svc-set1 next-hop-service
outside-service-interface sp-0/1/0.2
user@AFTR# set service-set dslite-svc-set2 syslog host local services any
user@AFTR# set service-set dslite-svc-set2 software-rules dslite-rule
user@AFTR# set service-set dslite-svc-set2 stateful-firewall-rules sfw-r1
user@AFTR# set service-set dslite-svc-set2 nat-rules dslite-nat-rule2
user@AFTR# set service-set dslite-svc-set2 next-hop-service inside-service-interface
sp-1/3/0.1
user@AFTR# set service-set dslite-svc-set2 next-hop-service
outside-service-interface sp-1/3/0.2

```

5. Configure stateful firewall and software rules.

```

[edit services]
user@AFTR# set stateful-firewall rule sfw-r1 match-direction input
user@AFTR# set stateful-firewall rule sfw-r1 term t1 from applications junos-http
user@AFTR# set stateful-firewall rule sfw-r1 term t1 from applications junos-ftp
user@AFTR# set stateful-firewall rule sfw-r1 term t1 from applications junos-rtsp
user@AFTR# set stateful-firewall rule sfw-r1 term t1 from applications junos-icmp-all
user@AFTR# set stateful-firewall rule sfw-r1 term t1 then accept
user@AFTR# set stateful-firewall rule sfw-r1 term t1 then syslog

```

6. Configure the services interfaces.

```

[edit interfaces]
user@AFTR# set sp-0/1/0 services-options syslog host local services any
user@AFTR# set sp-0/1/0 unit 0 family inet
user@AFTR# set sp-0/1/0 unit 0 family inet6
user@AFTR# set sp-0/1/0 unit 1 family inet6
user@AFTR# set sp-0/1/0 unit 1 service-domain inside
user@AFTR# set sp-0/1/0 unit 2 family inet6
user@AFTR# set sp-0/1/0 unit 2 service-domain outside
user@AFTR# set sp-1/3/0 services-options syslog host local services any
user@AFTR# set sp-1/3/0 unit 0 family inet
user@AFTR# set sp-1/3/0 unit 0 family inet6

```

-
- ```
user@AFTR# set sp-1/3/0 unit 1 family inet6
user@AFTR# set sp-1/3/0 unit 1 service-domain inside
user@AFTR# set sp-1/3/0 unit 2 family inet6
user@AFTR# set sp-1/3/0 unit 2 service-domain outside
```
7. Configure the interface between the home router running the B4 and the AFTR.  

```
[edit interfaces]
user@AFTR# set ge-0/0/2 unit 0 family inet
user@AFTR# set ge-0/0/2 unit 0 family inet6 address 2001::1/16
```
  8. Configure the interface between the AFTR and the Internet.  

```
[edit interfaces]
user@AFTR# set ge-0/0/3 unit 0 family inet address 200.200.200.1/24
```
  9. Configure load-balancing options for the Packet Forwarding Engine to determine how the traffic is load-balanced between the two Services PICs.  

```
[edit]
user@AFTR# set policy-options policy-statement load-balancing-policy then
 load-balance per-packet
user@AFTR# set routing-options forwarding-table export load-balancing-policy
```
  10. Configure routing options to install a route with high priority to the anycast address for both Services PICs.
    - a. Configure the static route destination address.
    - b. Configure the next hops to the destination address. Include the Services PICs (sp-1/3/0.1 sp-0/1/0.1) in the list of next hops.

```
[edit routing-options]
user@AFTR# set rib inet6.0 static route 1001::1/128 next-hop sp-1/3/0.1
user@AFTR# set rib inet6.0 static route 1001::1/128 next-hop sp-0/1/0.1
```
  11. Configure load-balancing options for the Packet Forwarding Engine.  

```
[edit forwarding-options]
user@AFTR# set hash-key family inet6 layer-3 destination-address
user@AFTR# set hash-key family inet6 layer-3 source-address
```

**Results** In configuration mode, confirm your configuration by entering the **show chassis**, **show services**, **show interfaces**, **show routing-options**, **show policy-options**, and **show forwarding-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@AFTR1# show chassis
fpc 1 {
 pic 1 {
 adaptive-services {
 service-package layer-3;
 }
 }
}

user@AFTR1# show services
service-set dslite-svc-set1 {
 syslog {
```

```
 host local {
 services any;
 }
}
software-rules dslite-rule;
stateful-firewall-rules sfw-r1;
nat-rules dslite-nat-rule1;
next-hop-service {
 inside-service-interface sp-0/1/0.1;
 outside-service-interface sp-0/1/0.2;
}
}
service-set dslite-svc-set2 {
 syslog {
 host local {
 services any;
 }
 }
 software-rules dslite-rule;
 stateful-firewall-rules sfw-r1;
 nat-rules dslite-nat-rule2;
 next-hop-service {
 inside-service-interface sp-1/3/0.1;
 outside-service-interface sp-1/3/0.2;
 }
}
stateful-firewall {
 rule sfw-r1 {
 match-direction input;
 term t1 {
 from {
 applications [junos-http junos-ftp junos-rtsp junos-icmp-all];
 }
 then {
 accept;
 syslog;
 }
 }
 }
}
software {
 software-concentrator {
 ds-lite ds1 {
 software-address 1001::1;
 mtu-v6 9192;
 }
 }
 rule dslite-rule {
 match-direction input;
 term t1 {
 then {
 ds-lite ds1;
 }
 }
 }
}
```



---

```
nat {
 pool dslite-pool1 {
 address-range low 33.33.33.1 high 33.33.33.32;
 port {
 automatic;
 }
 }
 pool dslite-pool2 {
 address-range low 44.44.44.1 high 44.44.44.32;
 port {
 automatic;
 }
 }
 rule dslite-nat-rule1 {
 match-direction input;
 term t1 {
 from {
 source-address {
 20.20.0.0/16;
 }
 }
 then {
 translated {
 source-pool dslite-pool1;
 translation-type {
 napt-44;
 }
 }
 }
 }
 }
 rule dslite-nat-rule2 {
 match-direction input;
 term t1 {
 from {
 source-address {
 20.20.0.0/16;
 }
 }
 then {
 translated {
 source-pool dslite-pool2;
 translation-type {
 napt-44;
 }
 }
 }
 }
 }
}

user@AFTR1# show interfaces
ge-0/0/2 {
 unit 0 {
 family inet;
 family inet6 {
```

```
 address 2001::1/16;
 }
}
ge-0/0/3 {
 unit 0 {
 family inet {
 address 200.200.200.1/24;
 }
 }
}
sp-0/1/0 {
 services-options {
 syslog {
 host local {
 services any;
 }
 }
 }
 unit 0 {
 family inet;
 family inet6;
 }
 unit 1 {
 family inet6;
 service-domain inside;
 }
 unit 2 {
 family inet6;
 service-domain outside;
 }
}
sp-1/3/0 {
 services-options {
 syslog {
 host local {
 services any;
 }
 }
 }
 unit 0 {
 family inet;
 family inet6;
 }
 unit 1 {
 family inet6;
 service-domain inside;
 }
 unit 2 {
 family inet6;
 service-domain outside;
 }
}

user@AFTR1# show routing-options
rib inet6.0 {
```

---

```
static {
 route 1001::1/128 next-hop [sp-1/3/0.1 sp-0/1/0.1];
}
forwarding-table {
 export load-balancing-policy;
}

user@AFTR1# show policy-options
policy-statement load-balancing-policy {
 then {
 load-balance per-packet;
 }
}

user@AFTR1# show forwarding-options
hash-key {
 family inet6 {
 layer-3 {
 destination-address;
 source-address;
 }
 }
}
```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

- [Verifying Load Balancing Between the Two Services PICs on page 23](#)
- [Verifying Redundancy Between the Two Services PICs on page 25](#)

### Verifying Load Balancing Between the Two Services PICs

**Purpose** Verify that traffic is load-balanced between the two Services PICs.

- Action** 1. Verify traffic flow between the IPv4 host on the home network and the IPv4 node on the Internet by using the **show services stateful-firewall flows** command.

```
user@AFTR> show services stateful-firewall flows
Interface: sp-0/1/0, Service set: dslite-svc-set1
Flow count
ICMP 10.0.10.1 -> 45.45.45.2 Watch I 3
 NAT source 10.0.10.1 -> 129.0.0.1
 Software 5002::12 -> 1001::1
DS-LITE 5002::12 -> 1001::1 Forward I 6
ICMP 45.45.45.2 -> 129.0.0.1 Watch 0 3
 NAT dest 129.0.0.1 -> 10.0.10.1
 Software 5002::12 -> 1001::1
```

The output shows ICMP statistics indicating the traffic flow between the IPv4 host on the home network to the IPv4 node on the Internet.

2. Issue the **show services software**, **show services stateful-firewall conversations**, **show services stateful-firewall flows count**, and **show services stateful-firewall statistics** commands to check the traffic flows.

```
user@AFTR> show services software
Interface: sp-0/1/0, Service set: dslite-svc-set2
Software 2001::3 -> 1001::1 Direction I Flow count 3

Interface: sp-1/3/0, Service set: dslite-svc-set1
Software 2001::2 -> 1001::1 Direction I Flow count 3
```

The output shows statistics for service set **dslite-svc-set2** associated with the services interface **sp-0/1/0** and service set **dslite-svc-set1** associated with the services interface **sp-1/3/0**.

```
user@AFTR> show services stateful-firewall conversations

Interface: sp-0/1/0, Service set: dslite-svc-set2

Conversation: ALG protocol: tcp
Number of initiators: 1, Number of responders: 1
Flow count
TCP 20.20.1.2:1025 -> 200.200.200.2:80 Forward I 189280
 NAT source 20.20.1.2:1025 -> 44.44.44.1:1024
 Software 2001::3 -> 1001::1
TCP 200.200.200.2:80 -> 44.44.44.1:1024 Forward 0 363675
 NAT dest 44.44.44.1:1024 -> 20.20.1.2:1025
 Software 2001::3 -> 1001::1

Interface: sp-1/3/0, Service set: dslite-svc-set1

Conversation: ALG protocol: tcp
Number of initiators: 1, Number of responders: 1
Flow count
TCP 20.20.1.2:1025 -> 200.200.200.2:80 Forward I 189280
 NAT source 20.20.1.2:1025 -> 44.44.44.1:1024
 Software 2001::3 -> 1001::1
TCP 200.200.200.2:80 -> 44.44.44.1:1024 Forward 0 363675
 NAT dest 44.44.44.1:1024 -> 20.20.1.2:1025
 Software 2001::3 -> 1001::1
```

```

count
TCP 20.20.1.2:1025 -> 200.200.200.2:80 Forward I 195847

 NAT source 20.20.1.2:1025 -> 33.33.33.1:1025
 Software 2001::2 -> 1001::1
TCP 200.200.200.2:80 -> 33.33.33.1:1025 Forward 0 391972

 NAT dest 33.33.33.1:1025 -> 20.20.1.2:1025
 Software 2001::2 -> 1001::1

Conversation: ALG protocol: tcp
Number of initiators: 1, Number of responders: 1
Flow
count
TCP 20.20.1.2:1025 -> 200.200.200.2:80 Forward I 219333

 NAT source 20.20.1.2:1025 -> 33.33.33.1:1024
 Software 2001::4 -> 1001::1
TCP 200.200.200.2:80 -> 33.33.33.1:1024 Forward 0 438848

 NAT dest 33.33.33.1:1024 -> 20.20.1.2:1025
 Software 2001::4 -> 1001::1

```

The output shows traffic flows for both services interfaces, **sp-0/1/0** and **sp-1/3/0**, indicating that both of the Services PICs are active.

```
user@AFTR> show services stateful-firewall flows count
```

```

Interface Service set Flow count
sp-0/1/0 dslite-svc-set2 3
sp-1/3/0 dslite-svc-set1 6

```

The output shows flow counts for both services interfaces, **sp-0/1/0** and **sp-1/3/0**, indicating that both of the Services PICs are active.

```
user@AFTR> show services stateful-firewall statistics
```

```

Interface Service set Accept Discard Reject Errors
sp-0/1/0 dslite-svc-set2 118991296 0 0 0
sp-1/3/0 dslite-svc-set1 237615050 0 0 0

```

**Meaning** The output shows traffic flows for both Services PICs, **sp-0/1/0** and **sp-1/3/0**. This indicates that the traffic is load-balanced between both of the Services PICs.

### Verifying Redundancy Between the Two Services PICs

**Purpose** Verify redundancy between the two Services PICs.

- Action** 1. Bring services PIC **sp-0/1/0** offline by issuing the **request chassis pic fpc-slot slot-number pic-slot pic-number offline** command.

```
user@host> request chassis pic fpc-slot 0 pic-slot 1 offline
fpc 0 pic 1 offline initiated, use "show chassis fpc pic-status" to verify
```

2. Issue the **show services stateful-firewall conversations** command again to check traffic flows through the redundant Services PIC **sp-1/3/0**.

Check the interface name and service-set name in the output.

```
user@host> show services stateful-firewall conversations
Interface: sp-1/3/0, Service set: dslite-svc-set1
```

Conversation: ALG protocol: tcp

Number of initiators: 1, Number of responders: 1

| Flow count                                   | State   | Dir | Frm    |
|----------------------------------------------|---------|-----|--------|
| TCP 20.20.1.2:1025 -> 200.200.200.2:80       | Forward | I   | 195847 |
| NAT source 20.20.1.2:1025 -> 33.33.33.1:1025 |         |     |        |
| Software 2001::2 -> 1001::1                  |         |     |        |
| TCP 200.200.200.2:80 -> 33.33.33.1:1025      | Forward | O   | 391972 |
| NAT dest 33.33.33.1:1025 -> 20.20.1.2:1025   |         |     |        |
| Software 2001::2 -> 1001::1                  |         |     |        |

Conversation: ALG protocol: tcp

Number of initiators: 1, Number of responders: 1

| Flow count                                   | State   | Dir | Frm    |
|----------------------------------------------|---------|-----|--------|
| TCP 20.20.1.2:1025 -> 200.200.200.2:80       | Forward | I   | 219333 |
| NAT source 20.20.1.2:1025 -> 33.33.33.1:1024 |         |     |        |
| Software 2001::4 -> 1001::1                  |         |     |        |
| TCP 200.200.200.2:80 -> 33.33.33.1:1024      | Forward | O   | 438848 |
| NAT dest 33.33.33.1:1024 -> 20.20.1.2:1025   |         |     |        |
| Software 2001::4 -> 1001::1                  |         |     |        |

**Meaning** The output indicates that all traffic is now routed through Services PIC **sp-1/3/0** when **sp-0/1/0** is deactivated. This indicates that redundancy is operational between the two Services PICs.

- Related Documentation**
- [Overview of Dual-Stack Lite on page 1](#)
  - [Redundancy and Load Balancing Using IPv6 Anycast Addresses on page 13](#)
  - [Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs on page 27](#)
  - [Example: Configuring Dual-Stack Lite for IPv6 Access on page 3](#)

---

## Example: Configuring AFTR Redundancy Using an IPv6 Anycast Address on Multiple AFTRs

---

This example shows how to configure redundancy with two or more DS-Lite Address Family Transition Routers (AFTRs) using a single IPv6 anycast address.

- [Requirements on page 27](#)
- [Overview on page 27](#)
- [Configuration on page 29](#)
- [Verification on page 38](#)

### Requirements

This example uses the following hardware and software components:


- Juniper Networks MX Series 3D Universal Edge Routers with Multiservices Dense Port Concentrators (DPCs)
- Juniper Networks® Junos® operating system (Junos OS) 10.4 or later running on the AFTRs

### Overview

You can provide redundancy using DS-Lite by configuring the same IPv6 anycast address on two or more AFTRs (software concentrators) as the software address. Basic Bridging Broadband Elements (B4s) only need to know this anycast address for the software endpoint, and the least-cost AFTR, per the routing updates, is used for the other software endpoint. If the least-cost AFTR goes down or the cost to get to this AFTR becomes higher than another AFTR, packets are redirected to the other AFTR. This is automatically handled by routing updates in the IPv6 cloud. You can also configure different Network Address Translation (NAT) pools at AFTRs and provide continuous service between IPv4 nodes in different domains.

[Figure 3 on page 28](#) provides a sample network topology for configuring IPv6 anycast address on two or more AFTRs.

**Figure 3: Sample Topology for DS-Lite Anycast Configuration Using Multiple AFTRs**



In [Figure 3 on page 28](#):

- The IPv4 client or host in the home network is configured with an IPv4 interface to the ISP and a static route to the IPv4 server on the Internet.
- The address of the NAT pool between AFTR1 and the Internet is 7.7.7.0/24. The address of the NAT pool between AFTR2 and the Internet is 8.8.8.0/24.
- The B4 or software initiator is configured with an IPv4 interface, an IPv6 interface, and an IPv4-in-v6 tunnel to an anycast address.
- The pure IPv6 node in the IPv6 cloud is configured with interfaces to the IPv6 interfaces and OSPFv3 for route updates.
- The AFTRs (AFTR1 and AFTR2) are configured with anycast address B001::1/128. If one of the links between the B4 and an AFTR fails, the other AFTR is used for traffic.
- The IPv4 node on the Internet is configured with an IPv4 interface and routes for reverse traffic.



---

## Configuration

**CLI Quick Configuration** To quickly configure this example, copy the following commands, paste them into a text file, remove any line breaks, change any details necessary to match your network configuration, and then copy and paste the commands into the CLI at the **[edit]** hierarchy level.

```
AFTR1 set chassis fpc 1 pic 1 adaptive-services service-package layer-3
set services service-set dsl-ss software-rules dsl-sw
set services service-set dsl-ss nat-rules dsl-nat1
set services service-set dsl-ss interface-service service-interface sp-1/1/0
set services software software-concentrator ds-lite dsl1 software-address b001::1
set services software software-concentrator ds-lite dsl1 mtu-v6 9192
set services software rule dsl-sw match-direction input
set services software rule dsl-sw term t1 then ds-lite dsl1
set services nat pool dsl-p1 address 7.7.0/24
set services nat pool dsl-p1 port automatic
set services nat rule dsl-nat1 match-direction input
set services nat rule dsl-nat1 term t1 from source-address 11.11.1.0/24
set services nat rule dsl-nat1 term t1 then translated source-pool dsl-p1
set services nat rule dsl-nat1 term t1 then translated translation-type napt-44
set services nat rule dsl-nat1 term t1 then syslog
set interfaces sp-1/1/0 unit 0 family inet
set interfaces sp-1/1/0 unit 0 family inet6
set interfaces ge-2/1/0 description B4-toward-AFTR
set interfaces ge-2/1/0 unit 0 family inet
set interfaces ge-2/1/0 unit 0 family inet6 service input service-set dsl-ss
set interfaces ge-2/1/0 unit 0 family inet6 service output service-set dsl-ss
set interfaces ge-2/1/0 unit 0 family inet6 address 8001::2/120
set interfaces ge-2/1/6 description AFTR-to-IPv4-node-on-the-Internet
set interfaces ge-2/1/6 unit 0 family inet address 88.88.1.1/24
set protocols ospf3 area 0.0.0.0 interface lo0.0
set protocols ospf3 area 0.0.0.0 interface ge-2/1/0.0

AFTR2 set chassis fpc 1 pic 1 adaptive-services service-package layer-3
set services nat pool dsl-p2 address 8.8.8.0/24
set services nat pool dsl-p2 port automatic
set services nat rule dsl-nat2 term t1 from source-address 11.11.1.0/24
set services nat rule dsl-nat2 match-direction input
set services nat rule dsl-nat2 term t1 then translated source-pool dsl-p2
set services nat rule dsl-nat2 term t1 then translated translation-type napt-44
set services software software-concentrator ds-lite dsl2 software-address b001::1
set services software software-concentrator ds-lite dsl2 mtu-v6 9192
set services software rule dsl-sw2 match-direction input
set services software rule dsl-sw2 term t1 then ds-lite dsl2
set interfaces sp-1/1/0 unit 0 family inet
set interfaces sp-1/1/0 unit 0 family inet6
set services service-set dsl-ss2 software-rules dsl-sw2
set services service-set dsl-ss2 nat-rules dsl-nat2
set services service-set dsl-ss2 interface-service service-interface sp-1/1/0
set interfaces ge-2/3/4 description V6-cloud-to-ipv6-node-in-v6-cloud
set interfaces ge-2/3/4 unit 0 family inet
set interfaces ge-2/3/4 unit 0 family inet6 service input service-set dsl-ss2
set interfaces ge-2/3/4 unit 0 family inet6 service output service-set dsl-ss2
set interfaces ge-2/3/4 unit 0 family inet6 address 9001::2/120
```

```
set interfaces ge-2/3/0 description to-ipv4-node-on-the-internet
set interfaces ge-2/3/0 unit 0 family inet address 89.89.1.1/24
set protocols ospf3 area 0.0.0.0 interface ge-2/3/4.0
set protocols ospf3 area 0.0.0.0 interface lo0.0
set routing-options static route 88.88.1.0/24 next-hop 89.89.1.2
```

---

### Configuring AFTR1

#### Step-by-Step Procedure

The following example requires you to navigate various levels in the configuration hierarchy. For information about navigating the CLI, see the CLI User Guide.

Router **AFTR1** is the primary router with an IPv6 interface to the ISP network (IPv6 cloud) and an IPv4 interface to the Internet. Configure the IPv4 interface, IPv6 interface, software endpoint, and NAT.

1. Configure the Layer 3 service package.

This example assumes that the PIC is in FPC 1, slot 1.

```
[edit chassis]
user@AFTR1# set fpc 1 pic 1 adaptive-services service-package layer-3
```

The service package with its associated **sp-** interface is for manipulating traffic before it is delivered to its destination. For details about configuring service packages, see the *Junos OS Services Interfaces Configuration Guide*.

2. Configure an IPv4 address and port for the NAT pool to specify the IPv4-to-IPv6 translation for packets traveling between the AFTR router and the Internet.

```
[edit services nat]
user@AFTR1# set pool ds1-p1 address 7.7.7.0/24
user@AFTR1# set pool dsl-p1 port automatic
```

3. Configure a NAT rule to translate the private IPv4 address from the home network to NAT pool **ds1-p1**.

NAT rules specify the traffic to be matched and the action to be taken when traffic matches the rule. In this example, only one rule is required to accomplish the address translation. The rule selects all traffic coming from the source address **11.11.1.0**.

```
[edit services nat]
user@AFTR1# set rule ds1-nat1 match-direction input
user@AFTR1# set rule ds1-nat1 term t1 from source-address 11.11.1.0/24
user@AFTR1# set rule ds1-nat1 term t1 then translated source-pool ds1-p1
user@AFTR1# set rule ds1-nat1 term t1 then translated translation-type napt-44
user@AFTR1# set rule dsl-nat1 term t1 then syslog
```

4. Configure the software concentrator, associate it with the IPv6 anycast address, and create a software rule.

The rule in this example specifies that any traffic destined for the software concentrator **dsl1** creates a new software. You can also configure more elaborate match conditions to perform as part of software initiator actions.

```
[edit services software]
user@AFTR1# set software-concentrator ds-lite dsl1 software-address b001::1
user@AFTR1# set rule dsl-sw match-direction input
user@AFTR1# set rule dsl-sw term t1 then ds-lite dsl1
```

- 
5. Configure the maximum transmission unit (ranging from 1280 to 9192 bytes) for the software for encapsulating IPv4 packets to IPv6.

This is the maximum packet size that can be sent on a tunnel from the AFTR to B4 without fragmentation. If the final length of the packet is greater than the MTU, the IPv6 packet would be fragmented.



**NOTE:** Including the `mtu-v6` statement is mandatory, and you cannot commit the example configuration unless this statement is configured.

[edit services software]

```
user@AFTR1# set software-concentrator ds-lite dsl1 mtu-v6 9192
```

6. Configure the services interface that contains the service set.

[edit interfaces]

```
user@AFTR1# set sp-1/1/0 unit 0 family inet
```

```
user@AFTR1# set sp-1/1/0 unit 0 family inet6
```

7. Configure a service set for the NAT and DS-Lite services using the **dsl-nat1** NAT rule and the **dsl-sw** software rule configured in Step 3 and Step 4.

In this example, the name of the service set is **dsl-ss**.

8. Associate the software and NAT rules and the service interface with the service set.

[edit services]

```
user@AFTR1# set service-set dsl-ss software-rules dsl-sw
```

```
user@AFTR1# set service-set dsl-ss nat-rules dsl-nat1
```

```
user@AFTR1# set service-set dsl-ss interface-service service-interface sp-1/1/0
```

9. Configure the interface between the home router running the B4 and the router in the ISP network running the AFTR, and include the IPv6 address of the AFTR router (software address).

In this example, the interface is **ge-2/1/0**.

[edit interfaces]

```
user@AFTR1# set ge-2/1/0 description B4-toward-AFTR
```

```
user@AFTR1# set ge-2/1/0 unit 0 family inet
```

```
user@AFTR1# set ge-2/1/0 unit 0 family inet6 address 8001::2/120
```

10. Associate the appropriate service set for the NAT and DS-Lite services.

[edit interfaces]

```
user@AFTR1# set ge-2/1/0 unit 0 family inet6 service input service-set dsl-ss
```

```
user@AFTR1# set ge-2/1/0 unit 0 family inet6 service output service-set dsl-ss
```

11. Configure the IPv4 interface between the AFTR and the Internet, and specify the IPv4 address connected to the Internet.

In this example, the interface is **ge-2/1/6**.

[edit interfaces]

```
user@AFTR1# set ge-2/1/6 description AFTR-to-V4-node-on-the-Internet
```

```
user@AFTR1# set ge-2/1/6 unit 0 family inet address 88.88.1.1/24
```

12. Configure OSPFv3 for route advertisements.

```
[edit protocols]
user@AFTR1# set ospf3 area 0.0.0.0 interface lo0.0
user@AFTR1# set ospf3 area 0.0.0.0 interface ge-2/1/0.0
```

**Results** In configuration mode, confirm your configuration by entering the **show chassis**, **show services**, **show interfaces**, and **show protocols ospf3** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@AFTR1# show chassis
fpc 1 {
 pic 1 {
 adaptive-services {
 service-package layer-3;
 }
 }
}

user@AFTR1# show services
service-set dsl-ss {
 software-rules dsl-sw;
 nat-rules dsl-nat1;
 interface-service {
 service-interface sp-1/1/0;
 }
}
software {
 software-concentrator {
 ds-lite dsl1 {
 software-address b001::1;
 mtu-v6 9192;
 }
 }
 rule dsl-sw {
 match-direction input;
 term t1 {
 then {
 ds-lite dsl1;
 }
 }
 }
}
nat {
 pool dsl-p1 {
 address 7.7.0/24;
 port {
 automatic;
 }
 }
 rule dsl-nat1 {
 match-direction input;
 term t1 {
 from {
 source-address {
```

```

 11.11.1.0/24;
 }
}
then {
 translated {
 source-pool dsl-p1;
 translation-type {
 napt-44;
 }
 }
}
syslog;
}
}
}
}

user@AFTR1# show interfaces
sp-1/1/0 {
 unit 0 {
 family inet;
 family inet6;
 }
}
ge-2/1/0 {
 description B4-toward-AFTR;
 unit 0 {
 family inet;
 family inet6 {
 service {
 input {
 service-set dsl-ss;
 }
 output {
 service-set dsl-ss;
 }
 }
 }
 address 8001::2/120;
 }
}
ge-2/1/6 {
 description AFTR-to-V4-node-on-the-Internet;
 unit 0 {
 family inet {
 address 88.88.1.1/24;
 }
 }
}

user@AFTR1# show protocols ospf3
area 0.0.0.0 {
 interface lo0.0;
 interface ge-2/1/0.0;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Configuring AFTR2

### Step-by-Step Procedure

Router **AFTR2** is the secondary router with an IPv6 interface to the ISP network (IPv6 cloud) and an IPv4 interface to the Internet. Configure the IPv4 interface, IPv6 interface, software endpoint, and NAT.

1. Configure the Layer 3 service package.

This example assumes that the PIC is in FPC 1, slot 1.

**[edit chassis]**

```
user@AFTR2# set fpc 1 pic 1 adaptive-services service-package layer-3
```

The service package with its associated **sp-** interface is for manipulating traffic before it is delivered to its destination. For details about configuring service packages, see the *Junos OS Services Interfaces Configuration Guide*.

2. Configure an IPv4 address and port for the NAT pool to specify the IPv4-to-IPv6 translation for packets traveling between the AFTR router and the Internet.

**[edit services nat]**

```
user@AFTR2# set pool dsl-p2 address 8.8.8.0/24
```

```
user@AFTR2# set pool dsl-p2 port automatic
```

3. Configure a NAT rule to translate the private IPv4 address from the home network to NAT pool **dsl-p2**.

NAT rules specify the traffic to be matched and the action to be taken when traffic matches the rule. In this example, only one rule is required to accomplish the address translation. The rule selects all traffic coming from the source address **11.11.1.0**.

**[edit services nat]**

```
user@AFTR2# set rule dsl-nat2 match-direction input
```

```
user@AFTR2# set rule dsl-nat2 term t1 from source-address 11.11.1.0/24
```

```
user@AFTR2# set rule dsl-nat2 term t1 then translated source-pool dsl-p2
```

```
user@AFTR2# set rule dsl-nat2 term t1 then translated translation-type napt-44
```

4. Configure the software concentrator, associate it with the IPv6 anycast address, and create a software rule.

The rule in this example specifies that any traffic destined for the **dsl2** software concentrator creates a new software.

**[edit services software]**

```
user@AFTR2# set software-concentrator ds-lite dsl2 software-address b001::1
```

```
user@AFTR2# set rule dsl-sw2 match-direction input
```

```
user@AFTR2# set rule dsl-sw2 term t1 then ds-lite dsl2
```

5. Configure the maximum transmission unit (ranging from 1280 to 9192 bytes) for the software for encapsulating IPv4 packets to IPv6.

This is the maximum packet size that can be sent on a tunnel from the AFTR to B4 without fragmentation. If the final length of the packet is greater than the MTU, the IPv6 packet would be fragmented.



**NOTE:** Including the **mtu-v6** statement is mandatory, and you cannot commit the example configuration unless this statement is configured.

```
[edit services software]
```

```
user@AFTR2# set software-concentrator ds-lite dsl2 mtu-v6 9192
```

6. Configure the services interface that contains the service set.

```
[edit interfaces]
```

```
user@AFTR2# set sp-1/1/0 unit 0 family inet
```

```
user@AFTR2# set sp-1/1/0 unit 0 family inet6
```

7. Configure a service set for the NAT and DS-Lite services using the **dsl-nat2** NAT rule and the **dsl-sw2** software rule configured in Step 3 and Step 4.

In this example, the name of the service set is **dsl-ss2**.

8. Associate the software and NAT rules and the service interface with the service set.

```
[edit services]
```

```
user@AFTR2# set service-set dsl-ss2 software-rules dsl-sw2
```

```
user@AFTR2# set service-set dsl-ss2 nat-rules dsl-nat2
```

```
user@AFTR2# set service-set dsl-ss2 interface-service service-interface sp-1/1/0
```

9. Configure the interface between the pure IPv6 node in the IPv6 cloud and the AFTR. In this example, the interface is **ge-2/3/4**.

```
[edit interfaces]
```

```
user@AFTR2# set ge-2/3/4 description V6-cloud-to-ipv6-node-in-v6-cloud
```

```
user@AFTR2# set ge-2/3/4 unit 0 family inet
```

10. Include the IPv6 address of the AFTR router (software address).

```
[edit interfaces]
```

```
user@AFTR2# set ge-2/3/4 unit 0 family inet6 address 9001::2/120
```

11. Associate the appropriate service set for the NAT and DS-Lite services.

```
[edit interfaces]
```

```
user@AFTR2# set ge-2/3/4 unit 0 family inet6 service input service-set dsl-ss2
```

```
user@AFTR2# set ge-2/3/4 unit 0 family inet6 service output service-set dsl-ss2
```

12. Configure the IPv4 interface between the AFTR and the Internet and specify the IPv4 address connected to the Internet.

In this example, the interface is **ge-2/3/0**.

```
[edit interfaces]
```

```
user@AFTR2# set ge-2/3/0 description to-ipv4-node-on-the-internet
```

```
user@AFTR2# set ge-2/3/0 unit 0 family inet address 89.89.1.1/24
```

13. Configure OSPFv3 for route advertisements.

```
[edit protocols ospf3]
```

```
user@AFTR2# set area 0.0.0.0 interface ge-2/3/4.0
```

```
user@AFTR2# set area 0.0.0.0 interface lo0.0
```

14. Configure a static route to the IPv4 node on the Internet.

```
[edit routing-options]
```

```
user@AFTR2# set static route 88.88.1.0/24 next-hop 89.89.1.2
```

## Results

---

In configuration mode, confirm your configuration by entering the **show chassis**, **show interfaces**, **show services**, **show protocols ospf3**, and **show routing-options** commands. If the output does not display the intended configuration, repeat the instructions in this example to correct the configuration.

```
user@AFTR2# show chassis
fpc 1 {
 pic 1 {
 adaptive-services {
 service-package layer-3;
 }
 }
}

user@AFTR2# show interfaces
sp-1/1/0 {
 unit 0 {
 family inet;
 family inet6;
 }
}
ge-2/3/0 {
 description to-ipv4-node-on-the-internet;
 unit 0 {
 family inet {
 address 89.89.1.1/24;
 }
 }
}
ge-2/3/4 {
 description V6-cloud-to-ipv6-node-in-v6-cloud;
 unit 0 {
 family inet;
 family inet6 {
 service {
 input {
 service-set dsl-ss2;
 }
 output {
 service-set dsl-ss2;
 }
 }
 address 9001::2/120;
 }
 }
}

user@AFTR2# show services
service-set dsl-ss2 {
 software-rules dsl-sw2;
 nat-rules dsl-nat2;
```



```

interface-service {
 service-interface sp-1/1/0;
}
}
software {
 software-concentrator {
 ds-lite dsl2 {
 software-address b001::1;
 mtu-v6 9192;
 }
 }
 rule dsl-sw2 {
 match-direction input;
 term t1 {
 then {
 ds-lite dsl2;
 }
 }
 }
}
nat {
 pool dsl-p2 {
 address 8.8.8.0/24;
 port {
 automatic;
 }
 }
 rule dsl-nat2 {
 match-direction input;
 term t1 {
 from {
 source-address {
 11.11.1.0/24;
 }
 }
 then {
 translated {
 source-pool dsl-p2;
 translation-type {
 napt-44;
 }
 }
 }
 }
 }
}
}

user@AFTR2# show protocols ospf3
area 0.0.0.0 {
 interface ge-2/3/4.0;
 interface lo0.0;
}

user@AFTR2# show routing-options
static {
 route 88.88.1.0/24 next-hop 89.89.1.2;
}

```

If you are done configuring the device, enter **commit** from configuration mode.

## Verification

Confirm that the configuration is working properly.

### Verifying Redundancy of the AFTRs

**Purpose** Verify that traffic flow is maintained using the secondary AFTR if an interface on one AFTR is brought offline.

**Action** 1. Verify traffic flow between the IPv4 host on the home network and the IPv4 node on the Internet.

Additionally, check the software flows for AFTR1.

**user@AFTR1> show services stateful-firewall flows**

Interface: sp-1/0/0, Service set: dsl-ss

| Flow                                   | State   | Dir | Frm |
|----------------------------------------|---------|-----|-----|
| count                                  |         |     |     |
| TCP 20.20.1.2:1025 -> 200.200.200.2:80 | Forward | I   |     |
| 107621                                 |         |     |     |
| NAT source 20.20.1.2:1025 -> 7.7.7.0   |         |     |     |
| Software 2001::3 -> 1001::1            |         |     |     |
| TCP 200.200.200.2:80 -> 7.7.7.0        |         |     |     |
| 208420                                 |         |     |     |
| NAT source 7.7.7.0 -> 20.20.1.2:1025   |         |     |     |
| Software 2001::3 -> 1001::1            |         |     |     |
| ICMP 10.0.10.1 -> 88.88.88.1.1         | Watch   | I   |     |
| 3                                      |         |     |     |
| NAT source 10.0.10.1 -> 129.0.0.1      |         |     |     |
| Software 8001::2 -> 1001::1            |         |     |     |
| DS-LITE 2001::3 -> 1001::1             | Forward | I   |     |
| 6                                      |         |     |     |
| ICMP 88.88.88.1 -> 129.0.0.1           | Watch   | O   |     |
| 3                                      |         |     |     |
| NAT dest 129.0.0.1 -> 10.0.10.1        |         |     |     |
| Software 8001::2 -> 1001::1            |         |     |     |

The output shows ICMP source and destination addresses indicating traffic flow between the IPv4 host on the home network and the IPv4 node on the Internet. The DS-Lite protocol statistics indicate the software flows.

2. Deactivate the interface **ge-2/1/0** on AFTR1.

**user@AFTR1# deactivate interfaces ge-2/1/0**

3. Commit the configuration.

4. Issue the **show services stateful-firewall flows** command on AFTR2 to verify the creation of software flows.

Additionally, verify traffic flows between the IPv4 host on the home network and the IPv4 node on the Internet.

**user@AFTR2> show services stateful-firewall flows**

Interface: sp-1/0/0, Service set: dsl-ss2

| Flow | State | Dir | Frm |
|------|-------|-----|-----|
|------|-------|-----|-----|

```

count
TCP 20.20.1.2:1025 -> 200.200.200.2:80 Forward I
107621
 NAT source 20.20.1.2:1025 -> 8.8.8.0
 Software 2001::3 -> 1001::1

TCP 200.200.200.2:80 -> 7.7.7.0
208420
 NAT source 8.8.8.0 -> 20.20.1.2:1025
 Software 2001::3 -> 1001::1
ICMP 10.0.10.1 -> 88.88.88.1.1 Watch I
3
 NAT source 10.0.10.1 -> 129.0.0.1
 Software 2001::3 -> 1001::1
DS-LITE 2001::3 -> 1001::1 Forward I
6
ICMP 88.88.88.1 -> 129.0.0.1 Watch 0
3
 NAT dest 129.0.0.1 -> 10.0.10.1
 Software 2001::3 -> 1001::1

```

**Meaning** The output shows NAT and software source and destination addresses for traffic flow between AFTR2 and the IPV4 node on the Internet. This indicates that AFTR2 is now operating as the secondary AFTR when AFTR1 is offline.

- Related Documentation**
- [Overview of Dual-Stack Lite on page 1](#)
  - [Redundancy and Load Balancing Using IPv6 Anycast Addresses on page 13](#)
  - [Example: Configuring Dual-Stack Lite for IPv6 Access on page 3](#)
  - [Example: Configuring Redundancy and Load Balancing Using a Single AFTR and Multiple Services PICs on page 14](#)

