

Configuring IPV6 Subscriber Services

SYSTEM ADMINISTRATOR GUIDE

Copyright

© Ericsson AB 2009–2011. All rights reserved. No part of this document may be reproduced in any form without the written permission of the copyright owner.

Disclaimer

The contents of this document are subject to revision without notice due to continued progress in methodology, design and manufacturing. Ericsson shall have no liability for any error or damage of any kind resulting from the use of this document.

Trademark List

SmartEdge is a registered trademark of Telefonaktiebolaget LM Ericsson.

NetOp is a trademark of Telefonaktiebolaget LM Ericsson.



Contents

1	Overview	1
1.1	General IPv6 Protocol Concepts	1
1.2	SmartEdge Implementation of IPv6	4
1.3	Overview of PPP Session Establishment	23
2	Configuration and Operations Tasks	25
2.1	Recommendations	25
2.2	Requirements	25
2.3	Restrictions	25
2.4	Configuring a SmartEdge Router to Provide IPv6 and Dual-Stack Subscriber Services	26
2.5	Conserving IPv4 Addresses for Dual-Stack Subscribers	39
2.6	IPv6 Subscriber Services Operations	40
3	Examples	43
3.1	End-to-End Solution Configurations	43
3.2	Configuring Dual-Stack Subscriber Support on a SmartEdge LNS	47
3.3	Detailed Configuration Examples for Individual Elements of an IPv6 Solution	52





1 Overview

When configured as a broadband remote access server (BRAS) or as an LAC (L2TP access concentrator) or LNS (L2TP network server), the SmartEdge® router supports the address assignment and management of Internet Protocol version 6 (IPv6) Point-to-Point Protocol (PPP) and PPP over Ethernet (PPPoE) subscribers. This document describes the configuration of IPv6 subscriber services for single (IPv6 only) and dual-stack (IPv6 and IPv4) PPP subscribers.

Note: This document describes the configuration and management of IPv6 subscriber services only. To configure IPv4 subscriber services on the SmartEdge router, see *Configuring Subscribers*.

To configure IPv6 subscriber services on the SmartEdge router, you must have enabled the IPv6 subscriber license with the *subscriber* command; dual-stack subscriber services also require a license for IPv4 subscribers. Depending on the subscriber licenses enabled on the SmartEdge router, a subscriber on a PPP- or PPPoE-encapsulated circuit connecting to a SmartEdge router can receive IPv4, IPv6 or dual-stack subscriber services.

Both the network (trunk) facing and access (subscriber) facing circuits carrying subscriber traffic can be bundled into link aggregation groups (LAGs) for resiliency to individual link failures and route redistribution. A network-facing LAG between two provider edge routers at each end of the core network provides improved traffic distribution. A subscriber-facing LAG between two customer edge routers secures the gateways for the L2VPN between them.

For information on:

- Enabling licenses in the SmartEdge router, see *Enabling Licensed Features*.
- Configuring the SmartEdge router as an LAC, see *Configuring L2TP*.
- Configuring LAGs, see *Configuring Link Aggregation*.
- Troubleshooting IPv6 subscriber services, see *Troubleshooting IPv6 and Dual-Stack Subscriber Services*.

1.1 General IPv6 Protocol Concepts

Before configuring IPv6 subscriber services on the SmartEdge router, you must be familiar with the differences between IPv4 and IPv6, address types supported by IPv6, and the IPv6 address format.



1.1.1 Differences Between IPv4 and IPv6

Table 1 describes the differences between IPv4 and IPv6.

Table 1 Differences Between IPv4 and IPv6

Element	IPv4	IPv6
Address size	32 bits	128 bits You do not need to type the full 128-bit address to pass a prefix to an end device.
Number of addresses supported	2 ³²	2 ¹²⁸
Types of addresses supported	Global unicast	Global unicast, link local, multicast, anycast.
PPP address assignment	/32 allocated through Internet Protocol Control Protocol version 4 (IPCPv4)	No. IPv6 supports Dynamic Host Configuration Protocol version 6 (DHCPv6) Prefix Delegation (PD) or Neighbor Discovery (ND). Address assignment is encapsulation independent.
Broadcast address	Yes	No; multicast is supported instead.
Consolidated OAM	No	Address Resolution Protocol (ARP) and Duplicate Address Detection (DAD).
Address auto-configuration through ND	No	Yes
Prefixes	No	The SmartEdge assigns a prefix to its PPP subscribers. Customer-premises equipment (CPE) can have one or more prefixes assigned to a wide-area network (WAN) link, and one or more delegated prefixes for its downstream nodes.
Fixed 40 bytes	No	Yes

1.1.2 IPv6 Address Types

IPv6 addresses are 128 bits long, and the first 64 bits are reserved for routing and network addressing. IPv6 supports the following types of addresses:



- **Global unicast**—Associated with a single interface only. The host uses router advertisements to determine globally unique addresses (GUAs). Global unicast addresses consist of a global routing prefix, a subnet ID, and an interface ID. The global routing prefix and subnet ID are the routing and networking part of the Global unicast address, and the interface ID is derived from the MAC address. See Table 2 for details about the components of a Global unicast IPv6 address.
- **Link-local unicast**—Used for communication among nodes on the same link (the link on which the link local address is assigned). A link local address is an IPv6 address that is unique on the link. The link is a network segment. For example, a link can be one Ethernet interface that is connected to ten different nodes. Link-local unicast addresses are automatically assigned by end-user equipment and require no external configuration. See Table 3 for details about the components of a link-local unicast IPv6 address.
- **Multicast**—Transmits a packet to all of the interfaces in a specified multicast group; each packet sent is replicated on every endpoint.
- **Anycast**—Transmits a packet to an address that is associated with multiple interfaces, but only one interface receives the packet.

Table 2 Components of Global IPv6 Address

Routing and Networking Part of the Address		Unique ID Derived from the Line Card MAC Address
Global routing prefix of size n bits, where n can be from 1 to 56 bits. Typically, the global routing prefix is 48 bits long.	Subnet ID of size $64 - n$ bits. The subnet ID can be from 8 to 16 bits, but is typically 16 bits.	64-bit interface ID

Table 3 Components of Link-local IPv6 Address

Routing and Networking Part of the Address	Unique Interface ID Derived from the Line Card MAC Address
Subnet prefix of size n bits, where n can be from 1 to 64 bits. Typically, the subnet prefix is 10 bits.	Interface ID of size $128 - n$ bits. Typically, the Interface ID is 118 bits.

With IPv6, an interface can have multiple IPv6 addresses of any type. For example, an interface can have three IPv6 multicast addresses, one IPv6 unicast address, and two anycast IPv6 addresses.

Some IPv6 addresses are reserved. Table 4 describes the reserved IPv6 addresses and their notation:

Table 4 Reserved IPv6 Address Notation

Address type	Binary prefix	IPv6 Notation
Unspecified	00...0 (128 bits)	::/128
Loopback	00...1 (128 bits)	::1/128
Multicast	11111111	FF00::/8



Address type	Binary prefix	IPv6 Notation
Link-local	1111111010	FE80::/10
Global Unicast	<p>All addresses are GUAs except for the following:</p> <ul style="list-style-type: none"> • Unspecified • Loopback • Multicast • Link-local 	<p><i>nnn:nnn:nnn:nnn</i> = routing prefix</p> <p><i>mmmmmmmmmm</i> = subnet ID</p> <p><i>128-n-m</i> = interface ID</p>

1.1.3 Address Format

IPv6 addresses are typically composed of two parts: a 64-bit network or subnetwork prefix, and a 64-bit interface ID (128 bits total). Typically, IPv6 addresses are written with hexadecimal digits and colon separators in the following format:

AAAA:BBBB:CCCC:DDDD:EEEE:FFFF:GGGG:HHHH

The IPv6 hexadecimal numbering system uses decimal digits 0 to 9 and letters A, B, C, D, E, and F (which represent the numbers 10, 11, 12, 13, 14, and 15). The decimal digit 16 is represented in hexadecimal by the number 10. Each section of hexadecimal characters represents 16 bits of the address and is separated by a colon. In the previous example, AAAA represents the first section of an IPv6 address, BBBB represents the second section, and so forth.

Following is an example of an IPv6 address. In this example, all 32 hexadecimal digits are represented:

ABCD:A162:1234:1234:ABCD:1234:5432:1010

By dropping nonsignificant and leading 0s, you can shorten an IPv6 address to eight hexadecimal digits. For example, the IPv6 address 1060:0000:0000:0000:0006:0600:800C:228A can be shortened to 1060:0:0:0:6:600:800C:228A. You can shorten an IPv6 address even further by replacing consecutive 0s with double colons. For example, the IPv6 address 1060:0:0:0:6:600:800C:228A can be shortened to 1060::6:600:800C:228A.

Note: Double colons are allowed only once in each IPv6 address.

For more information about IPv6 address formatting, see RFC 4291, *IP Version 6 Addressing Architecture*.

1.2 SmartEdge Implementation of IPv6

The sections that follow describe the specifications and configurations supported by the SmartEdge router.



1.2.1 Hardware Support Specifications

IPv6 subscriber services are supported on the following SmartEdge routers:

- SmartEdge 100
- SmartEdge 400
- SmartEdge 600
- SmartEdge 800
- SmartEdge 1200
- SmartEdge 1200H

IPv6 subscriber services are supported on the following traffic cards only:

- PPA2-based 5-port Gigabit Ethernet
- PPA2-based 10-port Gigabit Ethernet
- PPA2-based 2-port 60 Fast Ethernet–Gigabit Ethernet
- PPA2-based 1-port 10 Gigabit Ethernet
- PPA3-based 4-port 10 Gigabit Ethernet
- PPA3-based 20-port Gigabit Ethernet

Note: IPv6 subscriber services are not supported on PPA1-based traffic cards.

1.2.2 Subscriber Session Specifications

Subscriber sessions can be single-stack or dual-stack. Single-stack subscriber sessions have only one type of IP service configured (IPv4 or IPv6) and exclusively support one type of traffic (IPv4 or IPv6). Dual-stack subscriber sessions are authorized for both IPv4 and IPv6, and can simultaneously support both IPv4 and IPv6 traffic. Although dual-stack subscriber sessions are authorized to simultaneously support both IPv4 and IPv6 traffic, it is not necessary for both stacks to be active at the same time.

A dual-stack subscriber session consists of a single circuit bound to a single interface. Table 5 shows the number of dual-stack subscriber sessions the SmartEdge router supports for each card type:

Table 5 Number of Dual-Stack Subscriber Sessions per Card

Card Type	Sessions per System
XCRP3 Controller card	32,000
XCRP4 Controller card	64,000



Card Type	Sessions per System
PPA2-based 5-port Gigabit Ethernet	16,000
PPA2-based 10-port Gigabit Ethernet traffic card	16,000
PPA2-based 2-port 60 Fast Ethernet–Gigabit Ethernet traffic card	16,000
PPA2-based 1-port 10 Gigabit Ethernet traffic card	16,000
PPA3-based 4-port 10 Gigabit Ethernet	24,000
PPA3-based 20-port Gigabit Ethernet	24,000

1.2.3 Supported IPv6 Subscriber Configurations

The SmartEdge router supports IPv6 subscriber services for PPP and PPPoE subscribers. You can configure IPv6 prefixes for subscribers:

- Statically
- Using the Delegated-IPv6-Prefix attribute, which can be configured statically or by using the Delegated-IPv6-Prefix RADIUS attribute.
- Through a DHCPv6 PD prefix pool

The SmartEdge router uses Neighbor Discovery (ND) to assign an IPv6 prefix to the WAN link between the SmartEdge router and CPE router.

1.2.4 PPP Session Specifications

IPv6 Control Protocol (IPv6CP) negotiation is supported for authenticated IPv6 PPP subscribers authorized for IPv6. During IPv6CP negotiation, both ends of the PPP circuit exchange their interface IDs, in which the MAC address of the subscriber is embedded. If a subscriber cannot generate its own interface ID, the subscriber takes its interface ID from the subscriber record (if the record contains a client interface ID). In cases where the subscriber cannot generate an interface ID and no interface ID is available in the RADIUS database, PPP randomly generates an interface ID.

The SmartEdge OS then learns neighbor MAC addresses from PPP and installs those addresses in the RIB.

Dual-stack subscribers use IPv6CP for IPv6 subscribers and IPCP for IPv4 subscribers. IPCP and IPv6CP are independent of one another; if IPv6CP fails, IPCP still operates and vice-versa. Dual-stack PPP sessions are negotiated as follows:

- In configurations where the CPE is a router, the CPE initiates PPP sessions. After successfully negotiating the Link Control Protocol (LCP) between a



SmartEdge router and a CPE router, the CPE router runs IPCP negotiation for IPv4 subscribers and IPv6CP negotiation for IPv6 subscribers.

- In configurations where the CPE is a bridge, the host initiates the PPP sessions. After successfully negotiating the Link Control Protocol (LCP) between a SmartEdge router and a CPE bridge, the host runs IPCP negotiation for IPv4 subscribers and IPv6CP negotiation for IPv6 subscribers. The SmartEdge OS declines IPv6CP negotiation for subscribers not authorized for IPv6.

Dual-stack subscriber sessions remain active until either of the following events occur:

- LCP terminates
- All network control protocols (NCPs) terminate
- A DAD failure, (ND brings down the IPv6 session for a subscriber.)

When IPCP and IPv6CP report that a PPP session has terminated, the SmartEdge router terminates the subscriber session.

1.2.5 Multibind Interfaces

Multibind interfaces are the only interfaces that support IPv6 subscriber services; DHCPv6 server interfaces must be configured under a multibind interface. A multibind interface allows multiple circuits to be bound to a single interface and typically is used for subscriber circuits. You can also specify a multibind interface as a last-resort interface for any incoming subscriber circuit with a subscriber record that does not include an IP address that is assigned to any other interface. If a subscriber session is established, and no valid interface exists to which it can bind, the session binds to the last-resort interface.

The following restrictions apply when you configure a multibind interface for IPv6 subscriber services:

- A last-resort multibind interface must be configured as unnumbered, using the *ipv6 unnumbered* command (in interface configuration mode).
- The interface from which the IP address is borrowed for an unnumbered interface must be in the same context as the unnumbered interface.
- A standard multibind interface (not a last-resort interface) must have an IP address assigned explicitly, using the *ip address* command (in interface configuration mode).

For more information about multibind interfaces, see *Configuring Contexts and Interfaces*.



1.2.6 Subscriber Attributes

You can configure subscriber attributes:

- In a subscriber record, as described in *Configuring Subscriber Attributes in a Subscriber Record*. The subscriber record is stored locally (on the SmartEdge OS) or on a RADIUS server.
- In a subscriber profile, as described in *Configuring Subscriber Attributes in a Subscriber Profile*.

Note: The subscriber record takes precedence over subscriber profiles. As result, attributes in a subscriber profile are overridden when you configure identical attributes with different values in a specific subscriber record.

1.2.6.1 Configuring Subscriber Attributes in a Subscriber Record

The SmartEdge router uses subscriber records to configure a set of subscriber attributes that are applied to subscribers. Some examples of attributes that can be configured are the subscriber name, password, authentication, access control, rate limiting, and policing information. A record is specific to the context in which the subscriber is configured.

You can configure the following IPv6-specific subscriber attributes in a subscriber record:

- The delegated IPv6 prefix
- The neighbor discovery prefix (also called the framed IPv6 prefix)
- The subscriber IPv6 route (also called the framed IPv6 route)
- An ND profile used to apply ND attributes to the IPv6 subscriber
- Enabling or disabling source validation for IPv6

Note: To bring up an IPv6 stack, you must configure either the delegated IPv6 prefix or the neighbor discovery prefix (the framed IPv6 prefix).

You configure subscriber records in one of two ways:

- Locally, by using commands in the SmartEdge command line interface (CLI).

Subscriber records provide local authentication and authorization information whenever a remote authentication and authorization server, such as a RADIUS server, is not available or not required.

- Using attributes (authentication, accounting, or both) stored on a RADIUS server.



Note: If the RADIUS server is configured in the local context of the SmartEdge router, the RADIUS server can be used in all contexts. If the RADIUS server is configured in a nonlocal context, it can be used in that nonlocal context only. For more information about how the RADIUS server and the SmartEdge router interact, see *Configuring RADIUS*.

The following RADIUS attributes are supported for IPv6 subscribers:

- Framed-IPv6-Route—Provides an IPv6 route for the subscriber.
- Framed-Interface-Id—Provides an interface ID for PPP clients that do not generate their own interface ID.
- NAS-IPv6-Address—Identifies the IPv6 address of the Network Access Server (NAS) in RADIUS access-request and access-accounting messages. For more information about NAS and RADIUS, see *Configuring RADIUS*.
- Delegated-IPv6-Prefix—Identifies an IPv6 prefix that can be assigned to the subscriber using DHCPv6, so the subscriber can delegate the prefix to its downstream nodes.
- Framed-IPv6-Prefix—Used for stateless address autoconfiguration (SLAAC) and ND. If the Framed-IPv6-Prefix attribute is configured to be all 0s or all 1s, IPv6 prefixes are assigned from a shared IPv6 prefix pool.
- Framed-IPv6-Pool—Dictates that the specified subscriber obtain its IPv6 prefixes from an IPv6 pool configured under the same context as the subscriber.
- Delegated-Max-Prefix—Ericsson VSA used to configure the maximum number of DHCPv6 PD prefixes this subscriber can allocate to hosts. The range of values is 1 to 5.
- RB-IPv6-DNS—Ericsson VSA used to configure the IPv6 Primary and Secondary DNS of a subscriber. See *RADIUS Attributes* for more information about this VSA.
- RB-IPv6-Option—Ericsson VSA used to configure multiple ipv6 attributes for a single subscriber. See *RADIUS Attributes* for more information about this VSA.

Note: Use RADIUS filtering to configure individual attributes to be dropped from access and access accounting request messages.

1.2.6.2

Configuring Subscriber Attributes in a Subscriber Profile

In addition to the subscriber record, you can create and assign two types of subscriber profiles:

- Default subscriber profiles contain attributes shared by multiple subscribers in a single configuration. By configuring a default subscriber profile,



you do not need to apply the same attributes separately to each subscriber record. Attributes in the default subscriber profile apply to all IPv6 subscribers in the same context. To use the default profile, you must explicitly configure it by using the *subscriber default* command in context configuration mode; this accesses subscriber configuration mode, where you can configure attributes for the default subscriber profile.

Note: ECMP routing for subscribers, which is configured in the default subscriber profile, is not supported for IPv6 subscribers.

- Named subscriber profiles can be assigned to one or more subscribers. Unlike the default subscriber profile, which is assigned automatically to every subscriber record, you must explicitly assign a named subscriber profile to a subscriber record.

Attributes in the subscriber record take precedence over identical attributes configured in the named subscriber profile, and attributes in the named subscriber profile take precedence over identical attributes configured in the default subscriber profile.

1.2.7 AAA Support for IPv6 Subscribers

An IPv6 subscriber must be authorized through AAA before PPP negotiates connectivity and ND processes packets. If a protocol (for example, the IPv6 protocol) is not authorized, PPP does not negotiate that protocol with a client, even when the PPP negotiation process is initiated by a client.

The following AAA attributes are supported for IPv6 subscribers:

- IPv6 route tag
- IPv6-ND-profile
- IPv6 Option source-validation

For general information about how AAA works on the SmartEdge router, see *Configuring Authentication, Authorization, and Accounting*.

1.2.8 DHCPv6 Prefix Delegation

With IPv6, DHCPv6 obtains IPv6 prefixes from:

- The Delegated-IPv6-Prefix attribute in a subscriber record
- A DHCPv6 PD prefix that is statically mapped to a DHCPv6 Unique Identifier (DUID). This document refers to this type of prefix as a *statically mapped delegated prefix*.
- A DHCPv6 PD pool



When DHCPv6 has the IPv6 prefix, the DHCPv6 server then assigns that prefix to a subscriber. If the subscriber is a CPE router, it uses the prefix to derive a set of longer prefixes that are sent to its clients. Subscribers that are not CPE routers do not use delegated prefixes.

The Delegated-Max-Prefix attribute dictates the maximum number of IPv6 prefixes that can be delegated to a subscriber. Prefixes are assigned hierarchically; the Delegated-IPv6-Prefix attribute in a subscriber record takes precedence over statically mapped delegated prefixes, which take precedence over prefixes in the DHCPv6 PD pools. Use the `aaa accounting event dhcpv6` command to enable inclusion of DHCPv6 PD transition events such as PD host address assignment or release in accounting messages.

For example, consider a situation where a subscriber requests five IP addresses from a router that has the following configuration:

- The Delegated-Max-Prefix attribute is set to six (so that up to six IPv6 prefixes can be leased at any given time).
- One delegated IPv6 prefix is configured in the subscriber record.
- Two statically mapped delegated prefixes are configured.
- Three DHCPv6 PD pools are configured, and all three have IPv6 prefixes available.

In this instance, DHCPv6 assigns the IPv6 prefix from the subscriber record, both of the statically mapped delegated prefixes, and two prefixes from any of the three DHCPv6 PD pools. Those prefixes and the pools from which those prefixes are chosen are random and cannot be configured.

In addition to IPv6 prefix delegation, the DHCPv6 server provides additional information to a subscriber, such as the default domain and DNS name-server address.

When configuring DHCPv6, keep in mind that:

- A static delegated IPv6 prefix can be from 1 to 128 bits in length. In common configurations, DHCPv6 provides IPv6 prefixes that are between 48 to 64 bits.
- Static delegated IPv6 prefixes are used by a CPE for a specified lifetime.
- Static delegated IPv6 prefixes remain assigned to the CPE until their valid lifetimes expire, or until the CPE sends a DHCPv6 RELEASE message to the DHCPv6 server.
- Only one DHCPv6 server is allowed per context.
- A DHCPv6 server interface can be a last-resort or non-last-resort multibind interface.
- You can configure multiple static IPv6 prefixes in a subscriber record.



Note: Unlike framed IPv6 prefixes, DHCPv6 PD prefixes do not use route tags.

For faster IPv6 prefix delegation, you can configure DHCPv6 to use the RAPID COMMIT option. With the RAPID COMMIT option, only two messages (SOLICIT and REPLY messages) are exchanged between the DHCPv6 server and the CPE. You typically use the RAPID COMMIT option when the CPE can connect to only one server.

Note: For general information about how DHCP works on the SmartEdge router, see *Configuring DHCP*.

The SmartEdge router supports both stateful and stateless DHCPv6, which are described in the sections that follow.

1.2.8.1 Stateful DHCPv6

With stateful DHCPv6, the DHCPv6 server is used for DHCPv6 prefix delegation and maintains the dynamic state of each client. The IPv6 prefixes remain assigned to the CPE until their valid lifetimes expire, or until the CPE sends a DHCPv6 RELEASE message to the DHCPv6 server. The SmartEdge OS removes the affected routes and releases the IPv6 prefixes when any of the following occur:

- The CPE releases the IPv6 prefix.
- The time limit on the prefix expires.
- The PPP session terminates.
- The IPv6 stack is disabled.

The DHCPv6 server sends delegated IPv6 prefixes and the following DNS information to the CPE:

- IPv6 address of a DNS name server
- Domain name for DNS resolution
- Number of seconds a client waits before refreshing the configuration information received from the DHCPv6 server
- Preference that the client should use for this DHCPv6 server (in cases where multiple clients are requesting information from the router)

You can configure DNS information directly under a DHCPv6 server (in DHCPv6 server policy configuration mode) or inside a subnet configured under the DHCPv6 server (in DHCPv6 server policy subnet configuration mode). The DHCPv6 attributes configured inside a subnet are applicable to that subnet only. When you configure a subnet:



- Options that are configured for a particular subnet take precedence over options configured under the top-level DHCPv6 server (in DHCPv6 server policy configuration mode).
- Only those options administratively configured for a subnet differ from the options configured in the top-level DHCPv6 server policy. For example, if you configure all options except for the domain name for DNS resolution in a subnet, clients in that subnet use the domain name configured under the top-level DHCPv6 server.
- If you do not specify a particular DHCPv6 policy option for the subnet (in DHCPv6 server policy subnet configuration mode), the subnet configuration matches the top-level DHCPv6 server policy configuration (as specified in DHCPv6 server policy configuration mode).

1.2.8.2 Stateless DHCPv6

With stateless DHCPv6, the DHCPv6 server sends only the following DNS information to the CPE:

- IPv6 address of a DNS name server
- Domain name for DNS resolution
- Number of seconds a client waits before refreshing the configuration information received from the DHCPv6 server
- Preference that the client should use for this DHCPv6 server (in cases where multiple clients are requesting information from the router)

In a stateless configuration, the DHCPv6 server does not maintain dynamic state of each client or delegate IPv6 prefixes to clients.

Note: With stateless DHCPv6, only those DNS options specified in the top-level DHCPv6 server policy (in DHCPv6 server policy subnet configuration mode) are applicable; stateless DHCPv6 does not support subnets.

1.2.9 Neighbor Discovery Protocol for IPv6

The SmartEdge router uses the Neighbor Discovery (ND) protocol to assign an IPv6 prefix to the WAN link of the CPE router. ND obtains the IPv6 prefix from:

- The Framed-IPv6-Prefix attribute, which can be configured statically or by using the Framed-IPv6-Prefix RADIUS attribute
- A shared IPv6 prefix pool

Setting the Framed_IPv6_prefix to all 0s or all 1s indicates that the IPv6 prefixes come from a configured shared IPv6 prefix pool.

In addition to IPv6 prefix assignment, the CPE uses ND to:



- Determine the link-layer addresses of nodes on a link
- Find available routers
- Maintain reachability information about the paths to active neighbors

ND provides Duplicate Address Detection (DAD) and media-independent address resolution of on-link nodes.

For IPv6 subscriber services, the ND attributes are assigned in one of two ways:

- From the global default ND profile.
- From an administratively configured ND profile referenced in the subscriber profile or record. ND profiles are context-specific and cannot be applied directly to interfaces. A subscriber must be associated with the same context the profile is associated with before the profile can be used. If you do not reference an ND profile in a subscriber profile or record, the router automatically assigns a default ND profile (called the GLOBAL_DEFAULT_PROFILE) to the subscriber circuit.

Use the *show nd profile* command to see which profile a subscriber circuit is using for ND; use the *show nd profile GLOBAL_DEFAULT_PROFILE* command to see the default configuration used by the GLOBAL_DEFAULT_PROFILE.

Note: Router ND, which is configured under an individual interface and applies ND properties to the specified interface, is not supported for IPv6 subscriber services. Router ND is applicable for router-to-router connections only.

ND supports Stateless Address Autoconfiguration (SLAAC), which enables subscribing hosts to automatically configure global IPv6 addresses on their interfaces. SLAAC uses ND to advertise an IPv6 prefix or group of prefixes on-link. The host automatically configures its interface address by appending the host interface ID to the IPv6 prefix.

Note: SLAAC is automatic on any IPv6 prefix that is configured.

The SmartEdge OS uses its own interface ID to generate the link local-address on the WAN link.

The SLAAC process is as follows:

- 1 The host sends an ND Router SOLICIT multicast message soliciting an RA. The RA contains information about on-link prefixes and whether they are available or unavailable for SLAAC.
- 2 The router (which is listening for SOLICIT messages) responds to the host with a Router Advertisement (RA) message that contains the IPv6 prefix or group of prefixes identifying the interface. Any prefix advertised in an RA message has SLAAC enabled, and the host can use that IP prefix to auto-generate its IP address.



- 3 For IPv6 sessions, both ends of the PPP circuit exchange their interface IDs through IPv6CP negotiation. If a subscriber cannot generate its own interface ID, the subscriber takes its interface ID from the subscriber record in the RADIUS database (if the record contains a client interface ID). If the subscriber does not generate its own interface ID and an interface ID is not available in the RADIUS database, PPP randomly generates an interface ID. If the session also has an IPv4 stack, the router assigns an IPv4 address to the subscriber through IPCP.
- 4 Before assigning the IPv6 address to the interface, the host performs DAD on the candidate IPv6 address. If the SmartEdge OS detects a duplicate address, it logs an error message in the system log.

Note: How the CPE responds to duplicate-address detection depends on the type of equipment.
- 5 The SmartEdge OS installs the global IPv6 address prefixes (the framed IPv6 prefixes) in the RIB.

SLAAC is supported for all IPv6 (both subscriber and nonsubscriber) circuits, including access link aggregation groups .

Note: For more information about how ND works, see *Configuring ND*.

1.2.10

Statically Mapped DHCPv6 PD Prefixes

You can statically map one or more DHCPv6 PD prefixes to a specified subscriber DUID and, optionally, Identity Association Identifier (IAID). Use the *prefix duid* command in DHCPv6 server policy configuration mode to map a particular prefix to a DUID and, if desired, IAID. That prefix is delegated only to subscribers with a matching DUID (and IAID, if required).

When the router receives a request from a client, DHCPv6 PD checks whether the DUID and IAID of the client match the configuration for any statically-mapped IPv6 prefixes configured under the DHCPv6 server. If a match is present, the matching IPv6 prefixes are returned to the client. If no match is found, DHCPv6 PD attempts to assign prefixes for the client from other sources.

A DUID is a unique identifier included in DHCPv6 messages and used to identify a device. The IAID identifies a collection of addresses assigned to a subscriber. An individual subscriber can have multiple IAIDs assigned.

Note: The number of prefixes that can be assigned to a subscriber is limited by the Delegated-Max-Prefix value. If the number of matching prefixes is greater than the Delegated-Max-Prefix value, the SmartEdge router arbitrarily chooses which prefixes are assigned to the subscriber.

Consider the following when configuring statically mapped DHCPv6 PD prefixes:



- Statically mapped prefixes must not fall within the prefix range of any prefix pools configured on the system.
- Statically mapped prefixes are included in the Delegated-Max-Prefix value.

1.2.11 IPv6 Prefix Pools

Instead of statically configuring ND and DHCPv6 PD prefixes, you can configure subscribers to obtain ND and DHCPv6 PD prefixes from pools that lease IPv6 prefixes to subscribers. The SmartEdge OS supports two types of IPv6 prefix pools:

- DHCPv6 PD prefix pool (used to by DHCPv6 PD to assign IPv6 prefixes to subscribers).
- Shared IPv6 prefix pools (used by ND to assign IPv6 prefixes to the WAN link between the SmartEdge router and the CPE).

An individual last-resort multibind interface can support up to 1024 shared IPv6 or DHCPv6 PD prefix pools. A non-last-resort multibind interface supports a maximum of 16 shared IPv6 or DHCPv6 PD prefix pools under the primary IPv6 prefix of that interface.

Pool counters track prefix assignment for a context and for individual pools; counters are updated each time a prefix is assigned or released. Pool counters are checked against a predefined falling threshold. When the total number of available IP addresses in a particular pool or context equals the specified value, the router generates an alert (or *crossing event*) that is recorded as either or both of the following:

- A trap
- A log entry

If the number of available IPv6 prefixes becomes greater than the specified value before dropping again to the falling threshold value, a second crossing event is generated, and so on.

You can configure falling-threshold parameters for both shared IPv6 prefixes and DHCPv6 PD pools with the *ipv6 pool* command. To configure the falling threshold value for:

- A particular pool, use the command in interface configuration mode.
- All pools in a context, use the command in context configuration mode.

Consider the following when configuring falling threshold values for IP pools:

- Threshold values defined for a particular pool (in interface configuration mode) take precedence over threshold values defined for a context (in context configuration mode).



- If you configure falling threshold parameters for a context or pool that already has threshold parameters configured, the falling threshold parameters are set to the new values.
- Specifying a threshold value that is larger than the sum of all IPv6 prefixes in all IP pools in the context ensures that no threshold event ever occurs at the context level.
- Specifying a falling threshold value of **0** (in the `ipv6 pool` command string) removes the threshold configuration from a context or pool.
- An individual pool or pool group (all pools within a context) can have up to two falling thresholds defined. If two thresholds are defined, the second threshold must be less than the first threshold.

To see information related to threshold logs and traps:

- Use the `show ipv6 pool` command to determine whether a threshold crossing event has been generated.
- Use the `show dhcpv6 log` command to determine whether any pool-related traps have been generated.
- Use the `show log | grep pool` or `show log | grep threshold` command to view the logs generated for DHCPv6 PD and shared IP pools.

Note: Pool usage information is not stored in shared memory. After a router restart or switchover, the pool usage information must be rebuilt based on the subscriber information stored in shared memory.

1.2.11.1 DHCPv6 PD Pools

DHCPv6 PD pools contain a range of IPv6 prefixes leased to subscribers as needed. You configure a DHCPv6 PD pool on an individual server under a multibind interface. When building an IPv6 stack, a subscriber requests an IPv6 address from the DHCPv6 server. The DHCPv6 server searches the DHCPv6 PD pool to see if any IPv6 prefixes can be leased out to the subscriber. The DHCPv6 PD server then sends the client:

- One or more IPv6 prefixes. The number of entries sent back to the subscriber is determined by the Delegated-Max-Prefix attribute configuration for that subscriber.
- The prefix allocation length configured under the DHCPv6 PD pool.

DHCPv6 PD subscribers can specify a prefix hint in the REQUEST messages sent to the DHCPv6 PD server. A hint is an IPv6 prefix suggested by the client. If the suggested prefix is available, the DHCPv6 PD server delegates that prefix to the client. If the suggested prefix (the hint) is not available, the DHCPv6 PD server delegates another prefix from the pool to the client (if a prefix is available). If no prefixes are available in the pool, the DHCPv6 PD server delegates a prefix from another pool (if more than one DHCPv6 PD pool



is configured). A client can specify either of the following hints to the DHCPv6 PD server:

- The 128-bit part or length of the IPv6 prefix
- The number of IPv6 prefix entries allowed

Consider the following rules when configuring DHCPv6 PD pools:

- You can configure up to 1024 DHCPv6 PD pools under an individual last-resort interface.
- You can configure up to 16 DHCPv6 PD pools under an individual non-last-resort interface. To support DHCPv6 PD pools, you must configure the non-last-resort interface with a primary and secondary IPv6 prefix and prefix length.
- A pool must fall into the subnet of the interface the pool is bound to.
- Pools track assigned and available IPv6 prefixes; the same prefix cannot be leased out to two different subscribers at the same time.
- DHCPv6 PD pools are deleted when a subscriber interface is deleted.
- DHCPv6 PD pools are not deleted when an interface prefix is deleted.
- The DHCP-PD pool uses a round-robin algorithm to allocate prefixes.
- When DHCPv6 PD pools are deleted, subscribers using the deleted pools retain their subscriber sessions.
- For static DHCPv6 PD prefixes for the DUID under the DHCPv6 server:
 - If an IAID is not configured, clients matching the DUID obtain the DHCPv6 PD prefix.
 - If an IAID is configured, the client must match both the DUID and the IAID to obtain the DHCPv6 PD prefix.
 - Multiple IPv6 prefixes can be allocated for a single DUID.

The following restrictions apply to DHCPv6 PD pools:

- For dual-stack interfaces, the IPv6 stack does not come up if:
 - The Framed-IPv6-pool and Framed-ip-pool attributes that map an interface do not match.
 - The interfaces that are mapped by Framed-ip-pool and Framed-IPv6-pool attributes do not match.

In such instances, only the IPv4 stack comes up.

- Pools cannot overlap.



- A static DHCPv6 PD prefix that is configured under a DHCPv6 server policy is rejected if that prefix falls within the prefix range specified for a DHCPv6 PD pool.
- A RADIUS-configured static DHCPv6 PD prefix can fall within any range. However, if the RADIUS-configured PD prefix does not fall within the specified range, the SmartEdge router rejects the RADIUS-configured static PD prefix during subscriber authentication.

If a RADIUS-configured PD prefix falls into a specified DHCPv6 prefix range, and if that prefix is already used by another subscriber, the RADIUS-configured static PD prefix is used as a hint for obtaining a PD prefix from that PD pool.

- The primary and secondary IPv6 prefix of an interface are reserved and cannot be used by any subscribers.
- DHCPv6 PD pool prefix can be applied only to one active subscriber at a time. If no prefixes are available in a pool:
 - Subscribers requesting prefixes from that pool are torn down if *session-action dual-stack-failure force-down* is enabled.
 - The IPv6 stack is torn down if *session-action dual-stack-failure force-down* is disabled.
- Different types of pools cannot be shared; DHCPv6 PD pools and shared IPv6 prefix pools cannot overlap.

1.2.11.2 Shared IPv6 Prefix Pools

ND can obtain IPv6 prefixes from shared IPv6 prefix pools. A shared IPv6 prefix pool is configured directly under a multibind interface that has a primary IPv6 prefix and prefix length configured. The IPv6 prefix of a shared IPv6 prefix pool must fall within the primary prefix of the interface under which the pool is configured.

To delete a shared IPv6 prefix pool, you must delete the interface or the IPv6 prefix of the interface bound to that pool. When you delete an interface bound to a shared IPv6 prefix pool, all subscriber sessions on that interface are torn down.

To assign IPv6 prefixes from a shared IPv6 prefix pool, configure the `Framed_IPv6_prefix` attribute to be all 0s or all 1s.

You can optionally specify a name for a shared IPv6 prefix pool. If you do not administratively specify a name, ND automatically assigns the name of the parent interface to the shared IPv6 prefix pool. When configured under a last-resort interface, multiple shared pools can have the same name. Shared pools configured under different interfaces cannot have the same name if those interfaces exist in the same context.



Consider the following rules when configuring shared IPv6 prefix pools:

- You can configure up to 1024 shared IP pools under an individual last-resort interface.
- You can configure up to 16 shared IP pools under an individual non-last-resort interface. To support shared IPv6 pools, you must configure the non-last-resort interface with a primary and secondary IPv6 prefix and prefix length.
- A single shared ND pool can support up to 64,000 IPv6 prefixes.
- Shared IPv6 prefix pools support prefix lengths in the range from 32 to 64 bits.
- Different pools can have the same pool name. Pools that have the same name and exist under the same interface must have the same prefix length.
- A pool must fall into the subnet of the interface that the pool is bound to.
- Pools track assigned and available IPv6 prefixes; the same prefix cannot be leased out to two different subscribers at the same time.
- A prefix from a shared IPv6 prefix pool can be applied only to one active subscriber at a time. If no prefixes are available in a pool:
 - Subscribers requesting prefixes from that pool are torn down if *session-action dual-stack-failure force-down* is enabled.
 - The IPv6 stack is torn down if **session-action dual-stack-failure force-down** is disabled.
- A shared IPv6 prefix pool uses a round-robin algorithm to allocate prefixes.
- When shared IPv6 prefix pools are deleted, subscribers using the deleted pools retain their subscriber sessions.

The following restrictions apply to shared IPv6 prefix pools:

- Although last-resort multibind interfaces can simultaneously support multiple shared IPv6 prefix pools, a single non-last-resort multibind interface can support one shared IPv6 prefix pool only.
- For dual-stack interfaces, the IPv6 stack does not come up if:
 - The Framed-IPv6-pool and Framed-ip-pool attributes that map an interface do not match.
 - The interfaces that are mapped by the Framed-ip-pool and Framed-IPv6-pool attributes do not match.

In such instances, only the IPv4 stack comes up.

- Pools cannot overlap.



- Shared IPv6 prefix pools configured under different interfaces cannot have the same name if those interfaces exist in the same context.
- Do not delete an ND IPv6 prefix pool while there are still subscribers using that pool because the subscriber sessions will be dropped.
- Static ND prefixes that are configured under a RADIUS record can fall within any range. However, if a RADIUS-configured static ND prefix does not fall within the range allowed by any shared IPv6 prefix pool, and if that prefix is already being used by another subscriber, the SmartEdge router rejects the IPv6 stack at authentication time. If the RADIUS-configured static ND prefix falls within the range allowed by a shared IPv6 prefix pool, that prefix is used as a hint for obtaining an ND prefix from that shared IPv6 prefix pool. A static ND prefix that is configured under a RADIUS record can fall within any shared pool prefix range.
- The primary and secondary IPv6 prefix of an interface are reserved and cannot be used by any subscribers.
- Shared IPv6 prefix pools cannot overlap with DHCPv6 PD pools.

1.2.12

Duplicate Prefix and Address Errors

The SmartEdge OS detects any duplicate IPv4 addresses and IPv6 prefixes during session authentication. Duplicate address conflicts can occur between IPv4 addresses, IPv6 framed prefixes, and DHCPv6 PD prefixes.

With dual-stack subscribers, the IPv4 and IPv6 sessions function independently of one another. By default, if a duplicate IPv4 address or IPv6 prefix is detected during the authentication phase for a dual-stack subscriber:

- Authentication fails for the affected stack only.
- The unaffected stack is authenticated, and the subscriber remains in a state where only one stack is active for the life of the session.

You can use the *session-action dual-stack-failure force-down* command in subscriber configuration mode to modify the default behavior so that the entire dual-stack session fails if the router detects duplicate IPv4 addresses or IPv6 prefixes during session authentication.

Note: The SmartEdge OS prevents you from configuring duplicate addresses or prefixes with shared IP and DHCPv6 PD pools.

Table 6 describes the types of prefix conflicts that can occur during dual-stack session authentication, and the action the SmartEdge router takes in response to those errors. With these errors, only the session for the affected stack is brought down unless you use the **session-action dual-stack-failure force-down** command to bring down the entire dual-stack session (for both stacks).



Table 6 Dual-stack Session Authentication Errors and Actions Taken

Type of Conflict	Action
Authenticating subscriber is assigned a static IPv4 address that is already assigned to another subscriber	Authentication fails and the router sends a No-Accounting-Start message to the subscriber.
Authenticating subscriber is assigned a static IPv4 address that matches an assigned IPv4 address from an IP pool.	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber is assigned a static framed IPv6 prefix that is already assigned to another subscriber.	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber is assigned a framed IPv6 prefix that matches an assigned framed IPv6 prefix from an IP pool.	The framed IPv6 prefix is treated as a "hint" for the authenticating subscriber.
Authenticating subscriber is assigned a static delegated IPv4 prefix that is already assigned to another subscriber.	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber is assigned a delegated IPv6 prefix that is already assigned to another subscriber.	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber is assigned a delegated IPv6 prefix that matches an assigned delegated IPv6 prefix from an IP pool.	The delegated IPv6 prefix is treated as a "hint" for the authenticating subscriber.
The interfaces mapped by the Framed-ip-pool and Framed-IPv6-pool attributes do not match (the Framed-ip-pool and Framed-IPv6-pool must be mapped to the same interface).	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber has a matching ND and DHCPv6 PD prefix statically configured.	Authentication fails, and the router sends an Authentication-fail message to the subscriber.
Authenticating subscriber has a distinct framed IP prefix, but is assigned a static delegated prefix that is already assigned to another subscriber.	Authentication subscriber session fails.



When a session fails because duplicate IPv4 address is detected, the router sends an Accounting-Stop message to the RADIUS server with the following error codes:

- Session-Error-Code 236
- Session-Error-Message—Duplicate IP address detected

When a session fails because duplicate IPv6 prefixes are detected, the router sends an Accounting-Stop message to the RADIUS server with the following error codes:

- Session-Error-Code 237
- Session-Error-Message—Duplicate IPv6 address detected

1.2.13 QoS Support for IPv6 Subscribers

QoS is supported on IPv6 subscriber interfaces.

For information about how to configure QoS, see the following QoS documents:

- *Configuring Circuits for QoS*
- *Configuring Flow Admission Control*
- *Configuring Rate-Limiting and Class-Limiting*
- *Configuring Queuing and Scheduling*

1.2.14 Using IP ACLs for Traffic Control and IPv6 Protection

You can configure IP ACLs for IPv6 administrative protection on traffic card circuits, the Ethernet management port, and administrative traffic. Policy ACLs are also supported for IPv6 traffic. For information on how to configure IP ACLs to support IPv6, see *Configuring ACLs*.

1.3 Overview of PPP Session Establishment

When an IPv6 host or CPE initiates a PPP session with a SmartEdge router, the session establishment process is as follows:

- 1 A CPE initiates a PPP session with a subscriber network.
- 2 A SmartEdge router receives the request and creates a PPP session (single-stack or dual-stack) between the BRAS and the subscriber.
- 3 If the session has an IPv4 stack, the router assigns an IPv4 address to the subscriber through IPCP.



- 4 An ND RA advertises 0 or more IPv6 framed prefixes on the link.
- 5 The SmartEdge router installs a route for that IPv6 prefix on the link between the SmartEdge router and the CPE.
- 6 If the subscriber sends a DHCPv6 SOLICIT to the SmartEdge router, the router uses DHCPv6 PD to assign a delegated IPv6 prefix and DNSv6 to the subscriber.
- 7 IPv6 (and IPv4, if dual-stack) traffic is routed through the SmartEdge router.



2 Configuration and Operations Tasks

This section describes the requirements, restrictions, configuration tasks, and operations tasks for configuring IPv6 subscriber services on the SmartEdge router. For information about how to troubleshoot IPv6 subscriber services, see *Troubleshooting IPv6 and Dual-Stack Subscriber Services*.

2.1 Recommendations

If the subscriber is a router, we recommend assigning subscribers a /64, /56, or /48 PD prefix that can be further subdivided on downstream interfaces.

2.2 Requirements

The SmartEdge router and the CPE must each have at least one link local-address each.

2.3 Restrictions

- The SmartEdge router supports IPv6 subscriber services for PPP subscribers only.
- For routing IPv6 packets, the end system is responsible for fragmenting the packets, based on the MTU for the data path. If the size of an IPv6 control packet exceeds the Path Maximum Transmission Unit (PMTU) determined by Neighbor Discovery (ND) for the data path, the control or ASE card fragments the packet to match the PMTU value. If the size of an IPv6 data packet exceeds the MTU set on the egress port on the traffic card, the traffic card fragments the IPv6 data packet to match the MTU value. In the current release of the SmartEdge OS, IPv6 data packets are created on traffic cards for BGF over IPv6.



Warning!

Risk of data loss for IPv6 data packets created by a traffic card. IPv6 data packets created on a traffic card are fragmented based on the MTU set on the egress port. In the current release of the SmartEdge OS any ICMPv6 "Packet too big" message sent from any IPv6 router on the data path to the traffic card that created an IPv6 data packet if the fragment size exceeds the PMTU of the data path is ignored. As a result, the data packet traffic is dropped. (ICMPv6 "Packet too big" messages sent from any traffic card or IPv6 router on the data path to the control or ASE card that created an IPv6 control packet if the fragment size exceeds the PMTU are acted upon.)

To avoid this risk, ensure that the MTU set on the traffic card with the `mtu` command in port configuration mode is set to be less than the PMTU of the data path. If a network-wide PMTU policy is used, set matching port-level MTU and network-wide PMTU values.

- LAC and LNS IPv6 and LAC and LNS dual-stack (IPv4 and IPv6) traffic is supported on IPv4 L2TP tunnels only. IPv6 L2TP tunnels are not supported on the SmartEdge router. When PPP sessions are terminated on a SmartEdge LNS, IPv6 packets are not fragmented on the LNS. The IPv6 packet is encapsulated in the IPv4 tunnel and the IPv4 tunnel packets are fragmented.
- IPv6 subscriber services are supported only for connections between the SmartEdge router and a CPE router that has a GUA configured.
- With stateful DHCPv6, the DHCPv6 server is used for DHCPv6 PD only.
- With stateless DHCPv6, the DHCPv6 server sends only the following DNS information to the CPE:
 - IPv6 address of a DNS name server.
 - Domain name for DNS resolution.
 - Number of seconds a client waits before refreshing the configuration information received from the DHCPv6 server.

2.4 Configuring a SmartEdge Router to Provide IPv6 and Dual-Stack Subscriber Services

The steps that follow provide a high-level overview of the tasks required to configure IPv6 and dual-stack subscriber services. Detailed configuration procedures for each task follow.



To configure a SmartEdge router to provide IPv6 and dual-stack subscriber services:

- 1 Optional. If using RADIUS to authenticate a subscriber, you can configure the NAS-IPV6-Address to match the IPv6 address of the NAS. See *Configuring the NAS-IPV6-Address to Match the IPv6 address of the NAS (Optional)*.

If you are not using RADIUS to authenticate a subscriber or do not want to configure the NAS-IPV6-Address to match the IPv6 address of the NAS, do not perform this step.

- 2 Configure an interface with a GUA on the link between the SmartEdge router and the CPE. See *Configuring the WAN Link*.
- 3 If you do not want to use the default ND profile, create and configure an ND profile. To configure an ND profile for IPv6 subscribers, see *Configuring ND Attributes for IPv6 Subscribers*. To configure an ND profile for IPv4 subscribers, see *Configuring ND*.
- 4 Optional. Configure shared IPv6 prefix pools. See *Configuring Shared IPv6 Prefix Pools (Optional)*.
- 5 If using a DHCPv6 server to assign IPv6 prefixes to subscribers, create and configure the DHCPv6 server policy on the SmartEdge router, as described in *Configuring a DHCPv6 Server Policy*.
- 6 Configure one or more multibind interfaces to use the DHCPv6 server policy. These interfaces are called "DHCPv6 servers." See *Configuring a DHCPv6 Server*.

Note: To use a DHVPv6 server policy, the DHCPv6 server interfaces must be configured within the same context as the DHCPv6 server policy.

- 7 Optional. Configure DHCPv6 PD pools. See *Configuring a DHCPv6 PD Pool*.
- 8 Enable AAA subscriber authentication locally or through a RADIUS server. See *Enabling AAA Subscriber Authentication*.
- 9 If you are using the local database for subscriber authentication, configure the subscriber attributes in a subscriber record. You can also configure a subset of subscriber attributes in a default or named subscriber profile. See *Configuring the Subscriber Attributes*.

Note: If using a non-default ND profile, reference the ND profile you created in Step 5 in the subscriber record or profile.

If you are using a RADIUS server for subscriber authentication, skip this step and go to step 10.



- 10 Configure PPP or PPP over Ethernet (PPPoE) encapsulation on the WAN link and then bind the circuit using CHAP or PAP. The circuit is now ready to perform subscriber services.

For more information on configuring PPP and PPPoE, see *Configuring PPP and PPPoE*. To see how to configure the type of circuit you are using for your WAN link, see the appropriate section in *Configuring Circuits*.

- 11 Optional. Configure link access groups (LAGs) to bundle circuits carrying subscriber traffic. You can configure network-facing LAGs between pairs of provider edge routers at each end of the core network for redundancy, resiliency, and route distribution and subscriber-facing LAGs between pairs of customer edge routers for the same reasons plus securing the L2VPN gateways.

For more information about configuring LAGs, see *Configuring Link Aggregation*.

For information about how to troubleshoot IPv6 subscriber services, see the *Troubleshooting IPv6 and Dual-Stack Subscriber Services* document.

2.4.1 **Configuring the NAS-IPV6-Address to Match the IPv6 address of the NAS (Optional)**

This optional task is applicable to routers using RADIUS to authenticate subscribers. To configure the NAS-IPV6-Address to match the IPv6 address of the NAS:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.
- 3 Use the *radius attribute NAS-IPV6-Address interface* command to configure the IPv6 address of the NAS for RADIUS access-request and access-accounting messages.

2.4.2 **Configuring the WAN Link**

To configure an interface with a GUA on the link between the SmartEdge router and the CPE:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.
- 3 Use the *ipv6 address* command to specify an IPv6 GUA.

2.4.3 **Configuring a DHCPv6 Server Policy**

To configure DHCPv6 service policy attributes:



- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to access context configuration mode.
- 3 Use the *dhcpv6 server* command to create a DHCPv6 server policy and access DHCPv6 server policy configuration mode. Only one DHCPv6 server policy is allowed for a context.
- 4 Use the *option domain-name-server* command as follows to specify the IP address of a DNS name server:

```
option domain-name-server server-address
```

- 5 Use the *option domain-search* command as follows to specify a domain name for DNS resolution:

```
option domain-search domain-name
```

- 6 Use the *option information-refresh-time* command as follows to specify the number of seconds a client waits before refreshing the configuration information received from the DHCPv6 server:

```
option information-refresh-time seconds
```

Range is from 600 through 4294967295 seconds.

- 7 Use the *option preference* command as follows to configure the preference value for this DHCPv6 server:

```
option preference integer
```

A DHCPv6 server with a lower value is preferred over a server with a higher value.

Range is from 0 through 255.

- 8 Use the *option rapid-commit* command to enable Rapid Commit for faster IPv6 prefix delegation.

With the RAPID COMMIT option, only two messages (SOLICIT and REPLY messages) are exchanged between the DHCPv6 server and the CPE. We recommend using the RAPID COMMIT option when there is only one server for a client to connect to.

- 9 Use the *prefix duid* command to statically map a specified IPv6 prefix to a DUID or DUID and IAID.
- 10 Use the *prefix lifetime* command as follows to configure the length of time the subscriber router can use a delegated IPv6 prefix and a given DHCPv6 prefix:

```
prefix lifetime {preferred seconds valid seconds | infinite}
```



Set the prefix lifetime as follows:

- **preferred seconds**—Number of seconds the IPv6 addresses using the delegated IPv6 prefix are preferred. Range is from 600 through 4294967294 seconds.
- **valid seconds**—Number of seconds a delegated IPv6 prefix is valid and can be used by a client. Range is from 600 through 4294967294 seconds.
- **infinite**—Configures both the preferred and valid lifetimes to be infinite.

11 If required, configure a subset of DHCPv6 attributes that apply to a particular subnet only. Options configured for the subnet take precedence over options specified in the top-level DHCPv6 server policy:

- Use the *subnet* command as follows to access DHCPv6 server policy subnet configuration mode to configure DHCPv6 server attributes that are applicable only to subscribers in the specified subnet:

```
subnet ipv6-prefix/subnet-mask [name subnet-name ]
```

Only those options administratively configured for a subnet differ from the options configured in the top-level DHCPv6 server policy (in DHCPv6 server policy configuration mode). If you do not specify a particular DHCPv6 policy option for the subnet (in DHCPv6 server policy subnet configuration mode), the subnet configuration matches the top-level DHCPv6 server policy configuration (as specified in DHCPv6 server policy configuration mode).

Replace the *ipv6-prefix* argument with an IPv6 prefix. This IPv6 prefix cannot be the same as the prefix for any other interface.

- Use the *option domain-name-server* command as follows to specify the IP address of a DNS name server.

```
option domain-name-server server-address
```

- Use the *option domain-search* command as follows to specify a domain name for DNS resolution:

```
option domain-search domain-name
```

- Use the *prefix lifetime* command as follows to configure the length of time the subscriber router can use a delegated IPv6 prefix and a given DHCPv6 prefix:

```
prefix lifetime {preferred seconds valid seconds /  
infinite}
```

Set the prefix lifetime as follows:



- **preferred seconds**—Number of seconds the IPv6 addresses using the delegated IPv6 prefix are preferred. Range is from 600 through 4294967294 seconds.
- **valid seconds**—Number of seconds a delegated IPv6 prefix is valid and can be used by a client. Range is from 600 through 4294967294 seconds.
- **infinite**—Configures both the preferred and valid lifetimes to be infinite.

2.4.4 Configuring a DHCPv6 Server

To configure a multibind interface to be the DHCPv6 server:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.
- 3 Use the *interface* command as follows to configure a multibind interface, and access interface configuration mode:

```
interface name multibind [lastresort]
```

This is the interface you want to configure to be DHCPv6 enabled.

- 4 Use the *dhcpv6 server* command as follows to configure an interface to be a DHCPv6 server interface:

```
dhcpv6 server {ipv6-addr | interface}
```

You can configure the DHCPv6 server to use the primary IPv6 address of the interface as the server IP address or specify an IPv6 address for it.

2.4.5 Configuring a DHCPv6 PD Pool

To configure a DHCPv6 PD pool:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.
- 3 Use the *interface* command as follows to configure a multibind interface, and access interface configuration mode:

```
interface name multibind [lastresort]
```

This is the interface you want to configure to be DHCPv6 enabled.

- 4 Use the *ipv6 address* command to specify an IPv6 address for the multibind interface.



- 5 Use the `ipv6 pool dhcpv6` command to create pool of DHCPv6 PD prefixes under the multibind interface and, optionally, create a threshold value for which a crossing event occurs. The command syntax is:

```
ipv6 pool dhcpv6 {{starting-prefix/prefix_length  
last-prefix/prefix_length} [name pool-name] [threshold  
{absolute | percentage} falling first-threshold {trap [log] |  
log [trap]}] [second-threshold {trap [log] | log [trap]}]}
```

Note: DHCPv6 threshold configuration for a particular pool (in interface configuration mode) takes precedence over DHCPv6 PD threshold configuration in context configuration mode.

2.4.6 Configuring Pool Thresholds for a Context

To optionally configure pool thresholds that apply to all DHCPv6 PD or shared IPv6 pools in a context:

- 1 Use the `configure` command to access global configuration mode.
- 2 Use the `context` command to enter context configuration mode.
- 3 Use the `ipv6 pool` command to configure pool thresholds that apply to all DHCPv6 PD or shared IPv6 pools in a context. The command syntax is:

```
ipv6 pool {[dhcpv6] threshold {absolute | percentage} falling  
first-threshold {trap [log] | log [trap]}] [second-threshold {trap  
[log] | log [trap]}]}
```

Note: Threshold configuration for a particular pool (in interface configuration mode) takes precedence over threshold configuration in context configuration mode.

2.4.7 Enabling AAA Subscriber Authentication

To enable AAA subscriber authentication:

- 1 Use the `configure` command to access global configuration mode.
- 2 Use the `context` command to enter context configuration mode.
- 3 Use the `aaa authentication subscriber` command to enable AAA to authenticate subscribers through the SmartEdge router local database or RADIUS. The command syntax is:

```
aaa authentication subscriber [local | radius]
```

2.4.8 Configuring ND Attributes for IPv6 Subscribers

For IPv6 subscriber services, the SmartEdge router acquires ND attributes in one of two ways:



- From the global default ND profile
- From an administratively configured ND profile referenced in a subscriber profile or record

Note: If you do not reference an ND profile in a subscriber profile or record, the router automatically assigns a default ND profile (called the GLOBAL_DEFAULT_PROFILE) to the subscriber circuit. Use the *show nd profile* command to see which profile a subscriber circuit is using for ND; use the *show nd profile GLOBAL_DEFAULT_PROFILE* command to see the default configuration used by the GLOBAL_DEFAULT_PROFILE.

To create and configure an ND profile for IPv6 subscribers:

- 1 Access context configuration mode.
- 2 Access ND profile configuration mode.
- 3 Use the *ra-interval* command as follows to configure the interval between transmissions of RA messages:

```
ra-interval seconds
```

Note: Setting the RA interval to 0 suppresses the sending of RAs.

- 4 Use the *ra lifetime* command as follows to configure the router advertisement lifetime in seconds:

```
ra lifetime seconds
```

Replace *seconds* with the total number of seconds the prefix remains valid.

- 5 Use the *ra managed-config* command to configure the router advertisement to contain the managed address configuration flag. This flag is included in IPv6 RAs, indicating to hosts that they should use the managed (stateful) protocol for address autoconfiguration in addition to any addresses autoconfigured using stateless address autoconfiguration.
- 6 Use the *ra other-config* command to configure the router advertisement to contain the other stateful configuration flag. This flag is included in IPv6 router advertisements, indicating to hosts that they should use the administered (stateful) protocol to obtain autoconfiguration information other than addresses.
- 7 Use the *ns-retry-interval* command as follows to configure the Retrans Timer, which dictates the length of time between retransmitted Neighbor Solicitation (NS) messages:

```
ns-retry-interval milliseconds
```



- 8 Use the *dad-transmits* command to specify the number of Neighbor Solicitation (NS) messages the SmartEdge router sends to its peers for DAD:

```
dad-transmits num-dad-transmits
```

. Replace *num-dad-transmits* with the number of DAD NS messages to send; the range of values is 0 to 3. A value of 0 disables NS message transmission.

- 9 Use the *proto-down-on-dad* command to enable the SmartEdge router to send a request to bring down the IPv6 stack of the subscriber circuit in which a DAD failure is detected.
- 10 Use the *reachable-time* command as follows to configure the Reachable Time value, which is the length of time this ND router or ND router interface assumes that a neighbor is reachable:

```
reachable-time milliseconds
```

This attribute enables the router to detect unavailable neighbors. The reachable time value is advertised by the RA messages sent by the router.

- 11 Use the *preferred-lifetime* command as follows to configure the lifetime of the preferred router advertisement:

```
preferred-lifetime seconds
```

Replace *seconds* with the length of time (in seconds) an advertised prefix remains preferred.

- 12 Use the *valid-lifetime* command as follows to configure the router advertisement to list a specified prefix for a valid lifetime:

```
valid-lifetime seconds
```

Replace *seconds* with the length of time the addresses generated from the prefix remain valid.

Note: The SmartEdge router does not support the use of router ND (where ND is configured under a specific interface) for IPv6 subscriber services. Any router ND configuration that exists under an interface is ignored for subscribers bound to that interface.

2.4.9 Configuring Shared IPv6 Prefix Pools (Optional)

To configure a shared IPv6 prefix pool:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.



- 3 Use the *interface* command as follows to configure a multibind interface, and access interface configuration mode:

```
interface name multibind [lastresort]
```

This is the interface you want to configure to host the shared IP prefix pool.

- 4 Use the *ipv6 address* command to specify an IPv6 address for the multibind interface.
- 5 Use the *ipv6 pool* command to create pool of IPv6 prefixes under the multibind interface.

To optionally specify context-specific falling-threshold parameters that apply to all shared pools in the context:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to enter context configuration mode.
- 3 Create a threshold value for which a crossing event occurs. The command syntax is:

```
ipv6 pool {[dhcipv6]threshold {absolute | percentage} falling  
first-threshold {trap [log] | log [trap]} [second-threshold {trap  
[log] | log [trap]}
```

Note: Any threshold configuration for a particular pool (in interface configuration mode) takes precedence over threshold configuration in context configuration mode.

2.4.10 Configuring IPv6 Subscriber Attributes

Attributes that are applicable to the subscribers themselves are configured:

- Through a subscriber record
- In a subscriber profile that is referenced by a subscriber record
- In a default subscriber profile

Subscriber attributes are applied to an IPv6 subscriber in one of the following ways:

- From a subscriber record that is configured for a specific subscriber.
- From the default subscriber profile applied to all IPv6 subscribers in the same context. Subscribers inherit the configuration specified in the profile. If you want to use the default profile, you must explicitly configure it in default subscriber profile configuration mode.
- From a specific subscriber profile referenced in a subscriber record. You must explicitly assign a named subscriber profile to a subscriber record. When assigned to a subscriber record, the values of the attributes in a



named subscriber profile override the identical attributes in the default profile.

That tasks that follow describes how to configure various IPv6-specific subscriber attributes in a subscriber record or profile. Perform these tasks in one of the following modes:

Table 7 Subscriber Attribute Configuration Modes

To configure attributes for:	Perform these tasks in:
A subscriber record	Subscriber configuration mode
A default subscriber profile	Default subscriber profile configuration mode
A named subscriber profile	Subscriber profile name configuration mode

Note: Attributes in the subscriber record take precedence over identical attributes configured in the named subscriber profile, and attributes in the named subscriber profile take precedence over identical attributes configured in the default subscriber profile.

To configure various IPv6-specific subscriber attributes in a subscriber record or profile:

- 1 Use the *configure* command to access global configuration mode.
- 2 Use the *context* command to access context configuration mode.
- 3 Use the *subscriber* command as follows to access subscriber configuration mode for the specified IPv6 subscriber:

```
subscriber {name | default | profile}
```

- 4 If you are configuring stateful DHCPv6 PD, use the *ipv6 delegated-prefix* command as follows to specify the delegated IPv6 prefix to use for DHCPv6 PD. If you are using stateless DHCPv6, skip this step:

```
ipv6 delegated-prefix ipv6-prefix
```

Note: The DHCPv6 delegated prefix attribute is configurable only in a subscriber record.

- 5 Use the *ipv6 delegated-prefix maximum* command as follows to configure the Delegated-Max-Prefix attribute (the maximum number of IPv6 prefixes that can be delegated to a subscriber, either statically or from a DHCPv6 PD pool):

```
ipv6 delegated-prefix maximum number_of_prefixes
```

Range is from 1 to 5; default is 1.

- 6 Use the *ipv6 framed-prefix* command as follows to specify the prefix that will be advertised to subscribers using ND:



```
ipv6 framed-prefix ipv6-prefix
```

Replace the *ipv6-prefix* argument with a prefix that does overlap with any other interface prefix.

Note: This command is available in IPv6 subscriber record configuration mode only; you cannot configure the `ipv6 framed-prefix` command in a subscriber profile.

- 7 Use the *ipv6 framed-route* command as follows to specify a static IPv6 route that will be installed for the subscriber:

```
ipv6 framed-route ipv6-prefix next-hop metric
```

Note: This command is available in IPv6 subscriber configuration mode only; you cannot configure the `ipv6 framed-route` command in a subscriber profile.

- 8 Use the *ipv6 framed-pool* command as follows to specify that a subscriber obtains its prefix from the specified shared IPv6 prefix pool:

```
ipv6 framed-pool [name]
```

Replace *name* with the name of a shared IPv6 prefix pool that is configured under the same context as the subscriber.

- 9 Use the *ipv6 nd-profile* command as follows to assign an ND profile to be used with the given subscriber or subscriber profile:

```
ipv6 nd-profile name
```

- 10 Use the *dns6* command to specify the primary and secondary DNS IPv6 addresses:

```
dns6 {primary | secondary} ipv6-address
```

- 11 Use the *ipv6 source-validation* command to enable source validation for IPv6.

- 12 Use the *session-action* command to assign the actions taken when a subscriber reaches a timeout or traffic limit.

Table 8 describes the additional subscriber attributes you can configure that are not stack-specific. Configure the attribute commands in subscriber, default subscriber profile, or subscriber profile name configuration mode unless otherwise specified. For more information about these attributes and the configuration of subscriber records and profiles, see *Configuring Subscribers*.

Note: A subscriber record or profile may contain additional attributes that are not applicable to the stack of a subscriber. In such cases, only the applicable attributes are provisioned for the subscriber. For example, a profile applied to an IPv6 subscriber may contain IPv4 attributes that are not provisioned.



Table 8 Additional Subscriber Attributes In a Profile or Subscriber Record

Root Attribute Command	Description
<i>access-line adjust</i>	Uses information received from the DSLAM to adjust the rate.
<i>bulkstats schema</i>	Applies a bulkstats schema to the subscriber profile for this context.
<i>dns</i>	Specifies the primary and secondary DNS server IPv4 addresses This attribute is applicable to IPv4 and dual-stack subscribers only.
<i>flow</i>	Applies a flow policy.
<i>framed-route allow-ecmp</i>	Configures the framed-route attribute for this context.
<i>ip</i>	Applies IP attributes.
<i>nbns</i>	Sets the NBNS server address.
<i>port-limit</i>	Limits the number of sessions a subscriber can access simultaneously.
<i>ppp mtu</i>	Sets the MTU used by PPP for the subscriber circuit.
<i>pppoe client route</i>	Configures the PPPoE client for PPPoE subscribers.
<i>pppoe motm</i>	Creates the message of the minute (MOTM) that the subscriber sees when first logging on.
<i>pppoe url</i>	Sets the subscriber's PPPoE client to point the subscriber's browser to a specific location after the PPP session is established
<i>propagate qos from ip</i>	Modifies the internal classification settings of packets sent or received from the subscriber.
<i>qos node-reference</i>	Sets the QoS node reference.
<i>qos policy queuing</i>	Applies a QoS policy.
<i>rate</i>	Configures inbound and outbound policy circuit rates.
<i>rate-adjust dhcp pwfq</i>	Sets rate adjustment.
<i>session-action</i>	Sets the AAA session action.
<i>session-limit</i>	Sets a limit to the number of sessions allowed for each subscriber line identified by an agent circuit ID or agent remote ID.
<i>shaping-profile</i>	Assigns an ATM shaping profile.
<i>timeout</i>	Sets absolute or idle session timeout value.



Table 8 Additional Subscriber Attributes In a Profile or Subscriber Record

Root Attribute Command	Description
<i>tunnel domain</i>	Enables dynamic assignment of a subscriber PPP session to a L2TP peer that has the same domain alias as the subscriber domain alias.
<i>tunnel name</i>	Statically assigns the subscriber PPP session to a specified L2TP peer or group of L2TP peers.

2.5 Conserving IPv4 Addresses for Dual-Stack Subscribers

In Point-to-Point (PPP) dual-stack subscriber environments that use RADIUS, it is possible to configure the system to conserve IPv4 addresses. With IPv4 address save mode, when a subscriber needs an IPv4 address, the address is requested from the provider's public shared IPv4 address pool, controlled within RADIUS. When the subscriber is not using the address, it is released back to the pool. Providing for the release of unused addresses enables a provider to limit its public, shared IPv4 address pool size or multiplex a greater number of subscribers on a given public, shared IPv4 address pool.

When the subscriber session is initiated, the configuration available at initiation is valid for the duration of the session. If the authentication method changes from RADIUS to local while the session is in progress, IPv4 address re-requests and releases continue to be sent to RADIUS and are rejected. Similarly, if the authentication method changes from local to RADIUS while the session is in progress, IPv4 address re-requests and releases to RADIUS do not occur. The subscriber session must be reinitiated for authentication configuration changes to take effect.

This feature only applies to sessions that obtain their IPv4 addresses from a RADIUS-controlled address pool. A RADIUS server must be available and accessible, and the system must be configured for RADIUS accounting.

To conserve IPv4 addresses in PPP dual-stack subscriber environments using RADIUS, perform the tasks described in Table 9.

Table 9 IPv4 Address Conservation Tasks

Root Command	Task
<i>radius attribute vendor-specific</i>	Enable IPv4 address conservation in RADIUS using the <code>ipv4-address-release-control</code> keyword.
<i>aaa password</i>	Specify the password used to authenticate access re-requests for IPv4 addresses using the <code>ipv4-address-release-control</code> keyword.



Table 9 IPv4 Address Conservation Tasks

Root Command	Task
<i>aaa session rate-limit</i>	Configure for Denial of Service (DoS) protection by limiting Access-Request packets sent to RADIUS.
<i>show subscribers active</i>	Confirm the status of the subscriber session.

2.6 IPv6 Subscriber Services Operations

To manage IPv6 subscriber service functions, perform the appropriate tasks described in Table 9. Enter the **show** commands in any mode.

Table 10 IPv6 Subscriber Services Operations Tasks

Root Command	Task
<i>clear dhcpv6 statistics</i>	Clear DHCPv6 statistics.
<i>debug ipv6 policy</i>	Enable generation of debug messages for an IPv6 policy.
<i>debug ipv6 prefix-library</i>	Enable generation of debug messages for the IPv6 prefix library.
<i>debug ipv6 prefix-list</i>	Enable generation of debug messages for the maintenance of IP Version 6 (IPv6) prefix lists and for the comparison of IPv6 prefix entries to IPv6 prefix lists.
<i>debug ipv6 routing</i>	Enable generation of IP routing debug messages.
<i>show dhcpv6 log</i>	Display the DHCPv6-PD log. You can filter the log history by circuit, server or client DUID, or IPv6 prefix.
<i>show dhcpv6 server duid</i>	Display the DUID that the DHCPv6 server onboard the SmartEdge is using to communicate with its DHCPv6 clients.
<i>show dhcpv6 server host</i>	Display all the active DHCPv6 clients. Display more information with the detail keyword.
<i>show dhcpv6 server host circuit</i>	Display the active DHCPv6 clients on a circuit.
<i>show dhcpv6 server host prefix</i>	Display the active DHCPv6 clients that use a prefix.
<i>show dhcpv6 server host subnet</i>	Display the active DHCPv6 clients on a subnet.



Table 10 IPv6 Subscriber Services Operations Tasks

Root Command	Task
<i>show dhcpv6 statistics</i>	Display DHCPv6 Statistics. Include the <code>detail</code> keyword in the command string to display additional information pertaining to DHCPv6 statistics.
<i>show ipv6 all-host</i>	Display information about all IPv6 hosts stored in the local host table for the current context.
<i>show ipv6 dynamic-host</i>	Display IPv6 dynamic hostname and system ID mapping.
<i>show ipv6 host</i>	Display all static hostname-to-IPv6 address mappings stored in the local host table for the current context.
<i>show ipv6 interface</i>	Display information about IPv6 interfaces, including the interface bound to the Ethernet management port on the controller card.
<i>show ipv6 mroute</i>	Display the IPv6 Protocol Independent Multicast (PIM) routing table.
<i>show ipv6 access-list</i>	Display information about IPv6 subscriber policies configured in the current context.
<i>show ipv6 pool</i>	Display information about the IPv6 shared and DHCPv6 PD prefix pools configured under the current context.
<i>show ipv6 prefix-list</i>	Display information about configured IPv6 prefix lists.
<i>show ipv6 route</i>	Display information about all IPv6 routes.
<i>show nd profile</i>	Display ND profile information for a context.
<i>show nd-circuit</i>	Display ND circuit information for one or more ND circuits.
<i>show nd statistics</i>	Display global statistics for one or more ND router interfaces.
<i>show subscribers active</i>	Display the attributes of active IPv6 subscriber sessions.
<i>show subscribers summary</i>	Display the total number subscribers and their encapsulations in the current context.



Note: If subscribers are unable to obtain IPv6 prefixes even though there should be prefixes available, one or more IPv6 prefixes may not be associated with a valid owner. To verify that all assigned IPv6 prefixes have a valid owner, compare the output from the *show subscribers summary ipv6 all* and *show ipv6 pool summary* commands. When you subtract the number of statically assigned IPv6 prefixes from the total number of IPv6 prefixes in the *show subscribers summary ipv6 all* command, the resulting number should match the total number of prefixes shown in the *show ipv6 pool summary* command output.

You can use the *show subscribers active* command to see whether a prefix is assigned to a subscriber statically or from a DHCPv6 PD pool.



3 Examples

The examples that follow show how to configure a SmartEdge router to provide IPv6 subscriber services to PPP subscribers. For information about how to troubleshoot IPv6 subscriber services, see the *Troubleshooting IPv6 and Dual-Stack Subscriber Services* document.

3.1 End-to-End Solution Configurations

The examples that follow provide end-to-end configuration for a SmartEdge router in a BRAS solution. The examples presented show how to configure a BRAS to use stateful and stateless DHCPv6 to support dual-stack subscribers.

3.1.1 Configuring a BRAS for Dual-stack Subscriber Support Using Stateful DHCPv6

This example results in a configuration where:

- IPv6 prefixes are delegated through stateful DHCPv6.
- IPv4 addresses are assigned to a PPP dual-stack subscriber through IPCP.
- ND provides the IPv6 prefix for the WAN interface (between the BRAS and the CPE).

Figure 1 displays the network topology for this configuration example.




Figure 1 Sample Dual-Stack IPv6 Topology

In this topology:

- 1 A subscribing PC requests an IPv6 prefix from the CPE, which is a router.
- 2 The CPE initiates a PPP connection between the BRAS and the CPE, and LCP comes up.



- 3 The BRAS requests authorization of the subscriber through the RADIUS server.
- 4 On successful authorization, the CPE negotiates IPv6CP and IPCP between the BRAS and the CPE router:
 - IPv6CP exchanges interface IDs for each end of the WAN link.
 - The BRAS sends the IPv4 address to the CPE.
- 5 The BRAS advertises an IPv6 prefix to the CPE in an ND message.
- 6 The BRAS adds a route for the IPv6 prefix in its routing tables.
- 7 The CPE sends a DHCPv6 SOLICIT message to the BRAS to obtain the delegated prefixes and other information.
- 8 The BRAS returns a DHCPv6 ADVERTISE message to the CPE with a delegated IPv6 prefix and DNS information.
- 9 The CPE sends a DHCPv6 REQUEST message to the BRAS, confirming that the CPE accepts the delegated prefix.
- 10 The BRAS sends a DHCPv6 REPLY message to the CPE, confirming that the delegated prefix belongs to the CPE.
- 11 The BRAS adds the IPv6 prefix to the routing table, and the CPE uses the delegated prefix to derive a longer IPv6 prefix for the downstream interfaces.

The example that follows shows the configuration of the SmartEdge router only. For RADIUS and CPE configuration, see the documentation for those products.

Configure two interfaces between the BRAS and the CPE; each interface has its own IPv4 and IPv6 GUA address. One interface is a loopback interface, and the other is a non-loopback interface. A loopback interface is not required on the WAN link; this example shows one possible configuration:

```
[local]BRAS#configure
[local]BRAS (config)#context SJ1
[local]BRAS (config-ctx)#interface test-lb loopback
[local]BRAS (config-if)#ip address 155.13.1.1/24
[local]BRAS (config-if)#ipv6 address 2001:db8:b:4f::1/64
[local]BRAS (config-if)#exit
[local]BRAS (config-ctx)#interface to-cpe
[local]BRAS (config-if)#ip address 155.15.1.1/24
[local]BRAS (config-if)#ipv6 address 2001:db8:b:5f::1/64
```

Configure the DHCPv6 server policy:



```
[local] BRAS (config-ctx) #dhcpv6 server
[local] BRAS (config-dhcpv6-server) #option domain-name-server 2005:db8:b:3f::2
[local] BRAS (config-dhcpv6-server) #option domain-search SJ1.com
[local] BRAS (config-dhcpv6-server) #option preference 5
[local] BRAS (config-dhcpv6-server) #option information-refresh-time 3000000
[local] BRAS (config-dhcpv6-server) #option rapid-commit
[local] BRAS (config-dhcpv6-server) #prefix lifetime preferred 3600 valid 7200
[local] BRAS (config-dhcp-server) #subnet 2001:a:b:3f::/64
[local] BRAS (config-dhcpv6-subnet) #option-domain-name-server 2008:db8:b:3f::1
[local] BRAS (config-dhcpv6-subnet) #option domain-search NY1.com
[local] BRAS (config-dhcpv6-subnet) #prefix lifetime preferred 900 valid 1200
```

Configure a multibind interface to be the DHCPv6 server that uses the DHCPv6 server policy. In this example, the DHCPv6 server is a last-resort interface called `test-last`. Any subscriber circuit that attempts to come up binds to this interface. The `ipv6 unnumbered` command enables IP processing on the `test-1b` interface without assigning it an explicit IP address:

```
[local] BRAS (context) #interface test-last multibind lastresort
[local] BRAS (config-if) #ipv6 unnumbered test-1b
[local] BRAS (config-if) #dhcpv6 server interface
```

Enable AAA to authenticate subscribers through the SmartEdge router local database. Subscribers are authenticated according to parameters set in the subscriber profile for the current context:

```
[local] BRAS (context) #aaa authentication subscriber local
```

Note: To configure subscriber attributes in a subscriber profile, see [Configure the Subscriber Attributes](#). For more information about AAA subscriber authentication, see [Configuring Authentication, Authorization, and Accounting](#).

Create a user record for the subscriber `test`. The configuration specified in this example is applied to subscribers destined for the IP address 155.13.1.10. The `ipv6 framed-prefix` command specifies the IPv6 prefix (2001:db8:b:4f::/64) assigned to the subscriber (using ND or a static assignment). The `ipv6 delegated-prefix` command specifies the IPv6 prefix (2001:db8:1::/48) to be used for DHCPv6 PD. The `nd-profile` command assigns the `abc` profile to the subscriber `test`.

```
[local] BRAS (context) #subscriber name test
[local] BRAS (config-sub) #ip address 155.13.1.10
[local] BRAS (config-sub) #ipv6 framed-prefix 2001:db8:b:4f::/64
[local] BRAS (config-sub) #ipv6 delegated-prefix 2001:db8:1::/48
[local] BRAS (config-sub) #ipv6 nd-profile abc
```

Configure PPPoE encapsulation on an 802.1Q PVC and then bind the PVC using CHAP:

```
[local] BRAS (config) #port ethernet 12/1
[local] BRAS (config-port) #encapsulation dot1q
[local] BRAS (config-port) #dot1q pvc 1 encap pppoe
[local] BRAS (config-dot1q-pvc) #bind authentication chap
```

Create a second PVC with multiprotocol encapsulation (creating a child circuit), and set the protocol of the child circuit to PPPoE. Bind the PVC using CHAP:



```
[local]BRAS (config-port)#dot1q pvc 2 encapsulation multi
[local]BRAS (config-dot1q-pvc)#circuit protocol pppoe
[local]BRAS (config-dot1q-child-PROTO)#bind authentication chap
```

3.1.2 Configuring a BRAS for Dual-stack Subscriber Support Using Stateless DHCPv6

This example results in a configuration where:

- IPv6 DNS information is exchanged through stateless DHCPv6.
- IPv4 addresses are assigned to a PPP dual-stack subscriber through IPCP.
- ND provides the IPv6 prefix for the WAN interface (between the BRAS to the subscriber over the CPE bridge).

Figure 1 displays the network topology for this configuration example.




Figure 2 Sample Dual-Stack IPv6 Topology

In this topology, messages are exchanged between the BRAS and the subscriber through the CPE bridge as follows:

- 1 The subscribing client sends a DHCPv6 informational message request to obtain DNS parameters.
- 2 The BRAS returns a DHCPv6 Reply message to the subscribing client with the requested DNS information (all DNS options configured under the DHCPv6 server profile).

The example that follows shows the configuration of the SmartEdge router only. For RADIUS and CPE configuration, see the documentation for those products.

Configure an interface between the BRAS and the CPE; the interface has its own IPv4 and IPv6 GUA address:

```
[local]BRAS#configure
[local]BRAS (config)#context SJ1
[local]BRAS (config-ctx)#interface to-cpe
[local]BRAS (config-if)#ip address 155.15.1.1/24
[local]BRAS (config-if)#ipv6 address 2001:db8:b:5f::1/64
```

Configure the DHCPv6 server policy:

```
[local]BRAS (config-ctx)#dhcpv6 server
[local]BRAS (config-dhcpv6-server)#option domain-name-server 2005:db8:b:3f::2
[local]BRAS (config-dhcpv6-server)#option domain-search SJ1.com
[local]BRAS (config-dhcpv6-server)#option information-refresh-time 700
```



Configure a multibind interface to be the DHCPv6 server. In this example, the DHCPv6 server is a last-resort interface `test-last`. Any DHCPv6 subscriber circuit that attempts to come up binds to this interface. The `ipv6 unnumbered` command configures the `test-last` interface to use the IPv6 address from the `to-cpe` interface:

```
[local]BRAS(context)#interface test-last multibind lastresort
[local]BRAS(config-if)#ipv6 unnumbered to-cpe
[local]BRAS(config-if)#dhcpv6 server interface
```

Enable AAA to authenticate subscribers through the SmartEdge router local database. Subscribers are authenticated according to parameters set in the subscriber profile for the current context:

```
[local]BRAS(context)#aaa authentication subscriber local
```

Note: To configure subscriber attributes in a subscriber profile, see [Configuring the Subscriber Attributes](#). For more information about AAA subscriber authentication, see [Configuring Authentication, Authorization, and Accounting](#).

Create a user record for the subscriber `test`. The configuration specified in this example is applied to subscribers destined for the IP address 155.13.1.10. The `ipv6 framed-prefix` command specifies the IPv6 prefix (2001:db8:b:4f::/64) assigned to the subscriber (using ND or a static assignment). The `nd-profile` command assigns the `abc` profile to the subscriber `test`:

```
[local]BRAS(context)#subscriber name test
[local]BRAS(config-sub)#ip address 155.13.1.10
[local]BRAS(config-sub)#ipv6 framed-prefix 2001:db8:b:4f::/64
[local]BRAS(config-sub)#ipv6 nd-profile abc
```

Configure PPPoE encapsulation on an 802.1Q PVC and then bind the PVC using CHAP:

```
[local]BRAS(config)#port ethernet 12/1
[local]BRAS(config)#encapsulation dot1q
[local]BRAS(config-port)#dot1q pvc 1 encap pppoe
[local]BRAS(config-dot1q-pvc)#bind authentication chap
```

Create a second PVC with multiprotocol encapsulation (creating a child circuit), and set the protocol of the child circuit to PPPoE. Bind the PVC using CHAP:

```
[local]BRAS(config-port)#dot1q pvc 2 encapsulation multi
[local]BRAS(config-dot1q-pvc)#circuit protocol pppoe
[local]BRAS(config-dot1q-child-PROTO)#bind authentication chap
```

3.2 Configuring Dual-Stack Subscriber Support on a SmartEdge LNS

This configuration example illustrates how to configure the SmartEdge router as an L2TP Network Server (LNS) supporting dual-stack subscribers. In



addition, priority weighted-fair queuing (PWFQ) is applied to the subscribers. For detailed information on how to configure L2TP, including how to configure the SmartEdge router as an L2TP access concentrator (LAC), see *Configuring L2TP*.

The following example shows how to configure the IPv6 license:

```
[local] Redback (config) #software license
[local] Redback (config-license) #subscriber active ipv6 64000 encrypted
```

The following example show how to configure a domain alias for the context for use with a LAC peer:

```
[local] Redback (config) #context lns
[local] Redback (config-ctx) #domain isp
```

The following example show how to configure the interface for core traffic:

```
[local] Redback (config-ctx) #interface core
[local] Redback (config-if) #ip address 210.1.1.1/24
[local] Redback (config-if) #ipv6 address 1000:0:0:1::1/64
[local] Redback (config-if) #exit
```

The following example show how to configure the LNS interface:

```
[local] Redback (config-ctx) #interface lns-eth1
[local] Redback (config-if) #ip address 200.1.1.2/24
```

The following example show how to configure the multibind interface pool for IPv4 and IPv6 subscribers:

```
[local] Redback (config-ctx) #interface pool multibind
[local] Redback (config-if) #ip address 10.1.0.1/16
[local] Redback (config-if) #ipv6 address 3000:0:1::/48
[local] Redback (config-if) #ip pool 10.1.0.0/16
[local] Redback (config-if) #ipv6 pool 3000:0:1:1::/64 3000:0:1:8::/64
```

The following example show how to configure OSPF routing for the LNS interface:

```
[local] Redback (config-ctx) #router ospf1
[local] Redback (config-ospf) #fast-convergence
[local] Redback (config-ospf) #area 0.0.0.0
[local] Redback (config-ospf-area) #interface lns-eth1
[local] Redback (config-ospf-if) #passive
```



The following example show how to configure the Neighbor Discovery (ND) profile for IPv6 subscribers:

```
[local]Redback(config-ctx)#nd profile test
[local]Redback(config-nd-profile)#ra-interval 5
[local]Redback(config-nd-profile)#exit
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#ipv6 nd-profile test
```

The following example show how to configure subscribers and apply a PWFQ policy to them:



```
[local] Redback (config-ctx) #subscriber name joe1
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe2
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe3
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe4
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe5
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe6
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe7
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
[local] Redback (config-sub) #exit
[local] Redback (config-ctx) #subscriber name joe8
[local] Redback (config-sub) #password test
[local] Redback (config-sub) #ip address pool
[local] Redback (config-sub) #qos policy queuing my_pwfq
[local] Redback (config-sub) #ipv6 framed-pool pool
```

The following example show how to configure the LNS as a peer to the LAC:



```
[local] Redback (config-ctx) #l2tp-peer name lac media
udp-ip remote ip 200.1.1.1 local 200.1.1.2
[local] Redback (config-l2tp) #tunnel-auth key rbak
[local] Redback (config-l2tp) #function lns-only
[local] Redback (config-l2tp) #local-name lns
```

The following example show how to configure QoS queue mapping:

```
[local] Redback (config-ctx) #qos queue-map default
[local] Redback (config-queue-map) #num-queues 2
[local] Redback (config-num-queues) #queue 0 priority 0
[local] Redback (config-num-queues) #queue 1 priority 1 2 3 4 5 6 7
[local] Redback (config-num-queues) #exit
[local] Redback (config-queue-map) #num-queues 4
[local] Redback (config-num-queues) #queue 0 priority 0
[local] Redback (config-num-queues) #queue 1 priority 1 2
[local] Redback (config-num-queues) #queue 2 priority 3 4 5 6
[local] Redback (config-num-queues) #queue 3 priority 7
[local] Redback (config-num-queues) #exit
[local] Redback (config-queue-map) #num-queues 8
[local] Redback (config-num-queues) #queue 0 priority 0
[local] Redback (config-num-queues) #queue 1 priority 1
[local] Redback (config-num-queues) #queue 2 priority 2
[local] Redback (config-num-queues) #queue 3 priority 3
[local] Redback (config-num-queues) #queue 4 priority 4
[local] Redback (config-num-queues) #queue 5 priority 5
[local] Redback (config-num-queues) #queue 6 priority 6
[local] Redback (config-num-queues) #queue 7 priority 7
```

The following example show how to configure the PWFQ policy:

```
[local] Redback (config) #qos policy my_pwfq pwfq
[local] Redback (config-policy-pwfq) #queue 0 priority 0 weight 100
[local] Redback (config-policy-pwfq) #queue 1 priority 1 weight 100
[local] Redback (config-policy-pwfq) #queue 2 priority 2 weight 100
[local] Redback (config-policy-pwfq) #queue 3 priority 3 weight 100
[local] Redback (config-policy-pwfq) #queue 4 priority 4 weight 100
[local] Redback (config-policy-pwfq) #queue 5 priority 5 weight 100
[local] Redback (config-policy-pwfq) #queue 6 priority 6 weight 100
[local] Redback (config-policy-pwfq) #queue 7 priority 7 weight 100
```

The following example show how to configure the Ethernet cards, ports, interfaces, and bindings:



```
[local] Redback (config) #card ge-10-port 6
[local] Redback (config-card) #exit
[local] Redback (config) #port ethernet 6/6
[local] Redback (config-port) #no shutdown
[local] Redback (config-port) #bind interface core lns
[local] Redback (config-port) #exit
[local] Redback (config) #card ge-10-port 9
[local] Redback (config-card) #exit
[local] Redback (config) #port ethernet 9/1
[local] Redback (config-port) #no shutdown
[local] Redback (config-port) #bind interface lns-eth1 lns
```

3.3 Detailed Configuration Examples for Individual Elements of an IPv6 Solution

The sections that follow provide detailed, extended configuration examples for the individual elements of a BRAS IPv6 solution.

3.3.1 Configuring NAS IPv6 Address

The following example shows how to configure the NAS-IPv6-Address RADIUS attribute to match the primary IPv6 address of interface `if1`. The IPv6 address for the interface must already be configured. After this is configured, the NAS-IPv6-Address attribute value is included in RADIUS Access-Request and Accounting-Request packets sent by the SmartEdge router.

```
[local] BRAS #configure
[local] BRAS (config) #context SJ1
[local] BRAS (config-ctx) #radius attribute NAS-IPv6-Address interface if1
```

3.3.2 Configuring a Subscriber Profile

The following example shows how to create a subscriber profile `sj-sub-10`:

```
local] BRAS (config-ctx) #subscriber profile sj-sub-10
[local] BRAS (config-sub) #ipv6 delegated-prefix 2001:a:b:4f::1/128
[local] BRAS (config-sub) #ipv6 framed-prefix 2002:a:b:5f::1/128
[local] BRAS (config-sub) #ipv6 nd-profile abc
```

3.3.3 Configuring a Subscriber Record

The following example shows how to configure subscriber record `test`:



```
[local] BRAS(config-ctx)#subscriber name test
[local] BRAS(config-sub)#ipv6 delegated-prefix 2001:db8:b:4f::1/48
[local] BRAS(config-sub)#ipv6 framed-prefix 2002:a:b:5f::1/48
[local] BRAS(config-sub)#ipv6 nd-profile abc
[local] BRAS(config-sub)#ipv6 framed-route 2010:db8:b:5f::1/48 2002:db8:b:5f::1 1000
[local] BRAS(config-sub)#ipv6 source-validation
[local] BRAS(config-sub)#profile sj-sub-10
```

3.3.3.1 Configuring a DHCPv6 Profile

The following example shows how to configure the DHCPv6 server policy. In this example, the network administrator:

- Specifies a domain name server on the IPv6 address `2005:db8:b:3f::2`.
- Specifies the domain `sj.com` for DNS resolution.
- Configures the preference value for this DHCPv6 server to 5.
- Configures clients to wait 3000000 seconds before refreshing the configuration information received from DHCPv6 server.
- Enables RAPID COMMIT (only two messages are exchanged between the BRAS and the CPE).
- Specifies a preferred lifetime of 3600 seconds.
- Specifies a valid lifetime of 7200 seconds.
- Specifies a subset of DHCPv6 attributes for clients in the subnet `2001:db8:b:3f::/68`. For clients in this subnet, the following attributes are different from all other subscribers:
 - The preferred lifetime is 2000 seconds (instead of 3600 seconds).
 - The valid lifetime is 4000 seconds (instead of 7200 seconds).
- Specifies a subset of DHCPv6 attributes for clients in the subnet `2001:db8:2:2::/68`. For clients in this subnet, the following attributes are different from all other subscribers:
 - The domain name server is located at IPv6 address `2008:db8:4000:1::2` (all other subscribers use the domain name server located on `2005:db8:b:3f::2`).
 - Subscribers in this subnet use the domain `subnet.corp.com` for DNS resolution (all other subscribers use the domain `sj1.com`).
 - The preferred and valid lifetimes are set to be `infinite` (instead of 3600 seconds).



```
[local] BRAS (config-ctx) #dhcpv6 server
[local] BRAS (config-dhcpv6-server) #option domain-name-server 2005:db8:b:3f::
[local] BRAS (config-dhcpv6-server) #option domain-search SJ1.com
[local] BRAS (config-dhcpv6-server) #option preference 5
[local] BRAS (config-dhcpv6-server) #option information-refresh-time 3000000
[local] BRAS (config-dhcpv6-server) #option rapid-commit
[local] BRAS (config-dhcpv6-server) #prefix lifetime preferred 3600 valid 7200
[local] BRAS (config-dhcpv6-server) #subnet 2001:db8:b:3f::/68
[local] BRAS (config-dhcpv6-server) #prefix lifetime preferred 2000 valid 4000
[local] BRAS (config-dhcpv6-server) #subnet 2001:db8:2:2::/68
[local] BRAS (config-dhcpv6-subnet) #option domain-name-server 2008:db8:4000:1::2
[local] BRAS (config-dhcpv6-subnet) #option domain-search subnet.corp.com
[local] BRAS (config-dhcpv6-subnet) #prefix lifetime infinite
```

3.3.4 Configuring Shared IP Pools for ND

The example that follows shows how to create and configure two shared IP pools for ND to use for allocating IPv4 addresses and IPv6 prefixes.

First, create two shared IP pools under the multibind interface `ip_pools`:

- A shared IP pool of IPv4 addresses; this pool is identified by the IPv4 address `155.13.1.1/24`.
- A shared IPv6 prefix pool that contains IPv6 prefixes in the range from `2001:db8:b:1::/64` to `2001:db8:b:100::/64`. This pool has a falling threshold of 50%, and crossing events are recorded as traps.

```
[local] BRAS #configure
[local] BRAS (config) #context SJ1
[local] BRAS (config-ctx) #interface ip_pools multibind
[local] BRAS (config-if) #ip address 155.13.1.1/24
[local] BRAS (config-if) #ip pool 155.13.0.0/24
[local] BRAS (config-if) #ipv6 address 2001:db8:b::/48
[local] BRAS (config-if) #ipv6 pool 2001:db8:b:1::/64 2001:db8:b:100::/64 threshold percentage falling 50 trap
[local] BRAS (config-if) #exit
```

Then, specify that a subscriber (called `sub_1`):

- Obtains IPv4 addresses from the shared IP pool configured within the same context (`SJ1`).
- Obtains IPv6 prefixes from the shared IPv6 prefix pool configured within the same context (`SJ1`).
- Inherits the ND configuration parameters specified by the ND profile `abc`.

```
[local] BRAS (config-ctx) #subscriber name sub_1
[local] BRAS (config-sub) #ip address pool
[local] BRAS (config-sub) #ipv6 nd-profile abc
[local] BRAS (config-sub) #ipv6 framed-pool
```

3.3.5 Configuring a DHCPv6 PD Pool

The following example shows how to create and configure a DHCPv6 PD pool, and then configure a subscriber to obtain IPv6 prefixes from that pool. In this



example, the DHCPv6 PD pool inherits falling threshold values specified for all DHCPv6 PD pools configured within a context.

First, specify falling threshold values applicable to all DHCPv6 pools configured under the context `SJ1`:

```
[local]BRAS#configure
[local]BRAS(config)#context SJ1
[local]BRAS(config-ctx)#ipv6 pool dhcpv6 threshold percentage falling 20 log 10 trap
[local]BRAS(config-ctx)#exit
```

Configure a DHCPv6 PD pool under a multibind interface `test-2`. This pool contains IPv6 prefixes in the range from `ipv6 pool dhcpv6 2001:db8:1:100::/56` to `2001:db8:1:ff00::/56`:

```
[local]BRAS(config-ctx)#interface test-2 multibind
[local]BRAS(config-if)#ipv6 address 2001:db8:b::/48
[local]BRAS(config-if)#ipv6 pool dhcpv6 2001:db8:1:100::/56 2001:db8:1:ff00::/56
```

Configure the following attributes in a subscriber profile for the subscriber `sub_2`:

- The Delegated-Max-Prefix attribute (the maximum number of IPv6 prefixes that can be delegated to a subscriber) is 5.
- The subscriber obtains IPv6 prefixes from the shared IPv6 prefix pool configured within the same context (`SJ1`).
- The subscriber inherits the ND configuration parameters specified by the ND profile `abc`.

```
[local]BRAS(config-ctx)#subscriber sub_2
[local]BRAS(config-if)#ipv6 delegated-prefix maximum 5
[local]BRAS(config-if)#ipv6 framed-pool
[local]BRAS(config-if)#ipv6 nd-profile abc
```

3.3.6 Configuring Statically Mapped DHCPv6 Prefixes

The following example shows how to configure static mapping for IPv6 two prefixes. In this example:

- The IPv6 prefix `3001:db8:c/48` can be assigned to subscribers with a DUID of `00:01:00:01:00:04:93:e0:00:00:00:00:a2:a2`.
- The IPv6 prefix `3001:db8:c/48` can be assigned to subscribers with a DUID of `00:01:00:01:00:04:93:e0:00:00:00:00:a2:a2` and an IAID of `0xfedcba98`.

```
[local]BRAS(config-ctx)#dhcpv6 server
[local]BRAS(config-dhcpv6-server)#prefix 3001:db8:c/48 duid 00:01:00:01:00:04:93:e0:00:00:00:00:a2:a2
[local]BRAS(config-dhcpv6-server)#prefix 3001:db8:c/48 duid 00:01:00:01:00:04:93:e0:00:00:00:00:a2:a2 \
iaid 0xfedcba98
```