Troubleshooting IPv6 and Dual-Stack Subscriber Services

SmartEdge Router

FAULT TRACING DIRECT.

11/154 51-CRA 119 1170/1-V2 Uen E

Copyright

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.

Contents

1 Is This an IPv6 Issue?

The first step in troubleshooting subscriber problems is to perform a general triage of the problem.

Use the *BRAS Troubleshooting Guide* for information about the following issues and factors that could cause subscriber problems:

- System status
- Subscriber licenses
- Subscriber connection
- Configuration
- Interface connections
- Bindings
- Authentication
- Traffic flow
- Subscriber routes

Depending on the evidence, continue troubleshooting in the problem areas described in Table 1.

For more information about the commands used in this guide, see *Command List*.

The following focused troubleshooting procedures assume that you have collected data about problematic subscriber sessions by using the **`show subscriber active`** ***`sub-name`***, **`show subscriber active all`**, or **`show ppp all`** commands. Note the subscriber name and domain, circuit handle, context, IP pool, prefix, IP address, and any other relevant information.

Table 1 IPv6 and Dual-Stack Subscriber Issues

Dual-stack PPP sessions do not come up. Subscribers are not enabled for IPv6.	See Section 2 on page 2 for PPP issues.
Subscriber has the wrong type of IP address. Dual-stack subscribers cannot log on, are stuck, or go down. After a reload or switchover, dual-stack subscriber sessions do not come up.	See Section 3 on page 6 for AAA issues

Table 1 IPv6 and Dual-Stack Subscriber Issues

Subscriber does not receive a prefix from the IPv6 pool	See Section 4 on page 9 for IPv6 prefix pool issues.
An Neighbor Discovery (ND) neighbor is missing from the neighbor cache, or a ping to the neighbor fails. IPv6 subscribers addresses are not autoconfiguring correctly.	See Section 5 on page 12 for ND issues.
The DHCPv6 client does not receive a prefix. Addresses are not being assigned by the IP (PD) address pool. Traffic is not flowing towards a subscriber. After a reload or switchover, prefixes are not restored.	See Section 6 on page 16 for DHCPv6 issues.
IPv6 subscriber traffic is not reaching a destination, or IPv6 routes are missing from the RIB. A ping does not reach an IPv6 destination.	See Section 7 on page 20 for IPv6 routing issues.

2 Troubleshooting PPP Problems

For the procedures to check general PPP problems on the SmartEdge® router, see *BRAS Troubleshooting Guide*.

To troubleshoot PPP problems with IPv6 and dual-stack subscribers, perform the steps in Table 2.

Table 2 Troubleshooting IPv6 and Dual-stack PPP Subscribers

Step	**Command**	**Check?**
PPP Sessions Are Not Coming Up or Are Dropping.		

Check Subscriber Licenses.	`show subscriber license summary`
	`debug aaa all`	

Table 2 Troubleshooting IPv6 and Dual-stack PPP Subscribers

Step	Command	Check?

Investigate Subscriber Session Status (single subscribers).	**`show subscriber active`** *`sub-name`*
	`show subscriber active` *`cct-handle`*
	`show subscriber log session` *`cct-handle`*
	`show subscriber all | grep` *`sub-name`*	

Investigate Subscriber Session Status (multiple subscribers).	**`show subscriber summary`** [**`ipv4`** | **`ipv6`**] [**`all`**]	
Subscribers Not Enabled For IPv6.		
Check subscriber IPv4 and IPv6 status.	**`show ppp all`**	

Verify that dual stack or IPv6 is enabled in the AAA response by running the following debug commands for a short time.	**`debug circuit`** *`cct-handle`*
	`debug circuit ppp authentication`
	`debug circuit ppp packet`	

Determine which circuit is not working.	**`debug ppp exception`**
	`show circuit` *`cct-handle`* **`detail`**	

If IPCP or IPv6CP was not enabled (above in **`show ppp all`** output), check the configuration. Verify configured RADIUS attributes in RADIUS users file or local configuration (Framed-IPv6-Prefix and Delegated-IPv6-Prefix are required).	**`show configuration`** Check RADIUS users file.	
Check the subscriber information for configured RADIUS attributes. Check a summary of IPv4 or IPv6 information for all subscribers.	**`show subscriber active sub-name@domain`** **`show subscriber summary`** [**`ipv4`** | **`ipv6`**] [**`all`**]	
Verify subscriber interface information, such as the IPv6 link-local address or IPv6 address.	**`show ipv6 interface`**	

2.1 PPP Sessions Are Not Coming Up or Are Dropping

To investigate perform the following tasks.

2.1.1 Check Subscriber Licenses

To start investigating dual-stack subscribers that are not coming up, check the subscriber licenses.

1. To check the status of SmartEdge licenses, enter the **`show subscriber license summary`** command.

2. To examine IPv6 subscriber license events, enable AAA debugging with the **`debug aaa all`** command.

2.1.2 Investigate Subscriber Status

1. To investigate single subscriber sessions, use the **`show subscriber active`** command as shown in Step 7 in Section 3.2 on page 8. You can also use the **`show subscriber log session`** ***`cct-handle`*** command, or the **`show subscriber all | grep`** ***`sub-name`*** command.

2. To investigate multiple subscribers, use the **`show subscriber summary all`** command.

2.2 Are Subscribers Enabled for IPv6?

Verify and troubleshoot the IPv6 stack with the following steps:

1. Verify that IPv4 and IPv6 are enabled and up:

```
[local]Redback#show ppp all
Fri Mar 19 11:01:34 2010
                             LCP     IPCP   IPV6CP  NLCP MPLSCP
Port/Circuit          Unit  State    State   State  State    State
----------------------------------------------------------------
2/8 vlan-id 32 pppoe 6   0 Opened   Opened   Opened
Total circuits: 1  up: 1  down: 0
```

The expected session should be up and LCP, IPCP, and IPv6CP state should be `Opened`.

2. To verify that IPv6CP is enabled, you can use the **`show ppp circuit`** command to ensure that IPCP came up (yes) and IPv4 is enabled (yes).

   Also displays the interface-ID from AAA.

3. Verify that dual stack is enabled by running the following debug commands for a short time:

   `debug circuit ppp authentication`

   `debug circuit ppp packet`

For example, in the following segment of the output, look for IPv4 state of `Enable` and IPv6 state of `Enable` or `Disable`. `LCP ConfRej` displays if IPv4 or IPv6 are disabled. If authentication is failing, you see repeated AAA messages.

```
[local]Redback#debug circuit ppp packet
Mar 9 15:08:57.000: [2/8:1023:63/6/2/3813]: %PPP-7-AUTH: [0]
Sending authentication request to AAAd
Mar 9 15:08:57.011: [2/8:1023:63/6/2/3813]: %PPP-7-AUTH: [0] Authentication response status: Success
Mar 9 15:08:57.011: [2/8:1023:63/6/2/3813]: %PPP-7-AUTH: [0] Authentication response: IPv4 Enable
Mar 9 15:08:57.011: [2/8:1023:63/6/2/3813]: %PPP-7-AUTH: [0] Authentication response: IPv6 Enable
Mar 9 15:08:57.011: [2/8:1023:63/6/2/3813]: %PPP-7-AUTH: [0] IPv4 address 192.168.14.254 is from pool
Mar 9 15:18:24.641: [2/8:1023:63/6/2/3814]: %PPP-7-AUTH: [0] Authentication response status: Success
Mar 9 15:18:24.642: [2/8:1023:63/6/2/3814]: %PPP-7-AUTH: [0] Authentication response: IPv4 Enable
Mar 9 15:18:24.643: [2/8:1023:63/6/2/3814]: %PPP-7-AUTH: [0] Authentication response: IPv6 Disable
```

4. Determine which circuit is having problems by running the **`debug ppp exception`** command, which provides data about subscriber circuits with exceptions.

   Copy the circuit handle and use it in the **`show circuit`** ***`cct-handle`*** **`detail`** command. Look for reports that IPCP and IPv6CP came up.

If IPCP or IPv6CP was not enabled, check the configuration of the subscribers in the RADIUS users file or the router configuration.

5. To verify required RADIUS attributes (Framed-IPv6-Prefix and Delegated-IPv6-Prefix), check the RADIUS users file. You can also verify the IPv6-DNS value and IPv6 options in the users file, as in the following RADIUS example:

```
user1 User-Password := "test"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Framed-IPv6-Prefix := "6002:2:2:33::3/64",
Framed-IPv6-Route = "3000:1:1::1 5001:1:2::3 7",
Framed-IP-Address = 10.1.0.3,
Framed-IP-Netmask = 255.255.255.255,
Delegated-IPv6-Prefix = "6002:2:3:33::4/56”,
IPv6-DNS = "1=2000::106:a00:20ff:fe99:a998,2=2000::106:a00:20ff:fe99:a995",
Framed-Interface-Id="200:70ff:fe02:102",
IPv6-Option  ="source-validation=1",
IPv6-Option +="route-tag=22",
IPv6-Option   += “nd-profile=nd1”
```

6. If the attributes do not appear in RADIUS, they may assigned through the local router configuration. Check the subscriber configuration (IPv6 address under the PPP interface or the IPv6 framed prefix under the subscriber name) in the **`show configuration`** command output:

```
context local
 interface LOOPBACK0 loopback
  ip address 192.168.1.1/16
  ipv6 address 2001:a:b::1/8
!
 interface SUBSCRIBER multibind
  ip unnumbered LOOPBACK0
  ipv6 unnumbered LOOPBACK0
  dhcpv6 server interface
  ip clear-df
  ip address 192.168.1.1/16
     ipv6 address 12ab:1:1:1::1/48
     ip pool 192.168.0.0/16
!
 aaa authentication subscriber local
 aaa accounting subscriber radius
 radius accounting server 10.1.1.1 key redback
!
 radius attribute nas-ip-address interface LOOPBACK0
 radius attribute nas-ipv6-address interface to_traffic
!
 subscriber name test
  password test
  ip address pool
  ipv6 framed-prefix 12ab:1:1:10::/64
  ipv6 delegated-prefix  4ffe:1:1::/64
  ipv6 nd-profile nd1
  ipv6 source-validation
```

In this output, you can also determine the IPv6 address assigned to the subscriber interface.

7. You can verify the subscriber status using the **show subscribers active** ***sub-name@domain*** command:

```
[local]Redback#show subscriber active user1@local
        Session state Up
        Circuit   11/10 vlan-id 101:1 pppoe 1
        Internal Circuit   11/10:1023:63/1/2/120
        Interface bound  SUBSCRIBER
        Current port-limit unlimited
        Protocol Stack Dual
   dns primary 1.1.1.1 (applied from sub_default)
        dns secondary 1.1.1.2 (applied from sub_default)
        idle timeout direction in (applied from sub_default)
        timeout absolute 86400 (applied)
        ascend data filter  (applied in)
           ip in drop dstip 217.89.29.250/32
           ip in forward dstip 224.0.0.22/32 srcip 80.156.111.77/32
           ip in drop dstip 224.0.0.0/4
           ip in forward dstip 58.0.0.0/8 srcip 115.0.0.0/8
           ip in forward dstip 57.0.0.0/8 srcip 115.0.0.0/8
           ip in forward dstip 115.0.0.0/8 srcip 58.0.0.0/8
           ip in forward dstip 115.0.0.0/8 srcip 57.0.0.0/8
           ip in drop
        Framed-IPV6-Prefix  3001:1:2:9101::/64 (applied)
        Ipv6-ND-Profile p1 (applied)
        Delegated-IPV6-Prefix 6001:1:2:9101::/64 (applied)
        Ipv6-DNS  primary 1:1::1  (applied from sub_default)
        Ipv6-DNS secondary 1:1:: (applied from sub_default)
```

In this case, the subscriber has the required Framed-IPv6-Prefix and Delegated-IPv6-Prefix attributes assigned.

8. For an overview of subscriber status, use the **show subscriber summary** [**ipv4** | **ipv6**] [**all**] command.

3 Troubleshooting Subscriber Access (AAA)

For basic AAA troubleshooting steps, see *BRAS Troubleshooting Guide*.

To troubleshoot subscriber access problems in a dual-stack or IPv6 network, perform the steps in Table 3, which assume that you have turned on logging.

Table 3 Troubleshooting Subscriber Access (AAA)

Step	Command	Check?

Turn on AAA debugging.	`debug aaa all`
	`debug aaa acct`
	`debug aaa author`
	`debug aaa authen`	

Subscriber Has The Wrong Type of IP Address		
Examine recent AAA events for a subscriber.	**show subscriber active** ***sub-name@domain*** **show subscriber log session** ***session-id*** **show subscriber log handle** ***cct-handle***	

If a subscriber has the wrong IP address type, examine the access events in the AAA debug output logs.	`logging console`
	Or
	`terminal monitor`	

After Reload Or Switch-Over, Dual Stack Sessions Do Not Come Up		
Check the number of PPP and PPPoE subscribers under Active.	`show subscriber summary all`	
For a single subscriber, also view AAA events.	**show subscriber log** ***cct-handle***	
If no other remedy is found, you can clear a subscriber and bring it back up.	**clear subscriber username** ***sub-name***	
As a last resort, you can restart the AAA process and collect troubleshooting data for Customer Support.	`process coredump aaad`	

3.1 Enable AAA Debugging and Configure AAA Logging

1. If you have not already done so, turn on debugging with the **`debug aaa all`** command. After the problem is resolved, use the **`no`** form of the command to turn it off.

   If authentication fails for the session, you may also find the **`debug aaa acct`**, **`debug aaa author`**, or **`debug aaa authen`** commands useful.

3.2 Subscriber Has the Wrong Type of IP Address

Troubleshoot the IP address type with the following steps:

If a subscriber has the wrong IP address type:

1. Examine the access events in the AAA debug output logs for errors in prefix assignment.

2. Examine recent AAA events for the specific subscriber with the **`show subscriber log session`** ***`cct-handle`*** command:

```
[local]Redback#show subscriber log session 4/3:1023:63/6/2/11
---------------------------------------------------------
8       IN      Sat Mar 27 18:42:04.694112
        IPC_ENDPOINT = PPPd, MSG_TYPE = SESSION_UP,
---------------------------------------------------------
        IPC_ENDPOINT = ISM-IF, MSG_TYPE = IF-BIND,
        Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = ISM-CCT, MSG_TYPE = CCT-GEN-CFG,
        Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = PPPd, MSG_TYPE = PROTO_IPV6_UP,
        Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = ISM-CCT, MSG_TYPE = CCT-GEN-CFG,
        Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = PPPd, MSG_TYPE = SESSION_DOWN, term_ec = 140
           Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = ISM-IF, MSG_TYPE = IF-UNBIND,
        Username = user2668@local,
---------------------------------------------------------
        IPC_ENDPOINT = ISM-CCT, MSG_TYPE = CCT-GEN-CFG,
        Username = user2668@local,
        CCT_HANDLE = Unknown circuit
---------------------------------------------------------
        IPC_ENDPOINT = ISM-CCT, MSG_TYPE = CCT-SUB-SESS-DOWN-CPLT,
---------------------------------------------------------
```

Look for PROTO_IPV4_UP or PROTO_IPV6_UP in the output.

3.3 After Reload or Switchover, IPv6 Subscriber Sessions Do Not Come Up

This section assumes that you already ruled out hardware, network, configuration, and RADIUS problems using *BRAS Troubleshooting Guide*.

1. Examine the subscriber numbers with the **show subscriber summary all** command and look for clues in the output of the **show subscriber log session** ***cct-handle*** command.

2. If no other remedy is found, you can restart a module or clear a subscriber and bring it back up:

```
[local]Redback#clear subscriber username sub-name
```

3. As a last resort, restart the AAA process and collect troubleshooting data for Customer Support with the **process coredump aaad** command.

   For more information, see *Data Collection Guideline for the SmartEdge Router*.

4 Troubleshoot IPv6 Address Pool Problems

Table 4 IPv6 Pool Problems

Step	Commands	Check?
Subscriber Does Not Receive Prefix From IPv6 Pool		
Investigate the status of IPv6 address pools.	**show ipv6 pool** [**dhcpv6**] [***pool-name***] **show ipv6 pool** [**dhcpv6**] [***pool-name***] [**thresholds**] **show ipv6 pool** [**dhcpv6**] **summary**	
Check the DHCPv6 configuration.	**show configuration dhcpv6**	
Enable IPv6 prefix list and prefix library debugging messages.	**debug ipv6 prefix-list** **debug ipv6 prefix-library**	
You can also display information about the configured prefix lists.	**show ipv6 prefix-list**	

Table 4   IPv6 Pool Problems

Step	Commands	Check?
Check the status of current DHCPv6 hosts, including delegated prefix information, and whether the prefixes are from static configuration or from pools. Also, check DHCPv6 statistics.	**show dhcpv6 server host** **show dhcpv6 statistics detail**	
Check subscriber use of prefixes.	**show subscriber summary ipv6 all** **show ipv6 pool summary** **show subscriber active** command with the ***username@domain*** or ***cct-handle*** keywords	

4.1 Subscriber Does Not Receive Prefix From IPv6 Pool

If subscribers are unable to get IPv6 addresses from a pool, you may see the log message, "Error getting ip address from pool for subscriber".

To investigate the problem, perform the following steps:

1. Investigate the status of prefixes in IPv6 address pools by using one of the **show ipv6 pool** commands in the following examples:

```
[CORP]Redback#show ipv6 pool dhcpv6 123-pool

Interface "123-pool":
   2001:a:c:1::/64           2001:a:c:4::/64         0 in-use,
    4 free,   0 reserved
   4001:a:c::/64             4001:a:c:ff::/64        1 in-use,
  255 free,   0 reserved

[CORP]Redback#show ipv6 pool 123-pool

Interface "123-pool":
   2001:a:b:1::/64           2001:a:b:4::/64                          1 in-use,
    3 free,   0 reserved

[CORP]Redback#show ipv6 pool dhcpv6 123-pool thresholds

Interface "123-pool":
   2001:a:c:1::/64           2001:a:c:4::/64        threshold
percentage falling 50 log 25 log
   4001:a:c::/64             4001:a:c:ff::/64       threshold
percentage falling 50 log 25 log


[CORP]Redback#show ipv6 pool dhcpv6 thresholds

Interface "123-pool":
   2001:a:c:1::/64           2001:a:c:4::/64        threshold
percentage falling 50 log 25 log
   4001:a:c::/64             4001:a:c:ff::/64       threshold
percentage falling 50 log 25 log
```

To debug shared prefix pools (also known as AAA or ND prefix pools), use the command without the **dhcpv6** keyword.

The **dhcpv6** keyword indicates that you are requesting information about the **dhcpv6** PD prefix pools in this context, and omitting the keyword means that you are requesting information about the shared prefix pools in this context.

You can also see a summary of IPv6 pool information with the **show ipv6 pool dhcpv6 summary** command.

2. Examine the configuration of the DHCPv6 server host using the **show configuration dhcpv6** command; see the output segment in the following segment of the output:

```
interface 123-pool multibind
  ip address 10.1.1.1/16
  ipv6 address 2001:a:1:2::1/64
  dhcpv6 server interface
  ip pool 10.1.0.0/16
  ipv6 pool dhcpv6 8001:a:b:1::/64 8001:a:b:100::/64 name pd-pool
  threshold percentage falling 25 log 10 trap
```

3. To examine the status of current DHCPv6 hosts, use the **show dhcpv6 server host** command.

4. To see more prefix-related information, use the **show dhcpv6 statistics detail** command.

5. You can also enable IPv6 prefix list debugging messages with the **debug ipv6 prefix-list** command. Use the output to maintain prefix lists and compare IPv6 prefix entries to the lists.

   You can also display information about the configured prefix lists using the **show ipv6 prefix-list** command.

6. To check the assignment of prefixes from the IPv6 pool, compare the number of active subscribers in the output of the **show subscriber summary ipv6 all** command, with the number of IPv6 prefixes that are `in use` in the output of the **show ipv6 pool summary** command.

   If all IP prefixes are assigned from the pool, the total number of active subscribers and IP addresses in use in the pool should be the same.

   To determine if the prefixes are from static configuration or from pools, use the **show dhcpv6 server host** command.

   When you subtract the number of any statically assigned IPv6 prefixes from the total number of IPv6 prefixes in the output of the **show subscribers summary ipv6 all** command, the resulting number should match the total number of prefixes in use in the **show ipv6 pool summary** command.

```
[CORP]Redback#show subscriber summary ipv6 all
-----------------------------------------------------------
Total=16000

Type            Authenticating          Active     Disconnecting
PPP                   0                  0         0
PPPoE                 0                16000       0
DOT1Q                 0                  0         0
CLIPs                 0                  0         0
ATM-B1483             0                  0         0
ATM-R1483             0                  0         0
Mobile-IP             0                  0         0

[CORP]Redback#show ipv6 pool dhcpv6 summary
1          in use, 16000        free, 0      reserved
260        total, 99  available percentage
```

The output displays the total number of subscribers and their encapsulation type (context specific).

7. To determine if individual subscribers have been assigned static IPv6 prefixes, use the **show subscriber active** command with the ***username@domain*** or ***cct-handle*** arguments.

8. If the number of subscribers does not equal the number of IPv6 addresses in use, and you can't identify a cause, you can restart the AAA process and collect troubleshooting data for Customer Support with the **process coredump aaad** command.

9. If the numbers are similar, but subscribers are still not receiving IPv6 prefixes, continue with Section 5 on page 12.

5 Troubleshooting ND Problems

ND subscribers are bound to interfaces indirectly by being assigned a prefix matching the interface. The ND process advertises subscriber prefixes.

If you are experiencing IPv6 forwarding or routing failures, and suspect that ND is the problem, perform the tasks in Table 5.

Table 5 Troubleshooting ND Problems

Step	Command	Check?
ND neighbor Entry Missing From The Neighbor Cache Or Ping To The Neighbor Fails (for non-subscriber circuits)		
Examine the neighbor cache.	**show nd neighbor**	

Table 5 Troubleshooting ND Problems

Step	Command	Check?

Enable ND-related debugging messages and ping the missing address.
If the ping fails, examine the debug output logs to see how far the address resolution process went before failing.
	`debug nd circuit`
	`debug nd lc`
	`debug nd interface`
	`debug nd packet`
	`debug nd rib`
	`debug nd subscriber`
	`ping` *`ipv6-prefix`*	

Verify that the circuit associated with the neighbor is appropriately bound to the correct interface and in the Up state.	`show nd interface` *`if-name`*
	Or
	`show nd circuit` *`cct-handle`*	

View the ND send/receive packet counters for the interface.	`show nd statistics interface` *`if-name`*	
IPv6 Subscribers Do Not Receive Auto-Configured Prefixes Correctly (May Be Occurring After Reload or Switch-Over)		
Display the subscriber prefixes and verify that they are bound and up.	`show nd circuit` *`cct-handle`*	
Verify that the ND profile is advertising the prefix in RA messages; check the RA interval setting (context specific).	`show nd profile` *`prof-name`*	
Verify whether a Neighbor Solicitation message has been received. Displays global statistics for one or more ND router interfaces.	`show nd statistics interface` *`if-name`*	

Display the contents of the RA message to determine if it contains the subscriber prefix.	`debug nd packet`
	`debug nd packet [detail]`
	`debug nd interface`	

Verify that the status of the circuit, the IPv6 stack, and the interfaces are all Up.	`show nd circuit` *`cct-handle`*
	`show nd interface` *`if-name`*	

Verify that ND can send and receive messages over the interface.	`show nd statistics interface` *`if-name`*	

5.1 ND Neighbor Entry Missing from the Neighbor Cache or a Ping to the Neighbor Fails

To troubleshoot ND neighbors perform the following steps:

1. For nonsubscriber circuits, ND maintains a neighbor cache for all neighboring IPv6 addresses that have been successfully resolved, as well as for its own local IPv6 interface addresses. To display the cache, use the **show nd neighbor** command:

```
[local]Redback#show nd neighbor
IPv6 Address        Age  Link-layer Addr    State  Circuit
2:2:2:2::1                 0    00:30:88:23:23:7f  intf   4/11
2:2:2:2::2                 0    00:30:88:13:2e:9e  reach  4/11
16:1:1:1::1                0    00:30:88:03:01:14  intf   to-sierra-lg
16:1:1:1::2                20   00:30:88:15:b8:2e  reach  to-sierra-lg
2009::1                    0    00:30:88:03:01:0e  intf   2/3
2009::2                    66   00:30:88:03:01:10  probe  2/3
fe80::230:88ff:fe03:115    0    00:30:88:03:01:14  intf   to-sierra-lg
fe80::230:88ff:fe15:b82f   36   00:30:88:15:b8:2e  reach  to-sierra-lg
fe80::230:88ff:fe03:10e    0    00:30:88:03:01:0e  intf    2/3
fe80::230:88ff:fe03:110    77   00:30:88:03:01:10  probe   2/3
fe80::230:88ff:fe13:2e9e   50   00:30:88:13:2e:9e  reach   4/11
fe80::230:88ff:fe23:237f   0    00:30:88:23:23:7f  intf    4/11
```

Entries in which the State is marked as `Intf` are for the local IPv6 addresses defined on the SmartEdge router. All other entries are for remote IPv6 addresses learned from ND neighbor nodes.

2. If an expected IPv6 neighbor address is missing from the cache, use the **ping** command to verify that ND has lost reachability for this address.

3. Enable ND neighbor debugging with the following commands:

   - **debug nd lc**
   - **debug nd interface**
   - **debug nd rib**
   - **debug nd packet**

4. Ping the missing IPv6 address; for example, for address 2001::1, enter the **ping 2001::1** command.

5. Examine the debug logs; they should show the following:

   - The PPA has requested resolution for the address in a cache_miss message.
   - ND sent an NS message over the interface with a matching prefix, requesting resolution from the neighbor.
   - If the resolution succeeds, the debug output includes a received NA message.

   Use this information in the debug logs to answer the following questions:

- How far did the resolution attempt get before failure occurred?
- Did the PPA send a cache_miss request?
- Did the PPA provide the correct circuit that matches the interface with the correct prefix?
- Did ND send an NS message?
- Was an NA received? If so, was it valid?

6. To verify that the circuit received in the cache_miss matches the correct interface (matching prefix), use the **`show nd interface`** or **`show nd circuit`** commands.

   Verify that the correct prefix is assigned to the circuit, the circuit is appropriately bound and in the Up state, and an RA message is scheduled to be sent if the RA Interval is greater than 0.

7. To view the ND send/receive packet counters on the interface, use the **`show nd statistics interface`** ***`if-name`*** command.

   Use the **`show nd statistics`** command to see if NS messages are being sent and NA messages are being received. Examine the total number of subscriber circuits in the output. The `NS sent` counter should increment when a ping attempt is performed, and the `NA rcvd` counter should also increment if the neighbor is responding to the request.

5.2 Subscribers Are Not Receiving Auto-Configured Addresses Correctly

The AAA and ISM components handle all subscriber-related configuration information upstream of the ND process. AAA keeps the subscriber configuration and passes the details on to the ISM process, which in turn notifies ND (as well as other processes such as RIB, DHCPv6, and PPA) of all subscriber information for a particular circuit. When troubleshooting, verify that these components pass on the relevant subscriber information correctly. The most important subscriber information is the Framed-IPv6-Prefix attribute used by hosts to autoconfigure their addresses.

The scenario in which subscribers do not receive autoconfigured addresses may occur after a system reload or an XCRP switchover.

To troubleshoot, perform the following steps:

1. Verify that the subscriber circuit is in the Up state and bound correctly to its interface, and that the interface is also Up. To display the subscriber prefixes assigned to a subscriber circuit, use the **`show nd circuit`** ***`cct-handle`*** command.

2. Verify that the ND protocol is advertising the prefix in the RA message over the subscriber circuit.

Once it receives the IPv6 host prefix correctly, ND should advertise this prefix in RA messages to the subscriber. If ND is not advertising the prefix, verify that the RA-interval is not disabled. By default, it is enabled for each circuit, so at regular intervals ND retransmits an RA message containing the prefix. If you have disabled the RA interval by setting the value to 0, the messages are not sent.

IF the interval is not enabled, an RA message (with the prefix) may still be advertised to the subscriber, but only if the subscriber explicitly solicits the router for the prefix, using a Route Solicitation (RS) message. ND immediately responds to an RS with the RA message.

Verify that the RA interval is enabled with the **`show configuration nd`** or **`show nd profile`** ***`prof-name`*** command.

3. To verify that the prefix is assigned to the circuit correctly and find the ND profile assigned to the subscriber, use the **`show nd circuit`** ***`cct-handle`*** command.

4. Use the **`show nd statistics interface`** ***`if-name`*** command to verify whether the prefix has been solicited.

   The `RS received pkt` and `RA sent` counters should increment when the circuit comes up. If the interval is enabled, the `RA sent` counter should increment every time the RA timer elapses.

5. If it appears that an RA message is being sent over the interface but the prefix is not being received, use the following debug commands to display the contents of the RA message:

   - **`debug nd packet`**
   - **`debug nd packet detail`**
   - **`debug nd interface`**

If the RA is sent, a hex dump of the entire RA message should be displayed. Save the output and send it to Customer Support to analyze whether the prefix has been included in the message, and whether it was advertised with autoconfiguration enabled.

6 Troubleshooting DHCPv6-PD Problems

To troubleshoot problems with DHCPv6-PD, perform the steps in Table 6:

Table 6 Troubleshooting DHCPv6-PD Problems

Step	Command	Check?
DHCPv6 Client Does Not Receive A Prefix		
Check the DHCPv6 server host status.	**show dhcpv6 server host** To display the active DHCPv6 clients on a subnet, add the **subnet** keyword.	
Verify that the DHCPv6 server is enabled on an interface.	**show configuration interface**	
Check DHCPv6-pool prefix assignment.	**show ipv6 pool dhcpv6** ***pool-name***	
Turn on debugging and show the output. Verify that packets are reaching the server.	**debug dhcpv6 all** **show dhcpv6 log** **show dhcpv6 statistics detail**	
Traffic is Not Flowing Toward A Subscriber		
Check the DHCPv6 server host.	**show dhcpv6 server host**	
Display the DUID that the DHCPv6 server onboard the SmartEdge router is using to communicate with the DHCPv6 clients. Verify the client received the prefix and the route was installed.	**show dhcpv6 server duid** **show ipv6 route**	
After Reload or Switch-Over, IPv6 Prefixes Are Not Restored		
View DHCPv6 server statistics.	**show dhcpv6 statistics**	

6.1 DHCPv6 Client Does Not Receive a Prefix

1. Check the DHCPv6 server host status:

```
[local]Redback#show dhcpv6 server host
DHCPv6 Server Host Record:
--------------------------
DUID: 00:03:00:01:00:00:65:01:01:02
   IA Type: PD, IA ID: 0, T1: 6000, T2: 9600
   Prefix: 2001:a:1:1::/64
           preferred lifetime: 12000
           valid lifetime: 14000
           TTL: 9908
           expires at Fri Mar 12 01:27:55 2010
```

The output includes the DUID, Identity Association (IA) type (that prefixes are applied to), the IA ID, the prefix, and options for the server host.

2. If the DHCPv6 server is up, verify that the DHCPv6 server is enabled on an interface, using the **show configuration interface** command:

```
[local]Redback#show configuration interface
...
interface sub multibind
  ip address 10.1.1.1/16
  ipv6 address 2001:a:1:2::1/64
  dhcpv6 server interface
  ip pool 10.1.0.0/16
  ipv6 pool dhcpv6 8001:a:b:1::/64 8001:a:b:100::/64 name pd-pool threshold percentage falling 25 log 10/
 trap
...
```

Verify that the IPv6 addresses for the DHCPv6 server interface and the multibind subscriber interface match in the configuration. If they do not match, prefixes are not delegated to clients.

3. Examine the DHCPv6-pool prefix assignment using the **show ipv6 pool dhcpv6** ***pool-name*** command:

```
[CORP]Redback#show ipv6 pool dhcpv6 123-pool

Interface "123-pool":
   2001:a:c:1::/64          2001:a:c:4::/64         0 in-use,
    4 free,    0 reserved
   4001:a:c::/64            4001:a:c:ff::/64        1 in-use,
  255 free,    0 reserved
```

4. Turn on debug with the **debug dhcpv6 all** command.

```
[local]Redback# debug dhcpv6 all
Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-OPT: received Solicit
from fe80::200:65ff:fe01:102
Mar 12 10:47:09: %DHCPV6-7-PKT: Packet process handler: Process packet

   0     01 01 00 00 00 08 00 02 00 00 00 01 00 0a 00 03
  16     00 01 00 00 65 01 01 02 00 06 00 02 00 19 00 19
  32     00 0c 00 00 00 00 00 00 00 00 00 00 00 00

Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-PKT: XID: 1010000
Mar 12 10:47:09: %DHCPV6-7-OPT: get DHCP option elapsed time, len 2
Mar 12 10:47:09: %DHCPV6-7-OPT: elapsed time: 0
Mar 12 10:47:09: %DHCPV6-7-OPT: get DHCP option client ID, len 10
Mar 12 10:47:09: %DHCPV6-7-OPT: client DUID: 00:03:00:01:00:00:65:01:01:02
Mar 12 10:47:09: %DHCPV6-7-OPT: get DHCP option option request, len 2
Mar 12 10:47:09: %DHCPV6-7-OPT: requested option: IA_PD
Mar 12 10:47:09: %DHCPV6-7-OPT: get DHCP option IA_PD, len 12
Mar 12 10:47:09: %DHCPV6-7-OPT: IA_PD: ID=0, T1=0, T2=0
Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-PKT: client ID 00:03:00:01:00:00:65:01:01:02
Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-PKT: response XID: 2010000
Mar 12 10:47:09: %DHCPV6-7-PKT: set client ID (len 10)
Mar 12 10:47:09: %DHCPV6-7-OPT: client DUID: 00:03:00:01:00:00:65:01:01:02
Mar 12 10:47:09: %DHCPV6-7-PKT: set server ID (len 14)
Mar 12 10:47:09: %DHCPV6-7-OPT: server DUID: 00:01:00:01:13:2c:1c:c8:00:30:88:00:1c:61
Mar 12 10:47:09: %DHCPV6-7-PKT: set status code (len 2)
Mar 12 10:47:09: %DHCPV6-7-OPT: status code: no prefixes
Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-PKT:
Response source addr: fe80::230:88ff:fe00:1c61,
destination addr: fe80::200:65ff:fe01:102
Mar 12 10:47:09: [0001]: [9/3:1023:63/6/2/2]: %DHCPV6-7-PKT:
 Sent server response to client:
00:03:00:01:00:00:65:01:01:02 success.

   0     02 01 00 00 00 01 00 0a 00 03 00 01 00 00 65 01
  16     01 02 00 02 00 0e 00 01 00 01 13 2c 1c c8 00 30
  32     88 00 1c 61 00 0d 00 02 00 06
```

5. Show the output using the **show dhcpv6 log** command.

```
[local]Redback#show dhcpv6 log
Time   Evnt SubEvent         Key (ccth|ipv6addr|etc)   Details
------ ---- ---------------- -------------------------
347.42 Pkt  Pkt from Client  9/3:1023:63/6/2/3         Solicit
                    (cont'd) fe80::200:65ff:fe01:102
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.42 Pkt  Pkt to Client    9/3:1023:63/6/2/3         Advertise
                    (cont'd) fe80::230:88ff:fe00:1c61
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Pkt  Pkt from Client  9/3:1023:63/6/2/3         Request
                    (cont'd) fe80::200:65ff:fe01:102
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Int  IPv6 Add         9/3:1023:63/6/2/3
                    (cont'd) 2001:a:1:1::/64
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Int  Route Add        9/3:1023:63/6/2/3
                    (cont'd) 2001:a:1:1::/64
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Tmr  Timer Start      9/3:1023:63/6/2/3
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Hlpr Add              9/3:1023:63/6/2/3
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Int  Bind Add         9/3:1023:63/6/2/3
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.43 Pkt  Pkt to Client    9/3:1023:63/6/2/3         Reply
                    (cont'd) fe80::230:88ff:fe00:1c61
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.45 Hlpr Add-Resp         2001:a:1:1::/64
                    (cont'd) 00:03:00:01:00:00:65:01:01:02
347.45 ISM  R:Cct-Cfg        9/3:1023:63/6/2/3         CCT ethcfg
```

6. Verify that packets are reaching the DHCPv6 server, using the **show dhcpv6 statistics detail** command to display packet statistics.

```
[local]Redback#show dhcpv6 statistics
Current time: Fri Mar 26 08:51:14 2010
Last cleared: Never

PKT---------------------------------------------------
 Packets Rx      : 150    Packets Tx     : 150
 Solicit         : 2      Advertise      : 2
 Request         : 2      Confirm        : 0
 Renew           : 146    Rebind         : 0
 Reply           : 148    Release        : 0
 Decline         : 0      Reconfigure    : 0
 Relay Fwd       : 0      Relay Reply    : 0
 Info Req        : 0      Unknown Pkt    : 0

Dropped pkt-------------------------------------------
 Solicit          : 0     Advertise      : 0
 Request          : 0     Confirm        : 0
 Renew            : 0     Rebind         : 0
 Reply            : 0     Release        : 0
 Decline          : 0     Reconfigure    : 0
 Inform Req       : 0     Unknown Pkt    : 0
```

7. Save the output in case you need to send it to Customer Support. Turn off debugging and logging.

6.2 Traffic Is Not Flowing Toward a Subscriber

1. Check the DHCPv6 server host using the **show dhcpv6 server host** [**circuit** ***cct-handle*** | **duid** ***hex-string*** | **prefix** ***host-prefix*** | **subnet** ***prefix/length***] [**detail** | **summary**] command, as in the following example, which displays the status for a subscriber circuit (and active DHCPv6 clients on a circuit):

```
[local]Redback#show dhcpv6 server host circuit 9/3 pppoe 1
DHCPv6 Server Host Record:
--------------------------
DUID: 00:03:00:01:00:00:65:01:01:02
   IA Type: PD, IA ID: 0, T1: 6000, T2: 9600
   Prefix: 2001:a:1:1::/64
           preferred lifetime: 12000
           valid lifetime: 14000
           TTL: 8677
           expires at Fri Mar 12 01:27:55 2010
```

To determine a host DUID, use the **`show dhcpv6 server duid`** command.

2. Verify that the client received the prefix and the route was installed by entering the **`show ipv6 route`** command.

6.3 Prefixes Are Not Restored after Reload or Switchover

1. Check DHCPv6 server statistics with the **`show dhcpv6 statistics`** command; for sample output, see Step 6 in Section 6.1 on page 17.

7 Troubleshooting IPv6 Routes

To troubleshoot IPv6 routes in a dual-stack network, perform the steps in Table 7:

Table 7 Troubleshooting IPv4 and IPv6 Routes

Step	Command	Check?
A Ping Fails to Reach an IPv6 Circuit		

Identify the IPv6 address for the subscriber being tracked or debugged and the slots to which the SmartEdge OS downloaded next hops.	**`show subscriber address username`** ***`user@domain`***
	`show ipv6 route` ***`ipv6-address/ mask`*** **`detail`**	

Verify subscriber interface information. Could also display information about the interface bound to the Ethernet management port	**`show ipv6 interface`**	
IPv6 Subscriber Traffic Is Not Reaching A Destination Or Routes Are Missing		

Table 7 Troubleshooting IPv4 and IPv6 Routes

Step	Command	Check?
Enable IPv6 route debugging.	`debug ipv6 routing [fib-addition | fib-deletion | redist-addition | redist-deletion | route-addition | route-deletion]`	

Investigate IPv6 next hops.	`show ipv6 route NH-ID detail`
	Or
	`show ipv6 route next-hop ipv6-addr detail`	

Check ND detail.	`show nd circuit`
	`show nd circuit cct-handle detail`
	`show nd neighbor`	

For routing IPv6 packets, the end systems are responsible for fragmenting the packets on ingress and reassembling the packets on egress, 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. Any ICMPv6 "Packet too big" message 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 is acted upon.

- 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 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.

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.

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.

7.1 Ping Fails to Reach an IPv6 Circuit

To investigate an IPv6 route failure, start by collecting information about the subscriber circuit.

1. To identify the subscriber address being tracked or debugged, use the **`show subscriber address username`** ***`user@domain`*** command.

2. Using the subscriber IPv6 address, enter the **`show ipv6 route`** ***`ipv6-address/mask`*** **`detail`** command, as in the following example:

```
[local]Redback#show ipv6 route 2001:a:b:1::/64  detail

    Best match Routing entry for 2001:a:b:1::/64 is 2001:a:b:1::/64 , version 41
    Route Uptime 01:45:48
    Paths: total 1, best path count 1

    Route has been downloaded to following slots
     04/0, 06/0, 11/0

    Path information :

      Active path :
      Known via subscriber nd , distance 15, metric 0,
      Tag 0, NH-ID 0x34300003, Adj ID: 0xA00000A, Interface sub
      Circuit 11/5:1023:63/6/2/5
```

Use this command to determine the slots to which the route has been downloaded and the next-hop ID (NH-ID) and adjacency ID (Adj ID).

7.2 IPv6 Subscriber Traffic Is Not Reaching a Destination or Routes Are Missing from the RIB

Using the information collected in Section 7.1 on page 22, continue troubleshooting the route with the following steps:

1. Enable IPv6 route debugging with the **`debug ipv6 routing`** [**`fib-addition`** | **`fib-deletion`** | **`redist-addition`** | **`redist-deletion`** | **`route-addition`** | **`route-deletion`**] command.

2. To investigate a specific IPv6 next hop using the NH-ID, enter the **show ipv6 route next-hop NH-ID detail** command as in the following example:

```
[local]Redback#show ipv6 route next-hop 0x34300003 detail
    ** = Via interface
    Next Hop Tbl Version :         41
    Current Next Hops    :          4

NH-ID                 Ref Cnt NH-IP            Via-NH     Interface

0x34300003                3/0                             sub
Adj-id          : 0xA00000A
Mac Address     : 00:1f:3f:3b:7a:5e
Info-Version    : 38                 Node-Version    : 38
Fib Card bits   : 0x4000468          Nh Client bits : 0x0
Info flags      : 0x1                Lsp ifgrid      : 0x0
IF-GRID         : 0x10000001


Next-hop has been downloaded to following slots
 04/0, 06/0,11/0, 11/1
```

In the output, `04/0` is the ingress Packet Processing ASIC (IPPA) of slot 4 and `11/1` is the egress PPA (EPPA) of slot 11.

To look up the next hop by IPv6 address, use the optional ***ipv6-addr*** argument (fully qualified IPv6 address) instead of the next-hop ID.

3. To check ND detail, use the **show nd circuit** and **show nd circuit *cct-handle* detail** commands.

```
[local]Redback#show nd circuit
ND circuit           Grid       Mac Address       Status    Encap Type
11/5:1023:63/6/2/5   0x10000001 00:30:88:03:03:ae Up        eth pppoe ppp

[local]Redback#show nd circuit 11/5:1023:63/6/2/5 detail
Circuit handle    : 11/5:1023:63/6/2/5   Intf grid       : 0x10000001
Mac addr          : 00:30:88:03:03:ae    Status          : Up
Encap type        : eth pppoe ppp        Circuit type    : Subscriber
Interface name    : sub
Circuit MTU       : 1492
Router lifetime   : 1800
ND Profile        : nd1                  IPv6 Proto State: Up
Subscriber Prefix : 2001:a:b:1::/64
```

To determine the ***cct-handle*** argument, use the **show ipv6 route next-hop *grid* detail** command; see Step 2.

You can also use the **show nd neighbor** command to check ND neighbor status; for sample output, see Step 1.

8 Troubleshooting IPv6 Forwarding and ACLs

To troubleshoot IPv6 forwarding and access control lists (ACLs), perform the steps in Table 8.

Table 8 Troubleshooting IPv6 Forwarding and ACLs

Step	Command	Check?
Handling IPv6 ACL Counters		
Enable IPv6 ACL Counters For:		
Subscribers (in subscriber configuration mode):	`access-list count ipv6 ipv6-policy`	
ACLs attached to QoS metering policies (in one of the following modes): • ATM PVC • dot1q PVC • Frame Relay PVC • Link group • Port • Subscriber	**qos policy metering** *pol-name* **[ip | ipv6] acl-counters**	
ACLs attached to QoS policing policies (in one of the following modes): • ATM PVC • dot1q PVC • Frame Relay PVC • Link group • Port • Subscriber	**qos policy policing** *pol-name* **[ip | ipv6] acl-counters**	
Forward policies (in subscriber configuration mode):	**forward policy** *pol-name* **in [[ip] [ipv6] acl-counters**	

Table 8 Troubleshooting IPv6 Forwarding and ACLs

Step	**Command**	**Check?**
Clear IPv6 Counters	To clear policy ACL counters—**clear access-group ipv6 {forward | qos} cct-handle** To clear filtering ACL counters—**clear access-group ipv6 filter cct-handle**	
Troubleshooting IPv6 QoS Policy ACLs with Static Configuration		
Verify that configured policies were downloaded and applied.	**show qos policy *pol-name*** **show access-group ipv6 qos *cct-handle* {in | out} [detail]**	
Verifying or Troubleshooting Class-based Rate Limiting		
Observe traffic in classes.	Enter the **show access-group ipv6 qos *cct-handle* in count** command twice, approximately (2–3 minutes apart, using the subscriber circuit (***slot/port:ch:sub*[*subsub*]**) where the counters were enabled for the ***cct-handle*** argument.	
Troubleshooting IPv6 ACLs For Subscribers		
Troubleshoot Filter-ID and Ascend-Data-Filter Attributes.	Use the **show subscribers active** command to verify that `Protocol Stack Dual` indicates that both IPv4 and IPv6 are supported for this subscriber, and that DQP (`qos-dynamic-param`) and DPF (`dynamic policy filter` are enabled.	
	show access-group ipv6 filter *cct-handle* {in | out} all	

Table 8 Troubleshooting IPv6 Forwarding and ACLs

Step	Command	Check?
To verify traffic on the active XCRP controller card, enable IPv6 ACL rule counters for the circuit and then enter the `show access-group` command twice, 2–3 minutes apart.	`show access-group ipv6 qos` *`cct-handle`* `in count`	
Troubleshoot IPv6 Policy ACLs With RADIUS Guided Policies		
Display ACLs on active subscribers.	`show subscribers active`	
Verify that ACLs are applied for forward policies.	`show access-group ipv6 forward` *`cct-handle`* `{ in | out} all`	
Verify forward policy configuration.	`show forward policy` *`pol-name`*	
Troubleshooting IPv4 and IPv6 HTTP Redirect		
Show redirect traffic counters.	`show http-redirect circuit all`	
Enable HTTP redirect debug messages.	`debug hr all`	
Troubleshooting RSE Service Accounting on Subscribers		

Verify RSE configuration:	`show subscribers active`
	`show access-group ipv6 {qos | forward } circuit {in | out} detail`
	`show access-group` *`cct-handle`*	

Troubleshooting RSE Limitations on LACs		
Check the configuration details: • Only circuit level service accounting is supported. • QoS policies must be a non-RG policies, with no ACLs.		
Enable debugging for IPv6 policy ACLs.	*debug ipv6 policy access-list*	
Use other IPv6 ACL and forwarding troubleshooting commands.		

8.1 Handling IPv6 ACL Counters

8.1.1 Enable Counters

Warning!

Enabling IPv6 ACL and forwarding related counters can impact SmartEdge router performance. Use them carefully.

Before you can display IPv6 filtering and policy ACLs, you need to enable counters in one of the following ways:

- Subscribers—To enable IPv6 filtering ACL counters for a subscriber through the subscriber record, the default subscriber profile, or a named subscriber profile, enter the *access-list count ipv6 ipv6-policy* command in subscriber configuration mode. For both IPv4 and IPv6, this command is not applied to active subscribers. The following example enables ACL counters for a default subscriber policy:

```
[local]Redback(config-ctx)#subscriber default
[local]Redback(config-sub)#qos policy metering pol23
[local]Redback(config-sub)#access-list count ipv6-policy
```

- IPv6 policy ACLs—To enable IPv6 policy ACL rule counters for policing, metering, and forward policy classification, on a static circuit, add the **acl-counters** keyword to the policy configuration. To enable counters for an IPv6 ACL use the **ipv6 acl-counters** construct or to enable counters for dual stack ACLs use the **ip ipv6 acl-counters** construct. The following example enables counters for a QoS metering policy at the port level:

```
[local]Redback(config)#port ethernet 4/1
[local]Redback(config-port)#qos policy meterin
g pol23 ipv6 acl-counters
```

Note: You enter **show** commands at the level where the counters were configured to view the counters. For example, for subscriber circuits, enter the **show** command specifying the subscriber circuit. For ACL counters enabled on a port, specify the port in the **show** command. For example, if you enabled ACL counters on the QoS policy on a port, entering the **show circuit counters *slot/port* detail** command shows how many packets passed through the policy processing on that port.

8.1.2 Clear Counters

Circuit IPv6 ACL drop counters are persistent during circuit life per protocol. They are applied without configuration.

The best way to verify traffic is to enable rule counters. Use them if you understand traffic patterns and performance is not an issue.

Use the following guidelines for IPv6 ACL counters:

- IPv6 ACL counters impact traffic performance.
- IPv6 ACL counters are cleared with any configuration update.
- To clear IPv6 ACL counters at other times, use:
  - For policy ACLs—**`clear access-group ipv6 {forward | qos}`** ***`cct-handle`*** **`{in | out}`**
  - For filtering ACLs—**`clear access-group ipv6 filter`** ***`cct-handle`*** **`{in | out}`**

  The ***`cct-handle`*** argument is the ***`slot`***`/`***`port`***`:`***`ch`***`:`***`sub`***`[`***`subsub`***`]` for the circuit where the policy is configured.
- IPv6 ACL counters should be applied to subscribers before subscribers are brought up.
- IPv6 ACL rule counters cannot be applied if service accounting is configured.

8.2 Troubleshooting IPv6 QoS Policy ACLs with Static Configuration

To troubleshoot IPv6 policy ACLs with static configuration, use the commands in the following sections.

Note: Some hidden commands or parameters are used in the tasks described in this document provide useful information, but we recommend you use them only as described. They can change without notice, are not fully tested, and are not officially supported. To avoid unexpected results, use them only when necessary, with the correct parameters and syntax.

8.2.1 Verify that Configured Policies Were Downloaded and Applied

1. To display information about a configured QoS policy, including whether it was downloaded to a line card, use the **`show qos policy`** ***`pol-name`*** command. The following example shows that a QoS policing policy was applied to the line cards in slots 2 and 3:

```
[local]Redback#show qos policy policing pol1
Policy-Name                    Type      Grid     Qs Slots   Ports   Bound   DnLd   Status
pol1                           police    1        0  2       2       in             updt

Slot#:        1   2   3   4   5   6   7   8   9  10  11  12  13  14
iPPA dnld:
ePPA dnld:
iPPA ports:       1   1
ePPA ports:
```

2. To determine whether the IPv6 ACL was applied on the line card, use the **show access-group ipv6 qos** ***cct-handle*** **{in | out} [detail]** command.

   Note: The **show access-group** command is context specific.

   The following example displays the provisioning for the QoS ACL, `DPF QO6_8000001F` with status, `Applied`:

```
[local]Redback#show access-group ipv6 qos 3/3 vlan-id 3:1 pppoe 6 out detail

QOS ACL         : DPF QO6_8000001F
ACL context     : local
Circuit         : 3/3 vlan-id 3:1 pppoe 6
Direction       : Out          ACL status   : Applied
Count           : Service      Log          : N/A
Number of rules: 2
```

If the IPv6 ACL status in the output is `Applied`, the ACL was applied.

8.2.2 Verify or Troubleshoot Class-Based Rate Limiting

To verify or troubleshoot class-based rate limiting observe traffic counters.

8.2.2.1 Observe Traffic in Classes

To observe traffic in the classes:

Note: Perform this task only if the traffic pattern is known and performance is not an issue.

1. To enable IPv6 policy ACL counters on static circuits, use the **ipv6 acl-counters** construct when binding the QoS policy to the static circuit; for more information, see Section 8.1 on page 27.

2. To view the incremented counters during traffic flow, enter the **show access-group ipv6 qos** ***cct-handle*** **in count** command twice, approximately (2–3 minutes apart, using the subscriber circuit (***slot*****/*****port*****:*****ch*****:*****sub*****[*****subsub*****]**) where the counters were enabled for the ***cct-handle*** argument.

8.3 Troubleshooting IPv6 ACLs For Subscribers

This section describes how to troubleshoot the Filter ID and ADF attributes for subscribers.

8.3.1 Troubleshoot Filter-ID and Ascend-Data-Filter Attributes

Use the following steps to troubleshoot the Filter-Id and ADF attributes:

1. To verify that all ACLs and rules were configured correctly:

    - Use the **show subscriber active** command to determine the status of ACLs and attributes. In the following example, `Protocol Stack Dual` indicates that both IPv4 and IPv6 are supported for this subscriber, and that DQP (`qos-dynamic-param`) and DPF (`dynamic policy filter`l) are applied.

```
[local]Redback#show subscriber active
user1@local
        Session state Up
        Circuit   4/7 vlan-id 101 pppoe 40
        Internal Circuit   4/7:511:63:31/6/2/40
        Interface bound  SUBSCRIBER
        Current port-limit unlimited
        Protocol Stack Dual
        ip address 10.1.0.2 (applied)
        ppp mtu 1492 (applied from sub_default)
        dns primary 155.53.247.12 (applied from sub_default)
        dns secondary 155.53.12.12 (applied from sub_default)
        Use DSL downstream rate from ANCP only (applied from sub_default)
        Adjust QOS policy at the CVLAN circuit level (applied from sub_default)
        propagate qos from ip class-map PD-SET-L3-IN (applied from sub_default)
        propagate qos to ip [class-map] PD-SET-L3-OUT (applied from sub_default)
        idle timeout direction in (applied from sub_default)
        Dual-stack-failure force-down 1 (applied from sub_default)
        timeout absolute 86400 (applied from sub_default)
        timeout idle 900 (applied from sub_default)
        qos-policing-policy POL [svc mask: 0x0001] (applied)
        qos-metering-policy MET [svc mask: 0x0001] (applied)
        service  (applied)
           [svc id: 0] IPOne cip=1.1.1.1/32 vip=1.1.1.2/32 cip6=5000::1/128 vip6=5000::2/128 cp
ol=160 vpol=600 (acct enabled)
        qos-dynamic-param  [svc mask: 0x0001] (applied)
           [svc id: 0] police-class-rate VuB rate-absolute 600 (applied)
           [svc id: 0] police-class-rate VuC rate-absolute 160 (applied)
           [svc id: 0] police-class-conform VuB mark-dscp cs5 (applied)
           [svc id: 0] police-class-conform VuC mark-dscp cs7 (applied)
           [svc id: 0] meter-class-rate VdC rate-absolute 160 (applied)
           [svc id: 0] meter-class-rate VdB rate-absolute 600 (applied)
           [svc id: 0] meter-class-conform VdC mark-dscp cs7 (applied)
           [svc id: 0] meter-class-conform VdB mark-dscp cs5 (applied)
        service-acct (in)  [svc mask: 0x0001] (applied)
           [svc id: 0] qos class-mask 0x03
        service-acct (out)  [svc mask: 0x0001] (applied)
           [svc id: 0] qos class-mask 0x03
        service-interim-acct-interval  [svc mask: 0x0001] (applied)
           [svc id: 0] unlimited
        Framed-IPV6-Prefix  2001:a:b:1::/64 (applied)
        Delegated-IPV6-Prefix 3001:2:2:55::/64 (applied)
        Framed Interface Id 200:70ff:fe02:102 (applied)
        Ipv6-ND-Profile nd-prof (applied)
        Ipv6-DNS  primary 2000::200:a00:20ff:fe99:a998  (applied from sub_default)
        Ipv6-DNS secondary 2000::201:a00:20ff:fe99:0 (applied from sub_default)
        dynamic policy acl ipv6  [svc mask: 0x0001] (applied)
           [svc id: 0] ipv6 out forward srcip 5000::1/128 class VdC qos
           [svc id: 0] ipv6 out forward srcip 5000::2/128 class VdB qos
           [svc id: 0] ipv6 in forward dstip 5000::1/128 class VuC qos
           [svc id: 0] ipv6 in forward dstip 5000::2/128 class VuB qos
          ipv6host entries installed by PD: (cur_entries 1)
                3001:2:2:55::/64
```

You can also use the **show access-group ipv6 filter *cct-handle* {in | out} all** command.

8.3.2 Troubleshoot IPv6 Policy ACLs With RADIUS Guided Policies (DPF and DQP Attributes)

Use the following procedures to troubleshoot DQP and DPF attributes.

To verify IPv6 ACL and forward policy configuration:

1. To show ACLs on active subscribers, use the **show subscriber active** command; for an example, see Step 1.

2. To verify that ACLs are applied for forward policies, use the **show access-group ipv6 forward *cct-handle* { in | out} all** command; for an example, see Step 2.

   And for QoS policies, use the **show access-group ipv6 qos *cct-handle* {in | out} detail** command.

3. To verify forward policy configuration, use the **show forward policy *pol-name*** command, as in the following example, which displays a forward policy for redirecting subscribers, `pol-red`, that is applied to the line cards in slots 2 and 3:

```
[local]Redback#show forward policy pol-red
Policy-Name       Type     Grid    Qs Slots  Ports  Bound  DnLd  Status
pol-red           forward  2       0  2      2      in           updt

Slot#:       1   2   3   4   5   6   7   8   9  10  11  12  13  14
iPPA dnld:
ePPA dnld:
iPPA ports:      1   1
ePPA ports:

Class-Name          Action  Mode  IP-Addr/Option   Bound  Int,msec  Output-Name
redirect            redir   local

Total policy map: 1
```

4. To display the IPv6 ACL drop count, enable IPv6 ACL rule counters on subscribers (see Section 8.1 on page 27) and enter the **show access-group** command several times 2-3 minutes apart to display traffic status. For more information, see Step 2.

8.3.3 Troubleshooting IPv4 and IPv6 HTTP Redirect

To troubleshoot the HTTP server:

1. To show redirect traffic counters, use the **show http-redirect circuit all** command. The following examples show the output for dual stack, IPv6 only, and IPv4 only.

```
[isp2]Redback#show http-redirect circuit all
user1@isp2
        Circuit    : 5/8 pppoe 4
        URL        : http://4.4.4.4:80
        IPv6 URL   :
        Redir Count: 2               Drop Count: 0

[isp2]Redback#show http-redirect circuit all
IPv6 only
user1@isp2
        Circuit    : 5/8 pppoe 3
        URL        :
        IPv6 URL   : http://[3000:1:2::1]:80

        Redir Count: 2               Drop Count: 0
IPv4 only
[isp2]Redback#show http-redirect circuit all
user1@isp2
        Circuit    : 5/8 pppoe 4
        URL        : http://4.4.4.4:80
        IPv6 URL   :
        Redir Count: 2               Drop Count: 0
```

2. You can also enable HTTP redirect debug messages with the **`debug hr all`** command.

   In the following example, look for `redirect !!!!!!!!!` to verify HTTP redirect is occurring.

```
[isp2]Redback#debug hr all
[isp2]Redback#!!!!!!! v6 redirect !!!!!!!!!
[isp2]Redback#Jan 6 20:44:11: [0003]: %HR-7-SESS: 2001:a:b:1::dead:beef-1000:1:2::1 fd 15 state
READING cct 5/8:511:63:31/6/2/2
Jan 6 20:44:11: %HR-7-SUB: Notify AAA of temporary redirected subscriber ccts
Jan 6 20:44:12: %HR-7-EVNT: Handle work for cleaning up temporary redirected subscriber ccts

[isp2]Redback#!!!!!!! v4 redirect !!!!!!!!!
[isp2]Redback#Jan 6 20:45:23: [0003]: %HR-7-SESS: 17.1.1.2-1.1.1.2 fd 15 state
READING cct 5/8:511:63:31/6/2/2
Jan 6 20:45:23: %HR-7-SUB: Notify AAA of temporary redirected subscriber ccts
Jan 6 20:45:24: %HR-7-EVNT: Handle work for cleaning up temporary redirected subscriber ccts
```

There are no special counters for HTTP redirect. There are also keywords in this command to display debug messages about errors, events, policies, sessions, and subscribers.

8.4 Troubleshooting RSE Service Accounting on Subscribers

In this release, the service accounting commands are not changed. If you configure service accounting, it is applied independent of both stacks for dual stack subscribers; IPV4 and IPv6 service counters are aggregated in the same bucket on line cards.

8.4.1 RSE support for IPv6 ACL Attributes (Service Accounting)

To verify RSE configuration, use the following commands:

- **`show subscriber active`**; for an example, see Step 1.

- **`show access-group ipv6 {qos | forward } circuit {in | out} detail`** ; for an example, see Step 2.
- **`show access-group`** ***`cct-handle`***

8.4.2 Troubleshooting RSE Limitations on LAC

Check the following configuration details:

- Only circuit level service accounting is supported.
- QoS policies must be a non-RG policies, with no ACLs.