CHAPTER 13

Topology Hiding

The Cisco Unified Border Element (SP Edition) distributed model for the Cisco ASR 1000 Series Routers has a primary purpose in protecting the network and providing seamless interworking functions. Cisco Unified Border Element (SP Edition) can protect the network by hiding the network addresses and names for both the access (customer) side and the backbone (network core) side. Cisco Unified Border Element (SP Edition) also provides network protection for firewalls or home gateway users with private addresses.

When a user connects to the outside network, its IP address and port needs to be properly translated to protect its identity. The data border element (DBE) performs translation of IP addresses and port numbers via Network Address and Port Translation (NAPT) and Network Address Translation (NAT) Traversal functions in both directions.

The DBE implementation supports the H.248 NAPT package, the IP NAT Traversal Package, and the ETSI TS 102 333 specification for NAT Traversal, but only one package can be active. Latch and Relatch functions of the NAT Traversal are supported by the IP NAT Traversal package. Support for these packages help protect IP addresses of the endpoints going across the other side of the network.

The NAPT implementations on the DBE described in more detail in this chapter are summarized below:

- IPv4 Twice NAPT—Where both access side and backbone side addresses are protected. In Twice NAPT, both the IP address and port are translated to a local IP address and port; and both of the end points on each side see the SBC address as a destination address.
- IPv6 Single NAPT for signaling packets—This function is useful for protecting the signaling infrastructure part of the backbone side. The backbone side is able to identify the address of the customer; however, for the customer, only the interface address of the DBE is visible.
- IPv6 No NAPT for media packets—With this method, there is no privacy on the customer side or backbone side. Both sides know each other's address and the DBE transparently passes the packets.

Cisco Unified Border Element (SP Edition) was formerly known as Integrated Session Border Controller and may be commonly referred to in this document as the session border controller (SBC).

For a complete description of the commands used in this chapter, see *Cisco Unified Border Element (SP Edition) Command Reference: Distributed Model* at:

http://www.cisco.com/en/US/docs/ios/sbc/command/reference/sbc_book.html

Contents

This chapter provides information about the following topics:

NAPT and NAT Traversal

NAPT and NAT Traversal are described in Chapter 9, "Security in Cisco Unified Border Element (SP Edition) Distributed Model."

IP NAPT Traversal Package and Latch and Relatch Support

The IP NAPT Traversal Package and Latch and Relatch Support functions are described in Chapter 9, "Security in Cisco Unified Border Element (SP Edition) Distributed Model."

IPv4 Support for Twice NAPT

The DBE successfully forwards media through Twice Network Address and Port Translation (NAPT) pinholes that form coupled pairs. For Twice NAPT hairpinning, the DBE forwards media on demand. The SBE sees no differences between Twice NAPT hairpins and Twice NAPT non-hairpins.

When forwarding media, a hairpinned pair behaves the way two separate pinholes behave, except that a packet going through a coupled pair has its IP Time-to-Live counter decremented only once, not twice.

Note Twice NAPT is only supported on IPv4.

IPv6 Inter Subscriber Blocking

Inter subscriber blocking prevents a subscriber from connecting to other subscribers without first going through a successful signaling/call setup process and having a termination established for the stream.

When the SBC DBE is implemented in the IPv4 environment, the DBE supports Twice NAPT, which has well-defined local media IP addresses or IP address pools. In the IPv4 environment, the DBE drops all SBC traffic destined for SBC local media IP addresses if there is no in-service termination successfully retrieved.

However, in the IPv6 environment, the SBC DBE only supports No NAPT for media pinholes, which, unlike Twice NAPT, does not have well-defined local media IP addresses or IP address pools. Because the same DBE Router routes non-SBC IPv6 traffic (which does not have SBC termination flow entry

whatsoever), the default operation for IPv6 traffic that does not have a corresponding termination flow entry is to continue to switch these packets. This can result in a situation where subscribers are still able to connect to other subscribers through the SBC DBE Router without completing the signaling and call setup process.

To support inter subscriber blocking in the IPv6 environment, you must classify subscribers at the ingress interface so that non-SBC traffic and SBC traffic can be differentiated.

For example, you might configure QoS at the ingress interface to mark all subscriber traffic with an unused unique differentiated services code point (DSCP) value, and then configure QoS at the egress interface to drop all the packets with this unused unique DSCP value. For SBC traffic with a termination flow entry, a separate DSCP value should be used to replace the original DSCP for these SBC packets as part of the normal diffserv package processing. As a result of this configuration, SBC packets with a session established will be routed and forwarded though the egress interface without being dropped because they have an SBC DBE-updated DSCP value. Depending on the QoS classification, you also have the flexibility of blocking partial traffic between subscribers without a session established or blocking all the traffic between them.

IPv6 inter subscriber blocking can be implemented using two methods: Quality of Service (QoS) policy-map-based inter subscriber blocking, or access control list (ACL)-based inter subscriber blocking.

QoS Policy Map-Based Inter Subscriber Blocking Method

In the following example of the QoS policy map-based inter subscriber blocking method, all the packets entering the Router (DBE) (through 0/1.1101) are marked using the policy-map INPUT_POLICY with DSCP=default (0). Any packets leaving the DBE (gigabitEthernet 0/2) with DSCP=0 will be blocked by the class-map IPv6_intersubscriber in the policy-map CORE_OUT. IPv6_intersubscriber uses the ACL ipv6_dscp0_any.

```
Router# show run interface gigabitEthernet 0/1.1101
...
Current configuration : 711 bytes
!
interface GigabitEthernet0/1.1101
encapsulation dot1Q 1101
ip dhcp relay information option subscriber-id 1101
ip address 12.21.1.1 255.255.255.0
ip access-group InFilter_IPv4 in
ip access-group OutFilter_IPv4 out
ip verify unicast reverse-path
ip helper-address 12.1.99.2
pppoe enable group global
ipv6 address 2000:12:21:1::1/64
ipv6 address FE80::1 link-local
ipv6 traffic-filter InFilter_IPv6 in
ipv6 traffic-filter OutFilter_IPv6 out
ipv6 verify unicast reverse-path
ipv6 mld explicit-tracking
ipv6 mld access-group VLAN1
ipv6 dhcp relay destination 2000:12:1:99::2
snmp trap link-status
no cdp enable
service-policy input INPUT_POLICY
service-policy output PARENT_OUTPUT_POLICY
end
```

```
Router# show policy-map INPUT_POLICY
Policy Map INPUT_POLICY
  Class class-default
  set dscp default

Router# show policy-map PARENT_OUTPUT_POLICY
Policy Map PARENT_OUTPUT_POLICY
  Class class-default
  Average Rate Traffic Shaping
  cir 100000000 (bps)
  service-policy CHILD_OUTPUT_POLICY

Router# show policy-map CHILD_OUTPUT_POLICY
Policy Map CHILD_OUTPUT_POLICY
  Class EF
  set cos 5
  set dscp ef
  priority level 1 10000 (kbps)
  Class AF4
  set cos 4
  priority level 2 75000 (kbps)
  Class AF1
  set cos 1
  priority level 2 5000 (kbps)
  Class IPv6_intersubscriber
  police cir 8000 bc 1500
  conform-action drop
  exceed-action drop
  Class class-default
  set cos 0
  bandwidth 9990 (kbps)
  queue-limit 1 packets

Router# show class-map IPv6_intersubscriber
Class Map match-all IPv6_intersubscriber (id 16)
  Match access-group name ipv6_dscp0_any

Router# show ipv6 access-list ipv6_dscp0_any
IPv6 access list ipv6_dscp0_any
  permit ipv6 any any dscp default sequence 10
  deny ipv6 any any sequence 20

Router# show run interface gigabitEthernet 0/2
...
Current configuration : 505 bytes
!
interface GigabitEthernet0/2
description to AER1-1 gi0/0/0/2
ip address 12.11.21.2 255.255.255.252
ip access-group Core_InFilter in
load-interval 30
carrier-delay msec 5
media-type sfp
speed 1000
duplex full
negotiation auto
ipv6 address 2000:12:11:21::2/64
ipv6 traffic-filter Core_InFilter_IPv6 in
ipv6 traffic-filter OutFilter_IPv6 out
no ipv6 mld Router
snmp trap link-status permit duplicates
keepalive 1
service-policy output CORE_OUT
```

```
hold-queue 1000 in
hold-queue 1000 out
end

Router# show policy-map CORE_OUT
Policy Map CORE_OUT
 Class IPv6_intersubscriber
 police cir 8000 bc 1500
 conform-action drop
 exceed-action drop
 Class class-default
```

ACL-Based Inter Subscriber Blocking Method

In the following example of the ACL-based inter subscriber blocking method, packets entering the DBE from the access side are marked with DSCP=0 using the same INPUT_POLICY as the QoS method above, but packets leaving the DBE use the ACL OutFilter_IPv6 as follows:

```
Router# show ipv6 access-list OutFilter_IPv6
IPv6 access list OutFilter_IPv6
 permit icmp any any packet-too-big sequence 10
 deny icmp any any sequence 20
 deny ipv6 any any dscp default sequence 40
 permit ipv6 any any sequence 50
```

Restrictions for DBE ACL-Based Inter Subscriber Blocking Method

The following is a restriction pertaining to DBE support for IPv6 inter subscriber blocking:

Because the configuration of inter subscriber blocking in the IPv6 environment relies on Cisco IOS QoS to mark the DSCP value in the ingress feature process, the original DSCP value of the packets arriving at the DBE Router will not be preserved.

IPv6 Support

IPv6 support includes the following functionality:

- The DBE supports IPv6 pinholes for both media endpoints and signaling endpoints.

  See the "IPv6 Pinholes" section on page 13-6.

  Note *Pinhole* is an informal term for a pair of terminations in the same stream and same context.

- Media flows do not support Network Address and Port Translation (NAPT); they must be No NAPT.

  As a result, you cannot configure any media addresses under IPv6. Media flows may consist of voice or video.

- Signaling flows support Single NAPT.

  You are able to configure signaling addresses under IPv6.

The DBE examines all IPv6 packets that arrive from the network and determines which ones belong to authorized SBC media streams. The DBE normally uses the destination (and possibly the source) IP address and port for packet classification. The DBE identifies packets belonging to an authorized media stream as SBC packets and applies the appropriate traffic policing rules to the packets. The counter showing number of packets received is modified.

After that, SBC performs packet processing and updating. The packet is forwarded out of the specified interface. IPv6 packet forwarding works in the same way as IPv4 packet forwarding, except for a few differences in the IP header processing.

Single NAPT for signaling means that packets arriving from an endpoint are addressed to an SBC media address. When they are passed to the media gateway controller (MGC), also know as an SBE, the packets need to keep the endpoint's source IP address and port number. Therefore, only destination addresses and ports are translated in Single NAPT. When the MGC/SBE sends a reply back to the endpoint, the packet has the endpoint's IP address as the destination address, and the MGC/SBE IP address as the source address. In Single NAPT, the DBE changes the source address to use the DBE IP address. See the "IPv6 Single NAPT for Signaling" section on page 13-7.

No NAPT means the received SBC packets do not contain any DBE local addresses because the DBE does not translate any IP addresses and ports during packet forwarding. The DBE rewrites neither the source nor destination addresses and ports in both directions. See the "IPv6 No NAPT Support for Media Flows" section on page 13-7.

IPv6 Pinholes

DBE support for IPv6 pinholes includes the following functionality:

- The DBE supports forwarding of media from one IPv6 endpoint to another IPv6 endpoint.
- The DBE supports IPv4 and IPv6 endpoints simultaneously. However, no interworking between IPv4 and IPv6 endpoints is supported. IPv4 endpoints can only forward media to other IPv4 endpoints and IPv6 endpoints can only forward media to other IPv6 endpoints.
- The DBE supports configuration of IPv6 pinhole addresses and pinhole address pools.
- DBE supports signaling pinholes using IPv6 addresses.

Support is added for the MGC to specify the address and port in the Megaco local descriptor for terminations as one of the following:

- An address and port that are not owned by the SBC and not configured in a media address range on the SBC, but matching the remote address and port for the other termination in the stream.
- An address range, in the form of a classless interdomain routing (CIDR) mask (for example, 10.13.8.0/21) together with a 0 port number, that does not overlap with any address ranges owned by the SBC or any media address range configured on the SBC, but the address and port match the gm/rsam (Gate Management/remote source address mask) for the other termination in the stream.

SBC recognizes these "local" addresses as signifying Single NAPT pinholes. And if specified for both terminations in the stream, SBC recognizes these addresses as No NAPT pinholes. All pinholes only forward packets to a full destination address and port that was either specified in the remote descriptor or latched to (within a gm/rsam that matches the local address mask).

IPv6 No NAPT Support for Media Flows

To support IPv6 on the DBE deployment, media flows do not support NAPT. No NAPT support means that no IP addresses and ports are translated by the DBE from a private address to a public address (for multiple users to share a single public address).

Because media addresses and ports are not translated, media flows on both sides of the media address are programmed with private, local addresses and ports that do not belong to the DBE. These local addresses and ports are specified by the MGC to match the remote address and port on the opposite side of the media address. Traffic in both directions is addressed directly to the remote endpoint on the other side of the DBE. The DBE rewrites neither the source nor destination addresses and ports in both directions because the DBE does not translate any IP addresses and ports during packet forwarding. Neither the source address nor destination address contains any DBE local media addresses.

Figure 13-1 illustrates a No NAPT media flow through the DBE between user side A and user side B.

Figure 13-1 No NAPT Media Flow

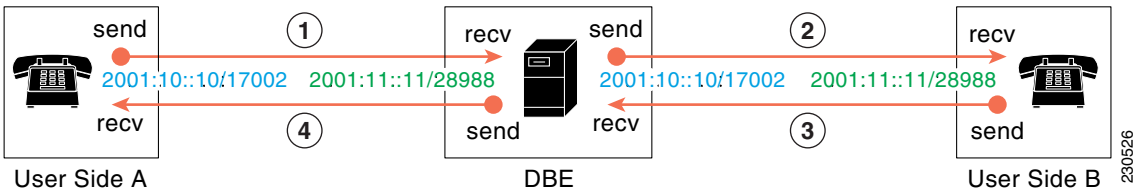

1. User side A sends a packet from IP address and port 2001:10::10/17002 to destination address and port 2001:11::11/28988 on side B. The DBE intercepts this packet and matches it to the side A flow.
2. The DBE applies QoS policing and forwards the packet to endpoint B without changing the destination address to a DBE local media address (as is done in Single NAPT). Under No NAPT processing, the DBE does not rewrite either source or destination IP addresses and ports.
3. Side B sends a packet from IP address and port 2001:11::11/28988 to originating source address and port 2001:10::10/17002. The DBE intercepts this packet and matches it to the side B flow.
4. The DBE applies QoS policing and forwards the packet to user side A without rewriting either source or destination IP addresses and ports.

IPv6 Single NAPT for Signaling

Support of IPv6 signaling flows requires Single NAPT.

The DBE is able to translate IP addresses and port numbers in both directions of a flow. However, Single NAPT means only one IP address and port is translated. In Single NAPT processing, the flow on one side of the pinhole is programmed with a local address and port that do not belong to the SBC. Instead, that local address and port of the flow are specified by the MGC to match the remote address and port on the other side of the pinhole. Thus, incoming traffic (downstream traffic of SIP server to access side) is addressed directly to the remote endpoint and the SIP server details are hidden from subscribers. Network topology must be used to route the downstream traffic through the DBE. In one sense, Single NAPT provides one-way topology hiding.

SBC rewrites destination IP address and port for packets received from the user. SBC does not rewrite source IP address and port of packets received from the user (they are unchanged from the IP address and port of the source endpoint). Correspondingly, SBC rewrites the source IP address and port of packets received from the MGC, but not the destination IP address or port.

Figure 13-2 illustrates a Single NAPT signaling flow through the DBE between user side A and user side B.

Figure 13-2 Single NAPT Signaling Flow

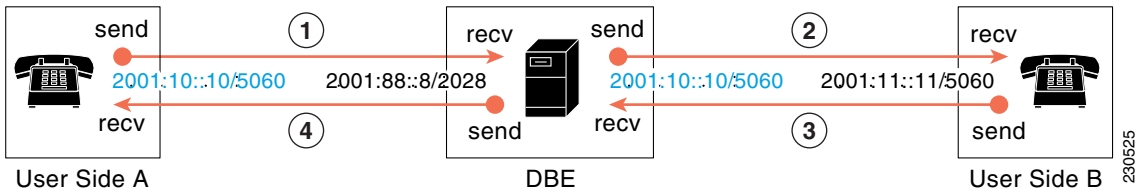

1. User side A sends a packet from IP address and port 2001:10::10/5060 to the DBE's local media address and port 2001:88::8/2028 for this pinhole. User side A only knows the DBE's local address and port 2001:88::8/2028. The source IP address is within the specified gm/rsam, so the DBE matches this packet to the flow.
2. The DBE applies QoS policing and forwards the packet to the MGC (user side B) without rewriting the source IP address and port. Under Single NAPT processing, the DBE changes the destination address and port to 2001:11::11/5060 on the MGC (side B) by replacing 2001:88::8/2028 with side B's address and port from the remote descriptor on side B. The MGC (side B) does not know about the 2001:88::8/2028 address and port on the DBE. After the DBE performs latching, the source address and port from side A becomes, in effect, the destination address and port in step 3 and step 4 for side B.
3. The MGC (side B) sends a packet to user side A with the destination address and port 2001:10::10/5060 copied from the source IP address and port of the packet it just received—that is, the address and port of side A. The DBE has intercepted the packet and matched it to the side B flow.
4. The DBE applies QoS policing and forwards the packet to side A without rewriting the destination IP address and port 2001:10::10/5060. However, under Single NAPT processing, the DBE rewrites the source IP address and port 2001:11::11/5060 to be 2001:88::8/2028, which is the local address and port of the side A flow.

Restrictions for DBE IPv6 Single NAPT for Signaling

The following are restrictions pertaining to DBE support for IPv6 pinholes:

- DBE does not support IPv6 for control communications with the SBE. H.248 communication with the controlling SBE is over IPv4 only.
- DBE does not support IPv6 addresses that are not global unicast addresses.
- DBE does not support IPv6 addresses that do not use the default zone.

- DBE does not use the IPv6 Flow Label to classify packets. It continues to use the transport protocol type (UDP/TCP) and local and remote ports, as with IPv4. Outgoing packets originating from the DBE, such as DTMF packets, have a Flow Label of 0.
- DBE does not support forwarding between IPv4 and IPv6 endpoints. In particular, 6 to 4 addresses (prefixed with 2002::/16) are treated as global unicast native IPv6 addresses.
- Intra-Site Automatic Tunnel Addressing Protocol (ISATAP) link-local addresses are not supported.

Related Commands

The commands related to IPv6 Single NAPT for Signaling include:

- The **ipv6 address (session border controller)** command sets or creates the IPv6 address prefix on an SBC interface.
- The **media-address ipv6** command adds an IPv6 address to the set of addresses that can be used by the DBE as a local media address.
- The **media-address pool ipv6** command creates a pool of sequential IPv6 media addresses that can be used by the DBE as local media addresses.
- The **port-range (ipv6)** command creates a port range associated with a single IPv6 media address or a pool of IPv6 media addresses. IPv6 addresses must be configured with the **signaling** keyword. The **any**, **voice**, **video,** and **fax** keywords supported in the IPv4 **port-range** command are not supported in IPv6.
- The **ipv6** {*ipv6-address*} keyword is added to the **debug sbc filter** command.
- The **ipv6** {*ipv6-address*} keyword is added to the **show sbc dbe media-flow-stats** and **show sbc dbe signaling-flow-stats** commands.

No NAPT Pinholes

No NAPT pinholes can form coupled pairs only under the following circumstances:

- Both pinholes are No NAPT.
- Each “internal termination” has local and remote addresses that are identical to those of the external termination on the associated pinhole.

Note The two terminations between which media loops back are called the “internal terminations” of their respective pinholes. Only external terminations directly receive packets from the network.

- Any remote source address masks (rsams) are duplicated. For example, if a termination with remote address A in one pinhole has an rsam of 1111:2222:3333:4444::/48, then the termination with remote address A in the other pinhole also has an rsam of 1111:2222:3333:4444::/48.

Restrictions for DBE No NAPT Pinholes

The following are DBE restrictions pertaining to the No NAPT Pinholes feature:

- The DBE chooses the internal terminations as follows:
  - The first specified termination is chosen to be internal.

  - The other termination is chosen accordingly from the other pinhole. If the termination with remote address A on one pinhole is internal, then the termination with local address A on the other pinhole is also internal.
  - The DBE does not support choosing internal terminations based on termination names.
- For No NAPT coupled pairs, any Network Address Translation (NAT) latching requests are duplicated. For example, if a termination with remote address A in one pinhole requests NAT latching, then the termination with remote address A in the other pinhole must also request NAT latching. The “request NAT latching” can be done using the ipnapt/latch H.248 signal.
- A hairpin of two pinholes in which both external terminations are provisioned with the NAT latching instruction cannot latch and cannot forward media. No NAPT pinholes are not allowed to (re)latch to the remote addresses on both sides.
- IPv6 hairpinning are supported on UDP and TCP.
- Coupling of Single NAPT pinholes is not supported.