Chapter 6: Troubleshooting Addressing Services

CCNP TSHOOT: Maintaining and Troubleshooting IP Networks

Cisco | Networking Academy®
Mind Wide Open™

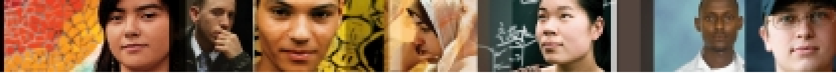

NAT/PAT Operation Review

- NAT was designed for IPv4 address conservation.
- Usually operates at the border of a network and translates private source address of exiting IP packets to public addresses.
- The packet header information and the corresponding translated IP address are kept in a NAT table.
- NAT does the reverse for the destination address of the responding IP packets based on the content of the NAT table.
- In VPN connectivity situations, NAT can overcome the connectivity issues that arise by translating overlapping address spaces to non-overlapping addresses.

NAT/PAT Operation Review – Cont.

- NAT operates at the border of the network and generally translates the source IP address of exiting packets and the destination IP address of incoming IP packets.

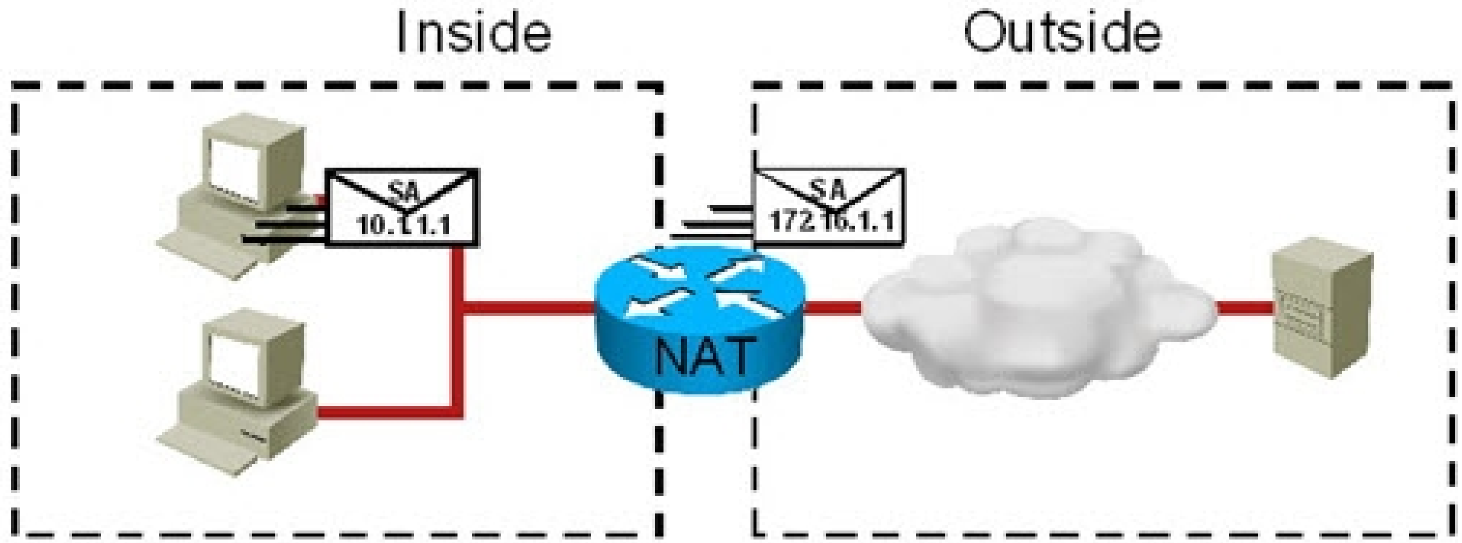

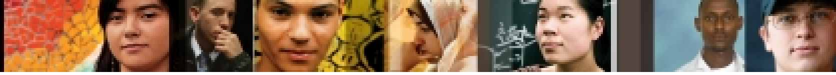

NAT/PAT Operation Review – Cont.

Three main categories of NAT:

- Static NAT:
  - Inside local and inside global addresses are mapped one to one.
  - Used when an inside device must be accessible from the outside network (such as a web server).
  - In troubleshooting, IP address changes might affect an existing static configuration.
- Dynamic NAT:
  - Translates addresses utilizing the same basic technology as static NAT.
  - Local addresses are translated to a group or pool of global addresses.
  - Issues can be related to the size of the global pool (requires one-to-one translation).
  - Some inside hosts may not obtain a valid global address, causing connectivity problems.
  - Subject to management, tracking, and audit issues due to the dynamic nature of the translation.
- NAT overloading:
  - Special type of dynamic NAT in which addresses are translated in a many-to-many fashion.
  - Also known as PAT, or Port Address Translation, because global addresses can be reused.
  - The differentiator for multiple inside local addresses sharing the same global address is a port number.
  - NAT overloading suffers from some application support issues.

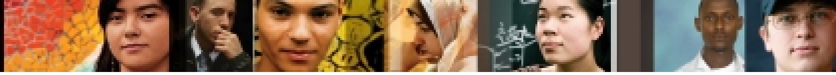

NAT/PAT Operation Review – Cont.

Advantages and disadvantages of implementing NAT:

Advantages	Disadvantages
Conserves registered addresses	Translation introduces processing delays
Hides the actual address of internal hosts and services	Loss of end-to-end IP reachability
Increases flexibility when connecting to Internet	Certain applications will not function with NAT enabled
Eliminates address renumbering as the network changes from one ISP to another	Considerations are needed when working with VPNs

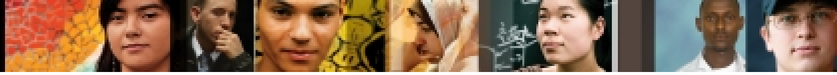

NAT/PAT Operation Review – Cont.

NAT-Sensitive protocols and their NAT-sensitive behavior:

Protocol	Behavior
IPsec	NAT changes certain IP header fields such as the IP address and the IP header checksum. This can conflict with IPsec integrity.
ICMP	Many ICMP packets, such as Destination Unreachable, carry embedded IP header information inside the ICMP message payload, not matching IP packet’s translated address.
Session Initiation Protocol (SIP)	Protocols such as SIP negotiate address and ports numbers at the application layer, which can become invalid through a NAT device.

NAT/PAT Operation Review – Cont.

Order of operations with reference to NAT for incoming and outgoing traffic:

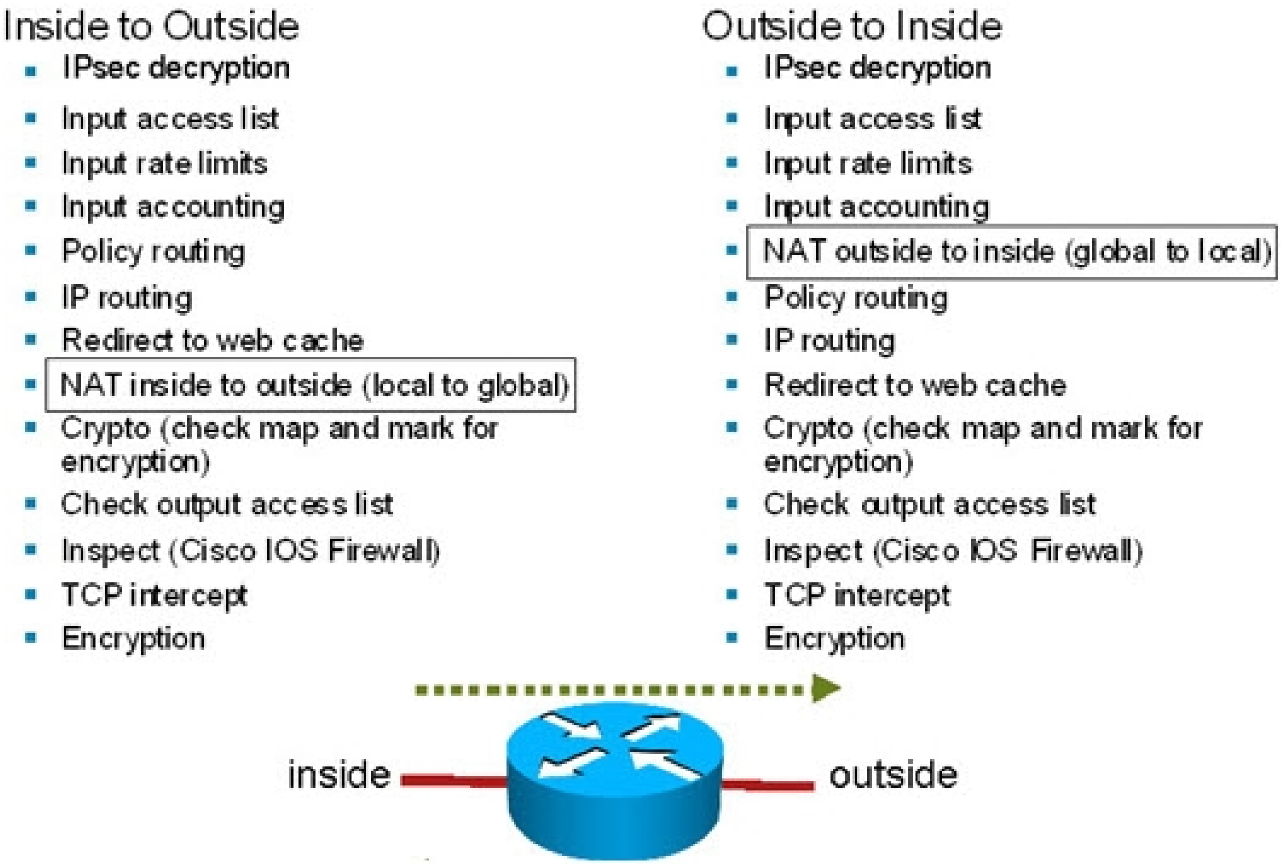

Inside to Outside

- IPsec decryption
- Input access list
- Input rate limits
- Input accounting
- Policy routing
- IP routing
- Redirect to web cache
- NAT inside to outside (local to global)
- Crypto (check map and mark for encryption)
- Check output access list
- Inspect (Cisco IOS Firewall)
- TCP intercept
- Encryption

Outside to Inside

- IPsec decryption
- Input access list
- Input rate limits
- Input accounting
- NAT outside to inside (global to local)
- Policy routing
- IP routing
- Redirect to web cache
- Crypto (check map and mark for encryption)
- Check output access list
- Inspect (Cisco IOS Firewall)
- TCP intercept
- Encryption

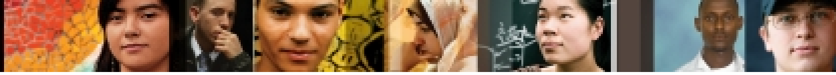

Troubleshooting Common NAT/PAT Issues

Some important NAT issues and considerations to keep in mind are:

- Diagrams for the NAT configuration are helpful and should be a standard practice.
- Do not start configuring without a diagram that shows or explains each item involved.
- ACLs are used to tell the NAT device "what source IP addresses are to be translated“
- IP NAT pools are used to specify the available global addresses.
- Marking the IP NAT inside interfaces and the IP NAT outside interfaces correctly is important.
- NAT packets still have to obey routing protocols and reachability rules.
- Make sure that every router knows how to reach the desired destinations.
- Make sure the public addresses to which addresses translate are advertised to the outside neighbors and autonomous systems.

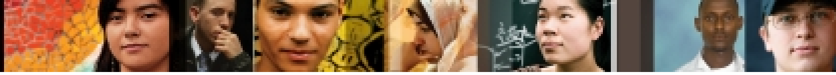

Troubleshooting NAT/PAT Issues – Cont.

- **`clear ip nat translation:`**
  - Removes NAT entries from the NAT table.
  - Specific entries can cleared with additional parameters.
  - Clearing all translations can cause disruption until new translations are re-created.
- **`show ip nat translations:`**
  - Displays all the translations (static and dynamic) that are currently installed and active on the router.
- **`show ip nat statistics:`**
  - Displays NAT statistics such as number of translations (static, dynamic, extended), number of expired translations, number of hits (match), number of misses (no match).

Troubleshooting NAT/PAT Issues – Cont.

- **`debug ip nat:`**
  - Displays information about each packet that the router translates.
- **`debug ip nat detailed:`**
  - Generates a description of each packet considered for translation.
  - Also displays information about certain errors or exception conditions, such as the failure to allocate a global address.
- **`debug ip packet [`*`access-list`*`]:`**
  - Displays general IP debugging information and IP security option (IPSO) security transactions.
  - If a communication session is closing when it should not be, an end-to-end connection problem can be the cause.
  - Useful for analyzing messages traveling between the local and remote hosts.
  - Captures packets that are process switched including received, generated, and forwarded packets.
  - IP packets that are switched in the fast path are not captured.
  - The *`access-list`* option allows you to narrow down the scope of debugging.

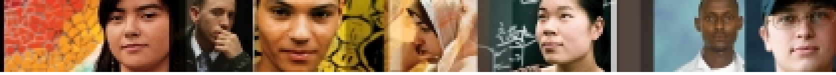

Troubleshooting NAT/PAT Issues – Cont.

Limiting debug output with the **`debug condition`** command:

- **`debug condition interface`** *`interface`*:
  - Called conditionally triggered debugging.
  - Generates debugging messages for packets entering or leaving on only the specified interface.
  - First define the condition with the **`debug condition`** command. For example, define a condition of **`interface serial 0/0`**.
  - The condition remains defined and applied until it is removed.
  - Check the active debug conditions using the **`show debug condition`** command.

NAT/PAT Troubleshooting Example 1: Routing Issue

- Router R1 can ping R4, but router R1 cannot ping R3.
- There are no routing protocols running in any of the routers.
- R1 uses R2 as its gateway of last resort.
- The objective is to restore end-to-end connectivity from R1 to all destinations.

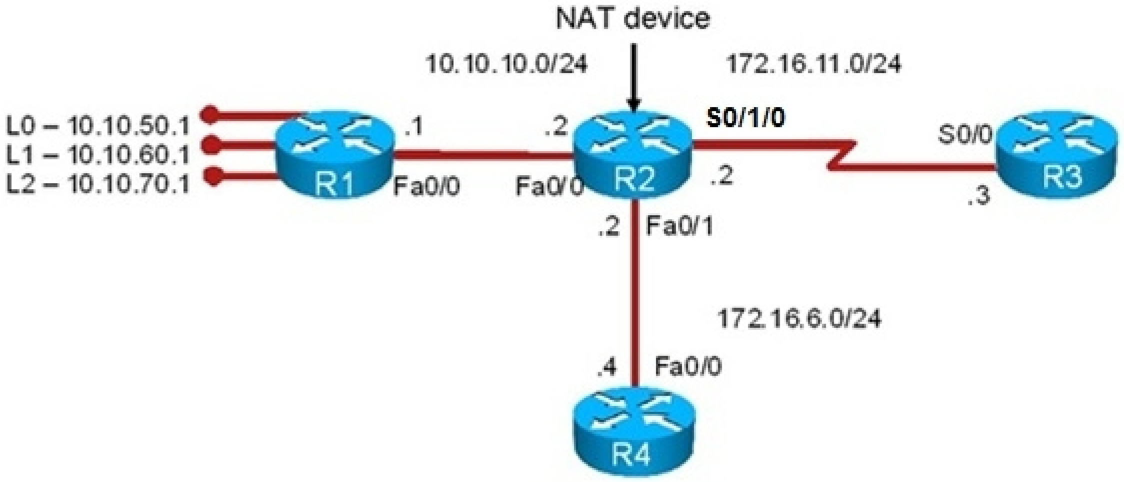

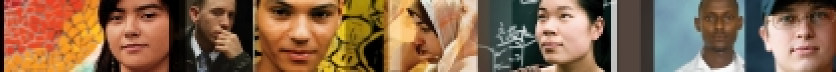

NAT/PAT Troubleshooting Example 1 – Cont.

```
R2# sh ip nat statistics
Total active translations: 1 (1 static, 0 dynamic, 0 extended)
Outside interfaces:
  FastEthernet0/1, Serial0/1/0

Inside interfaces:
  FastEthernet0/0

Hits: 39  Misses: 6
CEF Translated packets: 45, CEF Punted packets: 49
Expired translations: 6
Dynamic mappings:
-- Inside Source
[Id: 1] access-list 10 pool NAT_OUT refcount 0
 pool NAT_OUT: netmask 255.255.255.0
        start 172.16.6.129 end 172.16.6.240
        type generic, total addresses 112, allocated 0 (0%), misses 0
Appl doors: 0
Normal doors: 0
Queued Packets: 0
```

NAT/PAT Troubleshooting Example 1 – Cont.

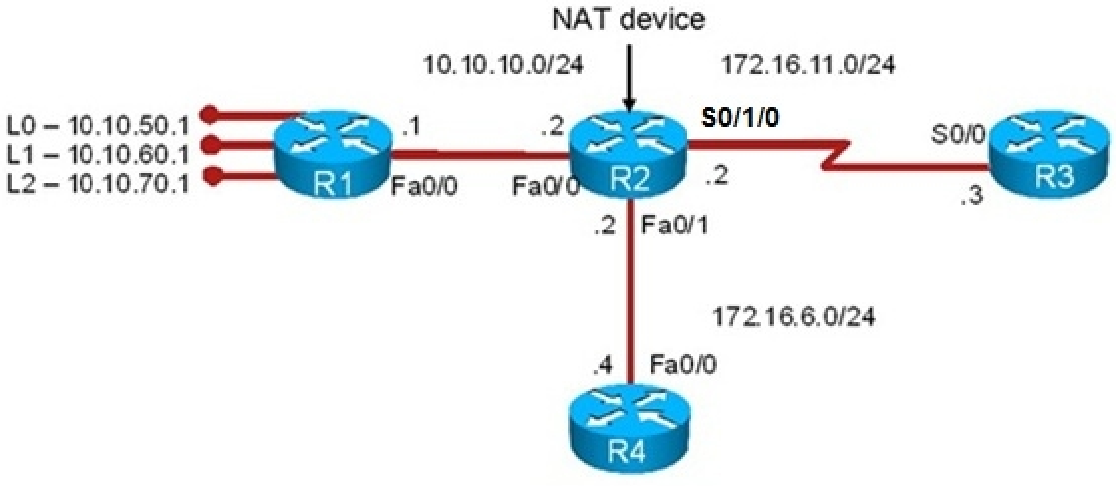

```
R2# sh ip nat translations
Pro  Inside global  Inside local     Outside local     Outside global
---  172.16.6.1     10.10.10.1         ---               ---
```

NAT/PAT Troubleshooting Example 1 – Cont.

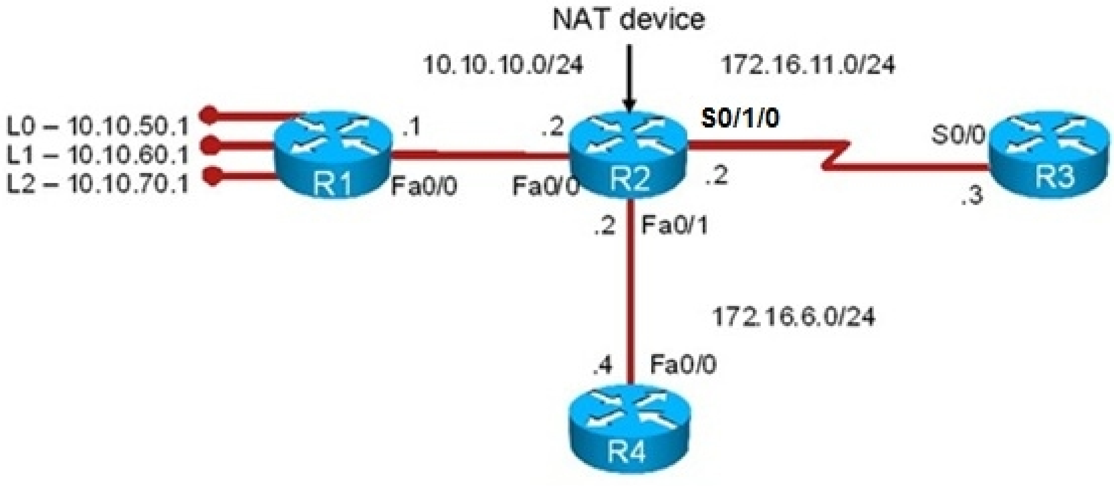

```
R3# debug ip icmp
ICMP packet debugging is on

R1# ping 172.16.11.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R3#
*Aug 23 13:54:00.556:  ICMP:  echo reply sent, src 172.16.11.3,  dst 172.16.6.1
*Aug 23 13:54:02.552:  ICMP:  echo reply sent, src 172.16.11.3,  dst 172.16.6.1
*Aug 23 13:54:04.552:  ICMP:  echo reply sent, src 172.16.11.3,  dst 172.16.6.1
*Aug 23 13:54:06.552:  ICMP:  echo reply sent, src 172.16.11.3,  dst 172.16.6.1
*Aug 23 13:54:07.552:  ICMP:  echo reply sent, src 172.16.11.3,  dst 172.16.6.1
```

NAT/PAT Troubleshooting Example: – Cont.

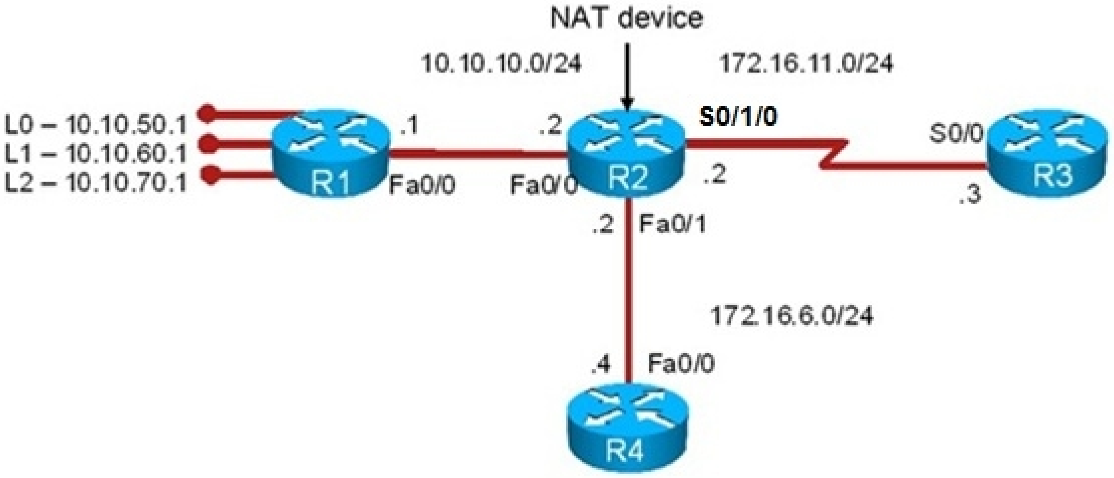

```
R3# show ip route 172.16.6.0 255.255.255.0
% Subnet not in table

R3# configure terminal
R3(config)# ip route 172.16.6.0 255.255.255.0 172.16.11.2
R3(config)# exit

R1# ping 172.16.11.3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/2/4 ms
R1#
```

NAT/PAT Troubleshooting Example 2: Incorrect Access List

- Administrators are unable to use SSH from the 10.10.10.0/24 network to routers R3 or R4, but they can connect to R1’s loopbacks.
- The risk management team recently performed an upgrade to router and firewall security policies.
- The routing protocol is single-area OSPF.
- Goal to restore end-to-end connectivity and make sure SSH is operational to support management processes.

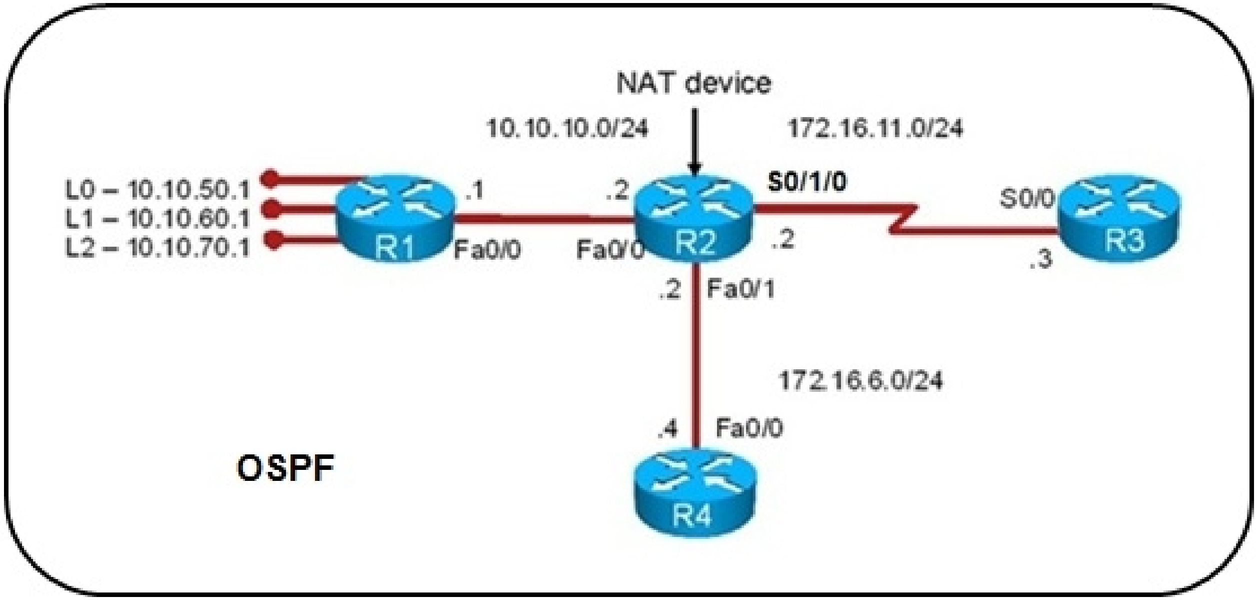

NAT/PAT Troubleshooting Example 2 – Cont.

- Extended ping and SSH results from R1 to R3

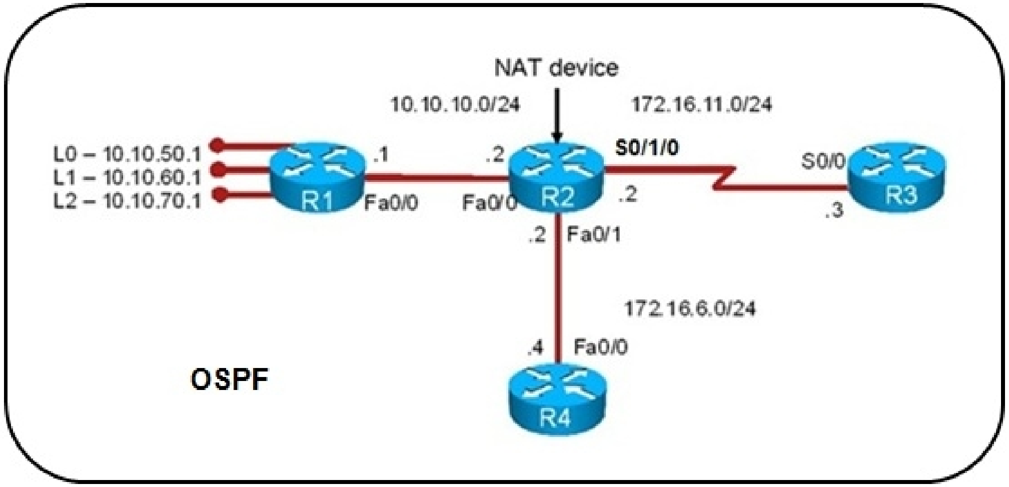

```
R1# ping 172.16.11.3 source 10.10.50.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.3, timeout is 2 seconds:
Packet sent with a source address of 10.10.50.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1# ping 172.16.11.3 source 10.10.10.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.3, timeout is 2 seconds:
Packet sent with a source address of 10.10.10.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

R1# ssh -l user 172.16.11.3
% Connection refused by remote host
```

NAT/PAT Troubleshooting Example 2 – Cont.

Using **`debug ip tcp transactions`** while attempting SSH

```
R1# debug ip tcp transactions
TCP special event debugging is on
R1# ssh -l user 172.16.11.3
% Connection refused by remote host
R1#
*Aug 23 14:59:42.636: TCP: Random local port generated 42115, network 1
*Aug 23 14:59:42.636: TCB63BF854C created
*Aug 23 14:59:42.636: TCB63BF854C bound to UNKNOWN.42115
*Aug 23 14:59:42.636: TCB63BF854C setting property TCP_TOS (11) 62AAF6D55
*Aug 23 14:59:42.636: Reserved port 42115 in Transport Port Agent for TCP IP type 1
*Aug 23 14:59:42.640: TCP: sending SYN, seq 1491927624, ack 0
*Aug 23 14:59:42.640: TCP0: Connection to 172.16.11.3:22, advertising MSS 536
*Aug 23 14:59:42.640: TCP0: state was CLOSED -> SYNSENT [42115 ->
172.16.11.3(22)]
*Aug 23 14:59:42.640: TCP0: state was SYNSENT -> CLOSED [42115 ->
172.16.11.3(22)]
*Aug 23 14:59:42.640: Released port 42115 in Transport Port Agent for TCP IP
type 1 delay 240000
*Aug 23 14:59:42.640: TCP0: bad seg from 172.16.11.3 — closing connection:
port 42115 seq 0 ack 1491927625 rcvnxt 0 rcvwnd 0 len 0
*Aug 23 14:59:42.640: TCP0: connection closed - remote sent RST
*Aug 23 14:59:42.640: TCB 0x63BF854C destroyed
```

NAT/PAT Troubleshooting Example 2 – Cont.

Checking the access list applied to the serial interface on R3

```
R3# sh ip int s0/0
Serial 0/0 is up, line protocol is up
 Internet address is 172.16.11.3/24
 Broadcast address is 255.255.255.255
 Address determined by nonvolatile memory
 MTU is 1500 bytes
 Helper address is not set
 Directed broadcat forwarding is disabled
 Multicast reserved groups joined: 224.0.0.5
 Outgoing access list is not set
 Inbound access list is FIREWALL-INBOUND
 Proxy ARP is enabled
 Local Proxy ARP is disabled
 Security level is default
 Split horizon is enabled
 ICMP redirects are always sent
 ICMP unreachables are always sent
 ICMP mask replies are never sent
 IP fast switching is enabled
 IP fast switching on the same interface is enabled
 IP Flow switching is disabled
 IP CEF switching is enabled
 IP CEF Feature Fast switching turbo vector
 IP multicast fast switching is enabled

R3# sh access-lists
Standard IP access list 11
  10 permit any
Extended IP access list FIREWALL-INBOUND
  10 permit tcp any host 172.16.11.3 eq www
  20 permit tcp any host 172.16.11.3 eq telent
  30 permit tcp any host 172.16.11.3 eq 22
  40 permit tcp any host 172.16.11.3 eq ftp
  50 permit tcp any host 172.16.11.3 eq ftp-data
  60 permit ospf any any (20 matches)
  70 deny ip any any (1 match)
```

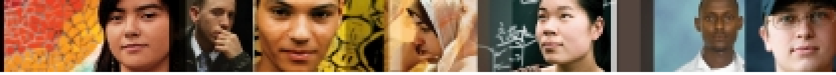

NAT/PAT Troubleshooting Example 2 – Cont.

Using **debug ip packet** while attempting SSH

```
R3# debug ip packet
IP packet debugging is on
R3#
R1# ssh -l user 172.16.11.3
% Connection refused by remote host
R1#
R3#
*Aug 23 16:32:42.711: IP: s=172.16.11.2 (Serial0/1/0), d=224.0.0.5,
len 80, rcvd 0
*Aug 23 16:32:49.883: %SEC-6-IPACCESSLOGP: list FIREWALL-INBOUND
denied tcp 10.10.10.1(29832) -> 172.16.11.3(2222), 1 packet
*Aug 23 16:32:49.883: IP: s=10.10.10.1 (Serial0/1/0), d-172.16.11.3,
len 44, access denied
*Aug 23 16:32:49.883: IP: tableid=0, s-172.16.11.3 (local),
d=10.10.10.1 (Serial0/1/0), routed via FIB
*Aug 23 16:32:49.883: IP: s=172.16.11.3 (local), d=10.10.10.1
(Serial0/1/0), len 56, sending
*Aug 23 16:32:50.067: IP: s=172.16.11.3 (local), d=224.0.0.5
(Serial0/1/0), len 80, sending broad/multicast
```

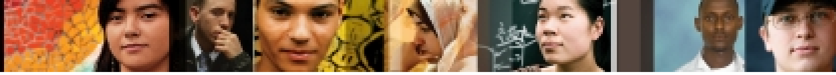

NAT/PAT Troubleshooting Example 2 – Cont.

Using **debug ip nat** while attempting SSH

```
R2# debug ip nat
IP NAT debugging is on
R2#
R2#
R2#
R2#
*Aug 23 16:28:31.731: NAT*: TCP s=555 55587, d=22->2222

R1# ssh -l user 172.16.11.3
% Destination unreachable; gateway or host down
R1#

R2# sh ip nat translations
Pro Inside global        Inside local        Outside local      Outside global
tcp ---                  ---                 172.16.11.3:22     172.16.11.3:2222
tcp 10.10.10.1:29832     10.10.10.1:29832    172.16.11.3:22     172.16.11.3:2222
tcp 10.10.10.1:43907     10.10.10.1:43907    172.16.11.3:22     172.16.11.3:2222
tcp 10.10.10.1:55587     10.10.10.1:55587    172.16.11.3:22     172.16.11.3:2222
tcp 10.10.10.1:60089     10.10.10.1:60089    172.16.11.3:22     172.16.11.3:2222
tcp 10.10.10.1:62936     10.10.10.1:62936    172.16.11.3:22     172.16.11.3:2222
```

NAT/PAT Troubleshooting Example 2 – Cont.

Correcting the ACL on R3 to allow SSH with a custom port.

```
R3# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# ip access-list exten FIREWALL-INBOUND
R3(config-ext-nacl)# permit tcp any host 172.16.11.3 eq 2222
R3(config-ext-nacl)# end
R3#


R1# ssh -l user 172.16.11.3
Password:
*Aug 23 16:30:42.604: TCP: Random local port generated 43884, network 1
*Aug 23 16:30:26.604: TCB63BF854C created
*Aug 23 16:30:26.604: TCB63BF854C bound to UNKNOWN.43884
*Aug 23 16:30:26.604: TCB63BF854C setting property TCP_TOS (11) 62AF6D55
*Aug 23 16:30:26.604: Reserved port 43884 in Transport Port Agent for TCP IP type 1
*Aug 23 16:30:26.604: TCP: sending SYN, seq 1505095793, ack 0
*Aug 23 16:30:26.604: TCP0: Connection to 172.16.11.3:22, advertising MSS 536
*Aug 23 16:30:26.608: TCP0: state was CLOSED -> SYNSENT [43884 ->
172.16.11.3(22)]
*Aug 23 16:30:26.608: TCP0: state was SYNSENT -> ESTAB [43884 ->
172.16.11.3(22)]
*Aug 23 16:30:26.608: TCP: tcb 63BF854C connection to 172.16.11.3:22, peer MSS
536, MSS is 536
*Aug 23 16:30:26.608: TCB63BF854C connected to 172.16.11.3.22
```

Reviewing DHCP Operation

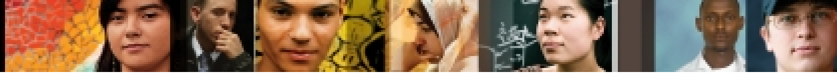

DHCP Overview

- DHCP is a client-server protocol. The DHCP client acquires an IP address, subnet mask, and default gateway, from a DHCP server.
- The DHCP server is typically centrally located and operated by the network administrator.
- The DHCP client will broadcast to find a DHCP server to obtain its IP configuration.
- Most enterprise networks are divided into VLANs and routers or multilayer switches route between the VLANs.
- Since routers block broadcasts, a DHCP server would be needed on each subnet.
- To address this issue, use the DHCP Relay Agent, the **ip helper-address** *server-address* command.
- The router will forward client DHCP messages to the configured server-address. When the server sends its reply back to the router interface, the router interface forwards it to the client subnet.

DHCP Overview – Cont.

- DHCP is used for automatic provisioning of IP parameters.
- Increased mobility and usage of laptop computers have made DHCP even more popular.
- The DHCP server offers more than just the IP address, subnet mask, and the default gateway.
- Cisco IP Phones require a TFTP server address to download their configuration files and become active in the network.
- IP phones obtain the TFTP server's IP address from the DHCP server through DHCP options and extensions.
- These options allow the protocol to expand the number and nature of parameters that can be delivered to hosts, including vendor specific parameters.

DHCP Overview: Client–Server Communication

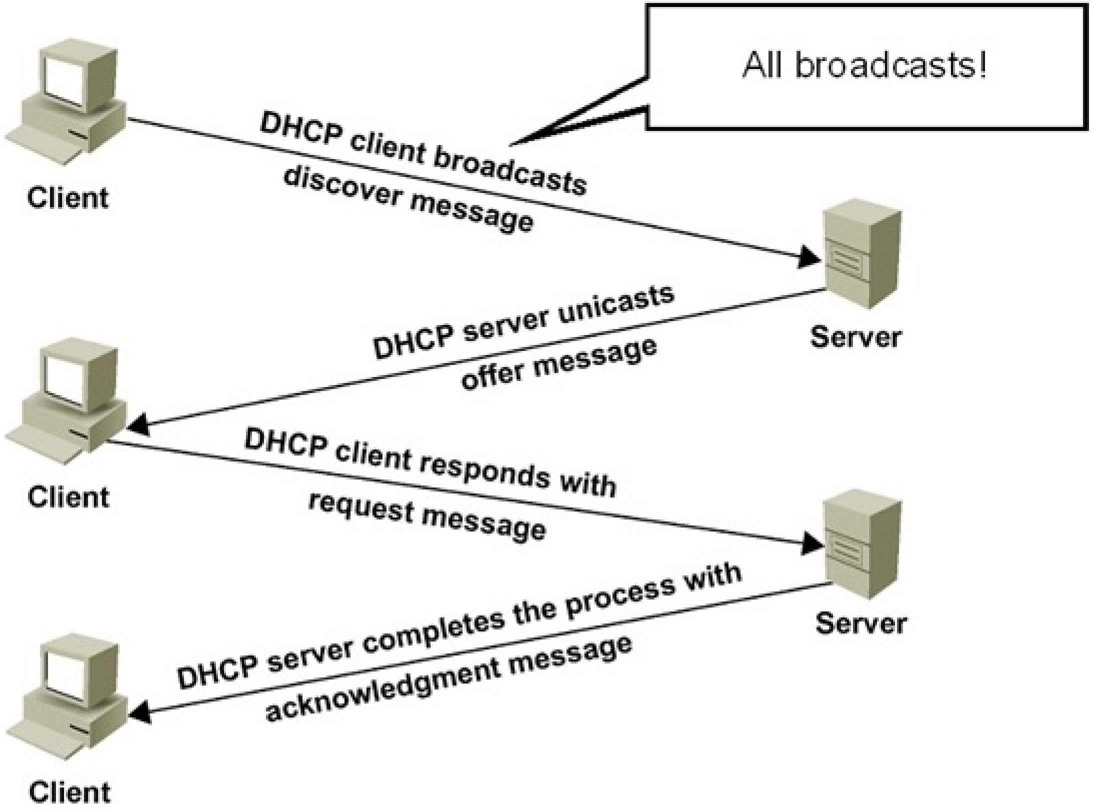
All broadcasts!
Client
DHCP client broadcasts discover message
Server
DHCP server unicasts offer message
Client
DHCP client responds with request message
Server
DHCP server completes the process with acknowledgment message
Client

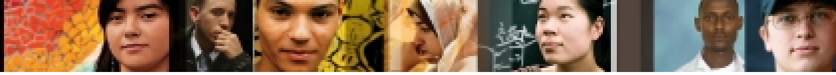

DHCP Overview: Client–Server Communication – DHCP Packets (Message Types)

Packet Type	Description
DHCP discover	Client looking for available DHCP servers. It is a UDP broadcast (source port is 68, and the destination port is 67).
DHCP offer	This is the server's response to the client's discover message. This is also a UDP broadcast (source port is 67, and the destination port is 68).
DHCP request	This is client's response to one specific DHCP offer.
DHCP decline	Client-to-server communication, indicating that the IP address is already in use.
DHCP ack	Server-to-client communication. This is the server's response to a client request. This message includes all configuration parameters.
DHCP nack	Server-to-client communication. This is the server's negative response to a client's request, indicating the original offer is no longer available.
DHCP release	Client-to-server communication. The client relinquishes its IP address and other parameters.
DHCP inform	Client-to-server communication. Using this message, the client asks for local configuration parameters such as DNS server's IP address, but it has its IP address externally configured.

Common DHCP Troubleshooting Issues: Three DHCP Roles a Router May Take

Router acting as DHCP server

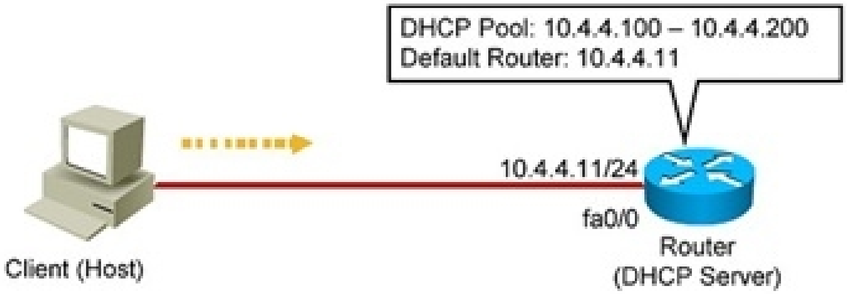

Router acting as DHCP client

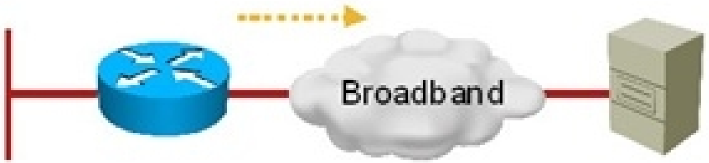

Router brokering DHCP transactions

(DHCP relay agent)

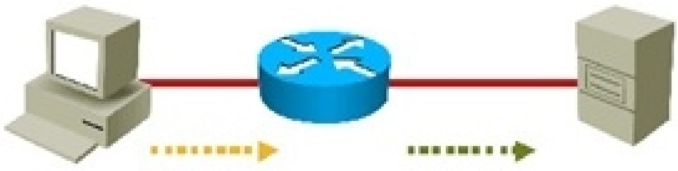

DHCP Troubleshooting Issues – Cont.

- Configuration issues can result in many symptoms:
  - Clients not obtaining IP information from the server
  - Client requests not reaching the server across a DHCP relay agent
  - Clients failing to obtain DHCP options and extensions
- Address pool issues:
  - Poor capacity planning and security issues might result in DHCP scope exhaustion.
  - When using static and dynamic IP address assignments, an IP address that is already in use can be granted.
  - Multiple DHCP servers, or even rogue DHCP servers can result in duplicate IP addresses
  - assigned to hosts.
- Management issues:
  - Due to the “pull” nature of DHCP.
  - There are no provisions in the protocol to allow the DHCP server to push configuration parameters or control messages to DHCP clients.
  - A good example, with critical implications in IP address renumbering, is that IP addresses must be renewed from the client side. There is no server-side, push-type renewal process.
  - This means that during renumbering, all clients would need to reboot or manually renew their IP addresses. Otherwise, you need to wait until the clients leases expire, which might not be a viable option.

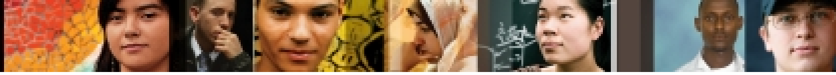

DHCP Troubleshooting Issues: DHCP Relay Agent

- The Cisco IOS command that makes a router a DHCP relay agent **`is ip helper-address`**.
- If the DHCP server's IP address changes, all interfaces of all routers must be reconfigured with the new IP helper-address (DHCP server's new IP address).
- Enabling a router interface with the **`ip helper-address`** command allows the interface forward UDP broadcasts for six protocols (not just DHCP) to the IP address configured.
  - TFTP (port 69)
  - DNS (port 53)
  - Time Service (port 37)
  - NetBIOS Name Service and Datagram Service (ports 137 and 138)
  - TACACS (port 49)
  - DHCP/BOOTP Client and Server (ports 67 and 68)
- If these other protocols do not require this service, forwarding their requests must be disabled manually on all routers using the Cisco IOS **`no ip forward-protocol udp`** *`port-number`* command in global configuration mode.

DHCP Troubleshooting Issues: DHCP Parameters and Options

DHCP Option	Code	Description
Subnet Mask	1	Specifies the subnet mask for the client to use (as per RFC 950)
Router	3	The list of routers the client can use (usually, in order of preference)
Domain Name Server	6	The list of DNS servers the client can use (usually, in order of preference)
ARP Cache Timeout	35	Specifies the timeout (seconds) for ARP cache entries
IP Address Lease Time	51	Specifies the period over which the IP address is leased (it must be renewed)
Relay Agent Information	82	Information about the port from which the DHCP request originates
TFTP Server IP Address	150	Typically used by devices such as IP Phones to download their configuration files

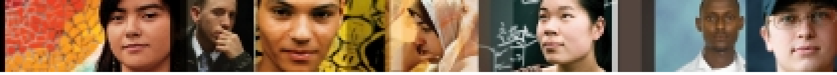

DHCP Troubleshooting Issues

- Automatic addressing is accomplished through DHCP.
- Security is accomplished through DHCP snooping.
- Some specific issues related to DHCP snooping:
    - Improper configuration of the DHCP snooping trust boundaries
    - Failure to configure DHCP snooping on certain VLANs
    - Improper configuration of the DHCP snooping rate limits
    - Performance degradation
- Poor planning of DHCP snooping can result in DHCP transactions being blocked or affected.

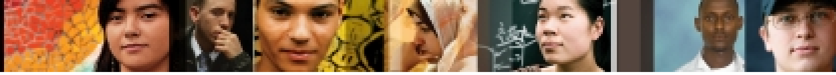

DHCP Troubleshooting Issues – Cont.

DHCP troubleshooting questions to ask:

- Where are the DHCP servers and clients located?
- Are DHCP relay agents configured?
- What are the DHCP pool sizes? Are they sufficient?
- Are there any DHCP option compatibility issues?
- Are there any ACLs or firewalls filtering UDP port 67 or 68?
- Are there any active DHCP DoS attacks?
- Is forwarding disabled on the router acting as DHCP Relay Agent for any UDP ports (using the Cisco IOS **`no ip forward-protocol udp`** *`port`* command)?
- Is the **`ip helper-address`** command applied to correct router interfaces?
- Is DHCP snooping configured?

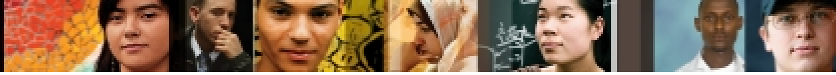

DHCP Troubleshooting Commands

- `show ip dhcp server statistics`
  - Displays counts for server statistics and messages sent and received for an IOS-based DHCP server.
- `show ip dhcp binding`
  - Displays DHCP binding information for IP address assignment and subnet allocation.
- `show ip dhcp conflict`
  - Displays address conflicts found by a Cisco IOS DHCP server when addresses are offered to the client.
- `show ip dhcp pool` *name*
  - Displays the subnets allocated and the current utilization level for the pool or all the pools if the name argument is not used.

DHCP Troubleshooting Commands

- .**`show ip dhcp database`**
  - Displays server database agent information:
    - **URL:** Specifies the remote file used to store automatic DHCP bindings
    - **Read/written:** The last date and time bindings were read/written from the file server
    - **Status:** Indication of whether the last read or write of host bindings was successful
    - **Delay:** The amount of time (in seconds) to wait before updating the database
    - **Timeout:** The amount of time (in seconds) before the file transfer is aborted
    - **Failures/Successes:** The number of failed/successful file transfers

DHCP Troubleshooting Commands – Cont.

- **`debug ip udp`**:
  - Displays UDP packets sent and received.
  - Can use considerable CPU cycles on the device.
- **`debug ip dhcp server [packets | events]:`**
  - Enables DHCP server debugging.
  - The **`events`** option reports server events such as address assignments and database updates.
  - The **`packets`** option decodes DHCP receptions and transmissions.
- **`clear ip dhcp binding {* |`** *`address`***`}:`**
  - Deletes an address binding from the DHCP server database.
  - The address denotes the IP address of the client.
  - If the asterisk (*) character is used as the address parameter, DHCP clears all automatic bindings.
- **`clear ip dhcp conflict {* |`** *`address`***`}:`**
  - Clears an address conflict for a specific entry with the *`address`* option.
  - Clears all address conflicts with the asterisk (**`*`**) option.

DHCP Troubleshooting Example 1: Problems After a Security Audit

- Router R1 provides DHCP services to clients in the 10.1.1.0 subnet.
- The DHCP clients are R2 and R3.
- A security audit has been recently performed in router R1.
- It is reported that R1 is no longer providing reliable DHCP services.
- The clients are unable to renew their IP addresses.

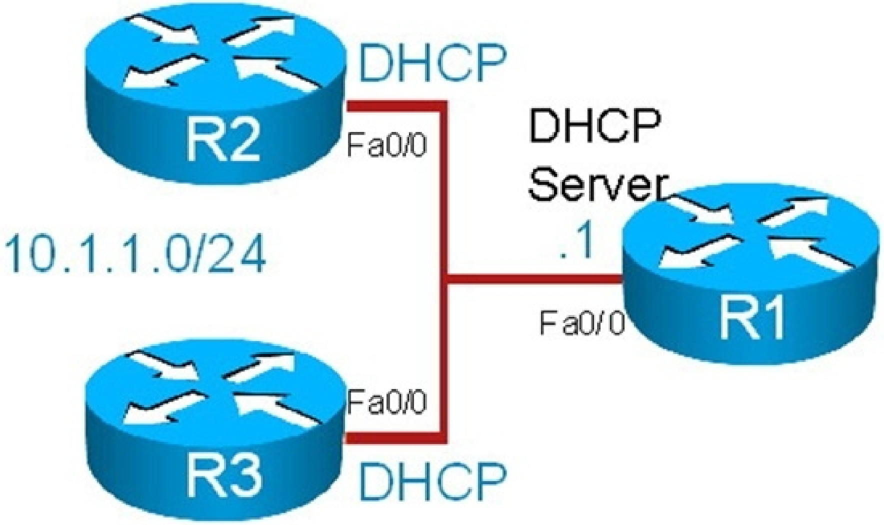

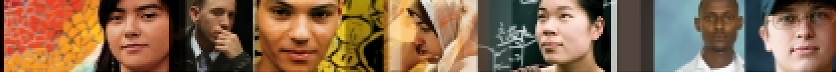

DHCP Troubleshooting Example 1 – Cont.

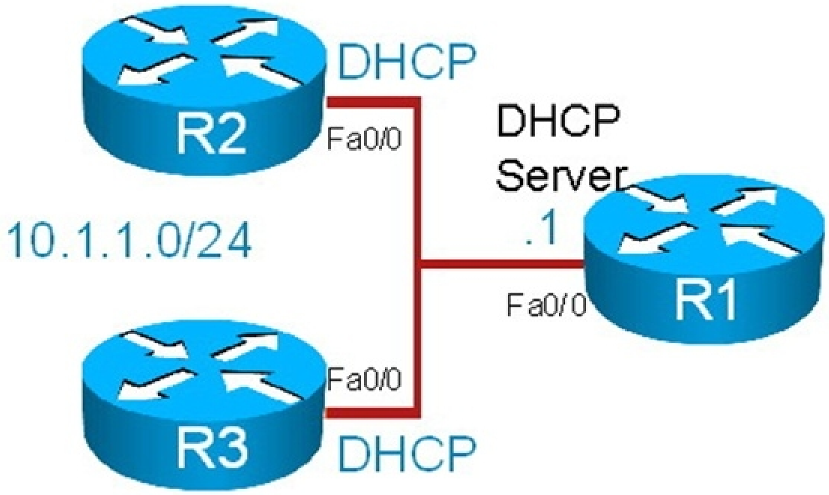

```
R2# show ip int brief
Interface         IP-Address   OK? Method Status                Protocol
FastEthernet0/0   unassigned   YES DHCP   up                    up
FastEthernet0/1   unassigned   YES NVRAM  administratively down down
Serial0/0/0       unassigned   YES NVRAM  administratively down down
Serial0/0/1       unassigned   YES NVRAM  administratively down down

R3# show ip int brief
Interface         IP-Address   OK? Method Status                Protocol
FastEthernet0/0   unassigned   YES DHCP   up                    up
FastEthernet0/1   unassigned   YES NVRAM  administratively down down
Serial0/0/0       unassigned   YES NVRAM  administratively down down
Serial0/0/1       unassigned   YES NVRAM  administratively down down
```

DHCP Troubleshooting Example 1 – Cont.

```
R3# debug dhcp detail
DHCP client activity debugging is on (detailed)
R3#

*Aug 23 17:32:37.107: Retry count: 1 Client-ID: cisco-0019.5592.a442-Fa0/0
*Aug 23 17:32:37.107: Client-ID hex dump: 636973636F2D303031392E353539322E
*Aug 23 17:32:37.107: 613434322D4551302F30
*Aug 23 17:32:37.107: Hostname: R3
*Aug 23 17:32:37.107: DHCP: SDiscover: sending 291 byte length DHCP packet
*Aug 23 17:32:37.107: DHCP: SDiscover 291 bytes
*Aug 23 17:32:37.107: B cast on FastEthernet0/0 interface from 0.0.0.0
*Aug 23 17:32:40.395: DHCP: SDiscover attempt #2 for entry:
*Aug 23 17:32:40.395: Temp IP addr: 0.0.0.0 for peer on Interface: FastEthernet0/0
*Aug 23 17:32:40.395: Temp sub net mask: 0.0.0.0
*Aug 23 17:32:40.395: DHCP Lease server: 0.0.0.0, state: 1 Selecting
*Aug 23 17:32:40.395: DHCP transaction id: 13BA
*Aug 23 17:32:40.395: Lease: 0 secs, Renewal: 0 secs, Rebind: 0 secs
*Aug 23 17:32:40.395: Next timer fires after: 00:00:04
*Aug 23 17:32:40.395: Retry count: 2 Client-ID: cisco-0019.5592.a442-Fa0/0
*Aug 23 17:32:40.395: Client-ID hex dump: 636973636F2D303031392E353539322E
*Aug 23 17:32:40.395: 613434322D4551302F30
<output omitted>
*Aug 23 17:32:44.395: Hostname: R3
*Aug 23 17:32:44.395: DHCP: SDiscover: sending 291 byte length DHCP packet
*Aug 23 17:32:44.395: DHCP: SDiscover 291 bytes
*Aug 23 17:32:44.395: B cast on FastEthernet0/0 interface from 0.0.0.0
*Aug 23 17:32:48.395: DHCP: Qscan: Timed out Selecting state
%Unknown DHCP problem... No allocation possible
*Aug 23 17:32:57.587: DHCP: waiting for 60 seconds on interface FastEthernet0/0
```

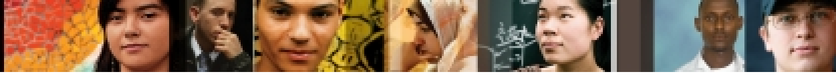

DHCP Troubleshooting Example 1 – Cont.

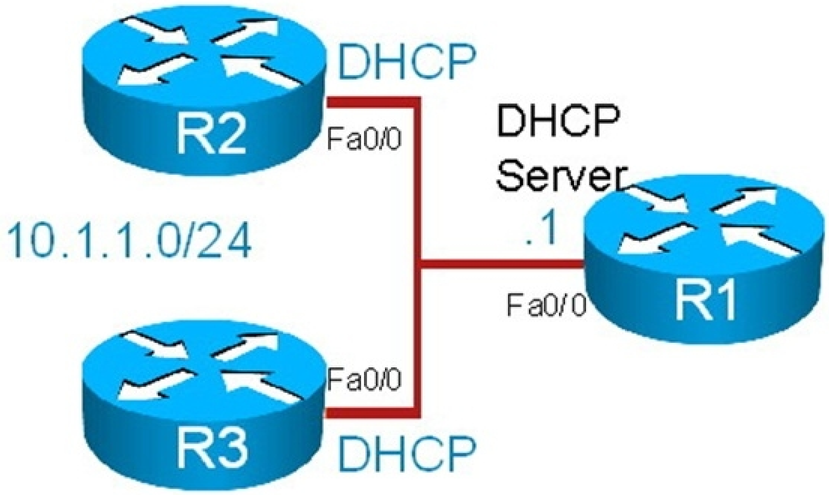

```
R1# show ip int brief
Interface          IP-Address   OK? Method Status                 Protocol
FastEthernet0/0    10.1.1.1     YES manual up                     up
FastEthernet0/1    unassigned   YES NVRAM  administratively down  down
Serial0/0/0        unassigned   YES NVRAM  administratively down  down
Serial0/0/1        unassigned   YES NVRAM  administratively down  down
```

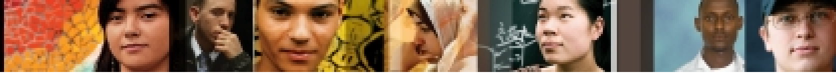

DHCP Troubleshooting Example 1 – Cont.

```
R1# show ip dhcp server statistics
Memory usage        9106
Address pools       1
Database agents     0
Automatic bindings  0
Manual bindings     0
Expired bindings    0
Malformed messages  0
Secure arp entries  0

Message             Received
BOOTREQUEST         0
DHCPDISCOVER        1
DHCPREQUEST         1
DHCPDECLINE         0
DHCPRELEASE         0
DHCPINFORM          0
Message Semt
BOOTREPLY           0
DHCPOFFER           1
DHCPACK             1
DHCPNAK             0

R1# sh ip dhcp pool
Pool vlan10 :
Utilization mark (high/low)  : 100/0
Subnet size (first/next)     : 0/0
Total addresses              : 254
Leased addresses             : 0
Pending event                : none
1 subnet is currently in the pool :
Current index    IP address range           Leased addresses
10.1.1.12        10.1.1.1 -10.1.1.254       0
```

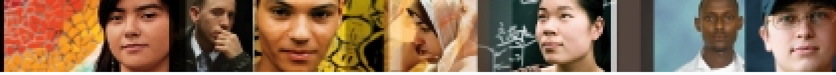

DHCP Troubleshooting Example 1 – Cont.

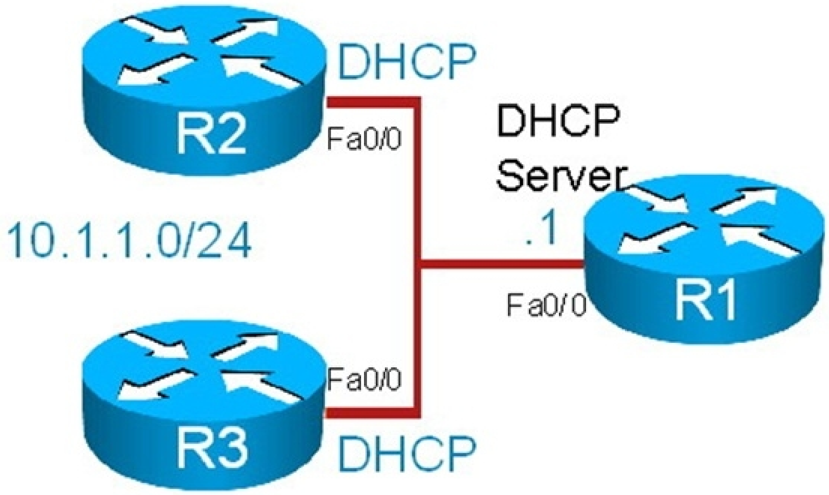

```
R1# show ip sockets
Proto  Remote     Port    Local      Port  In Out   Stat  TTY  OutputIF
88   --listen--      10.1.1.1        10  0     0      0    0
17   --listen--      10.1.1.1       161  0     0   1001    0
17   --listen--      10.1.1.1       162  0     0   1011    0
17   --listen--      10.1.1.1     57767  0     0   1011    0
17   --listen--      --any--        161  0     0   20001   0
17   --listen--      --any--        162  0     0   20011   0
17   --listen--      --any--      60739  0     0   20011   0
R1#
```

Note: There is no entry for UDP port 67 (DHCP server)

DHCP Troubleshooting Example 1 – Cont.

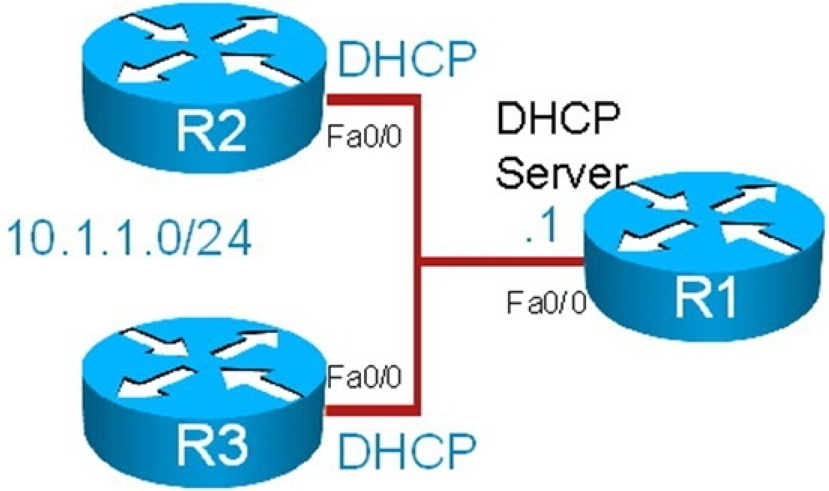

```
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# service dhcp
R1(config)# end
R1#

R1# show ip sockets
Proto  Remote    Port     Local     Port  In Out    Stat  TTY  OutputIF
88    --listen--       10.1.1.1       10  0    0       0    0
17    --listen--       10.1.1.1      161  0    0    1001    0
17    --listen--       10.1.1.1      162  0    0    1011    0
17    --listen--       10.1.1.1    57767  0    0    1011    0
17    --listen--       --any--       161  0    0    20001   0
17    --listen--       --any--       162  0    0    20011   0
17    --listen--       --any--     60739  0    0    20011   0
17 0.0.0.0          0 10.1.1.1        67  0    0     2211   0
R1#
```

DHCP Troubleshooting Example 2: Duplicate Client IP Addresses

- In this scenario, the IP address of router R1 Fa0/0 was previously 10.1.1.100.
- It has been changed to 10.1.1.1 to comply with a new network policy. This policy states that all branch routers will have the first IP address on any subnet
- After the change, some DHCP clients are reporting duplicated IP addresses. Users state that this happens sporadically, a few times a week.

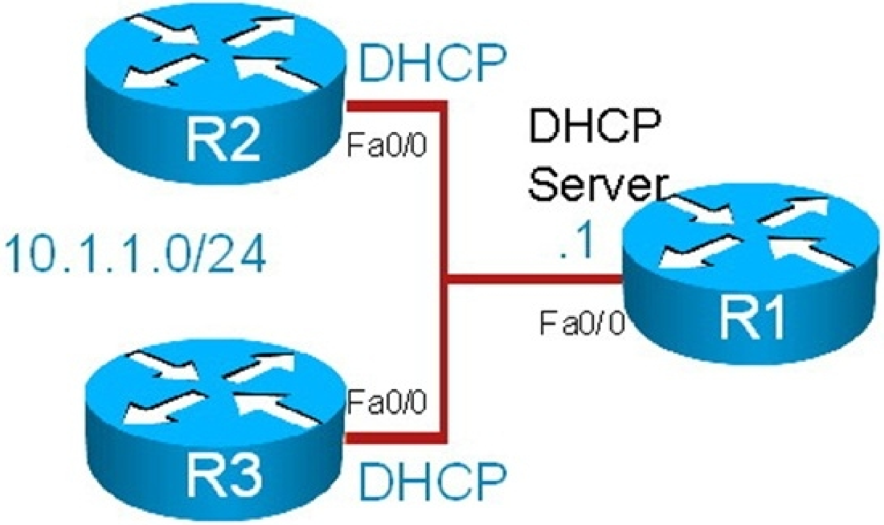

DHCP Troubleshooting Example 2 – Cont.

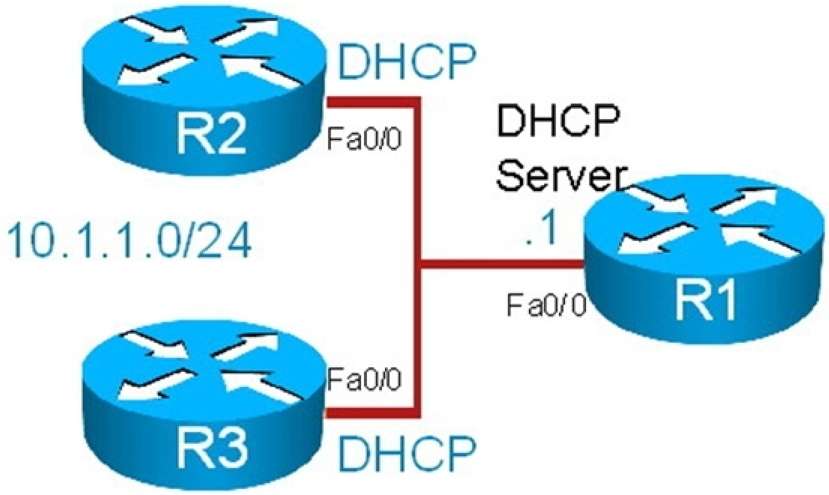

```
R1# show running-config | beg ip dhcp pool

ip dhcp pool vlan10
network 10.1.1.0 255.255.255.0
default-router 10.1.1.1
lease 3
```

DHCP Troubleshooting Example 2 – Cont.

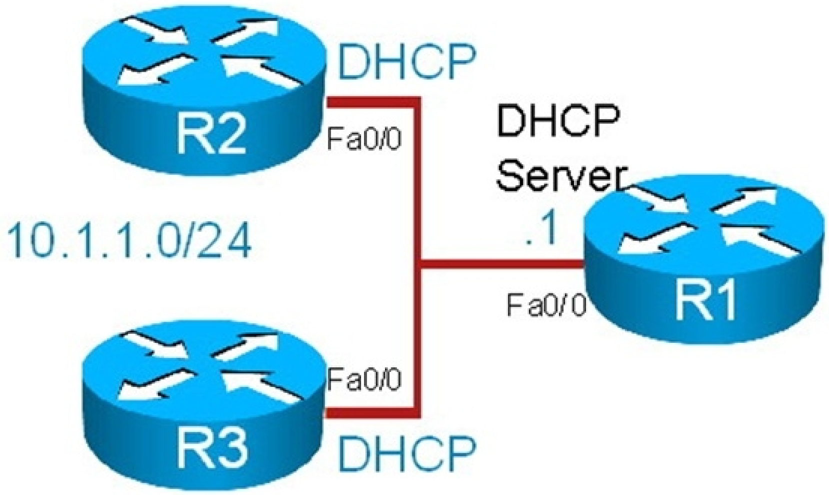

```
R1# show ip dhcp conflict
IP address         Detection method     Detection time              VRF
10.1.1.1           Gratuitous ARP       Aug 23 2009 06:28 PM
10.1.1.3           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.3           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.4           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.5           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.6           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.7           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.8           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.9           Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.10          Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.11          Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.12          Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.13          Gratuitous ARP       Aug 23 2009 06:29 PM
--More--
```

DHCP Troubleshooting Example 2 – Cont.

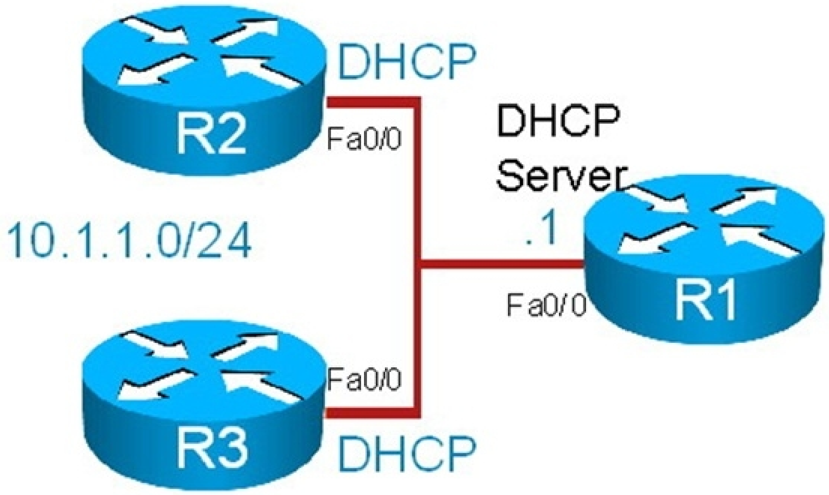

```
R1# sh run | inc excluded

ip dhcp excluded-address 10.1.1.100

R1#
```

DHCP Troubleshooting Example 2 – Cont.

Note: Configure R1 to exclude the range of addresses that are to be reserved for static assignment.

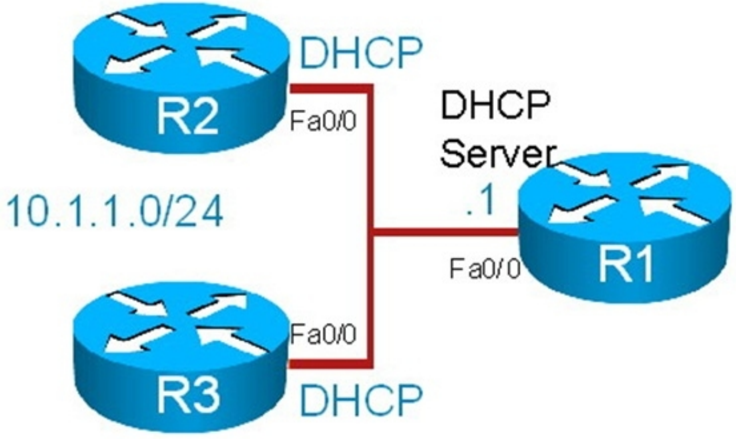

```
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.

R1(config)# no ip dhcp excluded-address 10.1.1.100
R1(config)# ip dhcp excluded-address 10.1.1.1 10.1.1.20
R1(config)# end
R1#
```

DHCP Troubleshooting Example 3: Relay Agent Issue

- R4 is a centrally located DHCP server.
- The DHCP clients in network segment 10.1.1.0 are unable to obtain IP address and other parameters.
- R2 is a DHCP client that is having trouble acquiring ip address.
- R1 is supposed to act as a relay agent and forward DHCP messages between local clients and the DHCP server (R4).

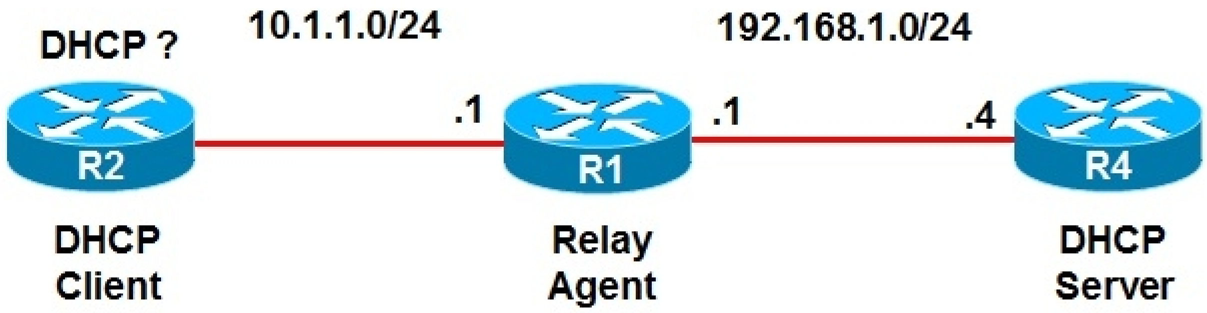

DHCP Troubleshooting Example 3 – Cont.

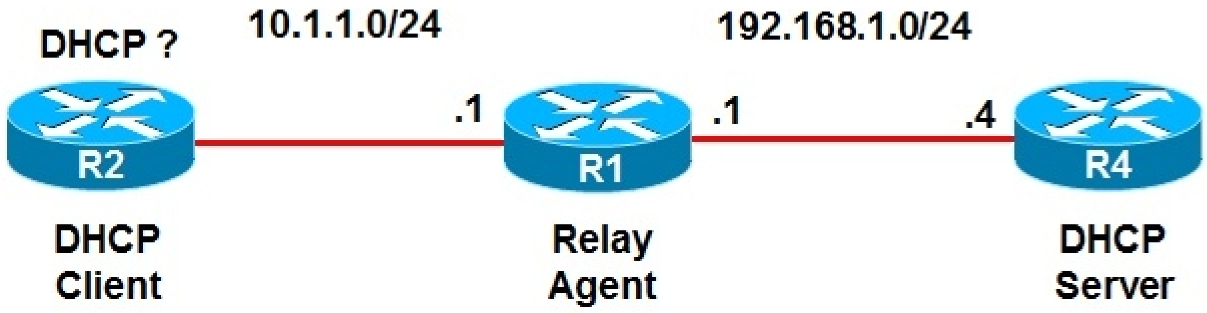

```
R1# debug ip udp
UDP packet debugging is on
R1#
R1#
*Aug 23 19:01:05.303: UDP: rcvd src-0.0.0.0(68), dst=255.255.255.255(67),
length=584
*Aug 23 19:01:05.303: UDP: broadcast packet dropped, src=0.0.0.0,
dst=192.168.1.255
*Aug 23 19:01:08.911: UDP: rcvd src-0.0.0.0(68), dst=255.255.255.255(67),
length=584
*Aug 23 19:01:08.911: UDP: broadcast packet dropped, src=0.0.0.0,
dst=192.168.1.255
*Aug 23 19:01:12.911: UDP: rcvd src-0.0.0.0(68), dst=255.255.255.255(67),
length=584
*Aug 23 19:01:12.911: UDP: broadcast packet dropped, src=0.0.0.0,
dst=192.168.1.255
<output omitted>
```

DHCP Troubleshooting Example 3 – Cont.

Note: Configure R1 with a helper address to forward DHCP requests to R4.

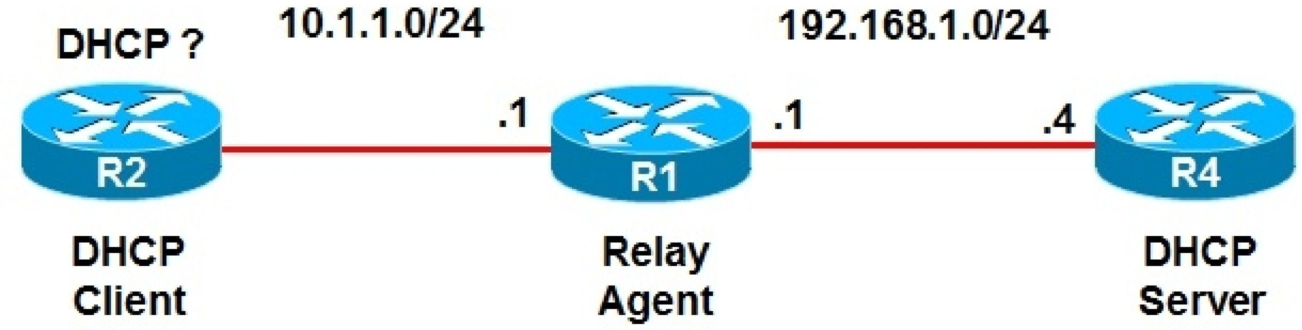

```
R1(config)# int fa0/0
R1(config-if)# ip helper-address 192.168.1.4
R1(config-if)# end
```

DHCP Troubleshooting Example 3 – Cont.

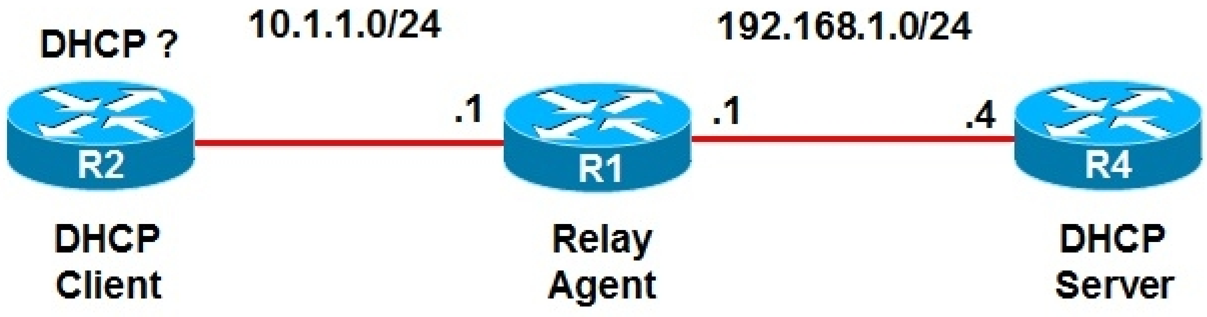

```
R4# debug ip udp
UDP packet debugging is on
R4#
*Aug 23 19:31:39.303: UDP: sent src=0.0.0.0(67), dst=255.255.255.255(68),length=308
*Aug 23 19:31:39.303: UDP: rcvd src=0.0.0.0(68), dst=255.255.255.255(67),length=584
*Aug 23 19:31:39.303: UDP: sent src=0.0.0.0(67), dst=255.255.255.255(68),length=308
*Aug 23 19:31:40.159: UDP: rcvd src=0.0.0.0(68), dst=192.168.1.4(67), length=584
*Aug 23 19:31:44.159: UDP: rcvd src=0.0.0.0(68), dst=192.168.1.4(67), length=584
*Aug 23 19:31:46.307: UDP: rcvd src=10.1.1.11(53470), dst=255.255.255.255(69),length=30
*Aug 23 19:31:49.307: UDP: rcvd src=10.1.1.11(53470), dst=255.255.255.255(69),length=30
<output omitted>
*Aug 23 19:32:28.439: UDP: rcvd src=10.1.1.11(53470), dst=255.255.255.255(69),length=29
*Aug 23 19:32:31.439: UDP: rcvd src=10.1.1.11(53470), dst=255.255.255.255(69),length=29
*Aug 23 19:32:35.439: UDP: rcvd src=10.1.1.11(53470), dst=255.255.255.255(69),length=29
*Aug 23 19:32:37.011: UDP: rcvd src=0.0.0.0(68), dst=192.168.1.4(67), length=584
```

Identifying Common IPv6 Routing Issues

IPv6 Routing Overview – IPv6 vs IPv4

- IPv6 has many similarities with IPv4 and that makes information gathering and analysis processes similar as well.
- Commands that contain the word “IP” maintain most of their syntax in terms of structure.
- IPv6 and IPv4 similarities represent a benefit but their differences can affect the troubleshooting process.
- There are no broadcasts in IPv6. Neighbors are discovered through ICMPv6 multicasts.
- The addressing structure is still hierarchical, and CIDR, or classless inter-domain routing rules and nomenclature still apply.
- The IPv6 subnet mask is potentially longer, for example /96, or even /128, as compared to /32 (at the most) with IPv4.
- OSPFv3 is still a link state protocol and neighbor sessions, LSAs, hierarchical areas, etc. still exist.

IPv6 Routing Overview – IPv6 vs IPv4 – Cont.

- For almost every IPv4 command there is an IPv6 counterpart.
  - The **`show ip route`** command is becomes **`show ipv6 route`**.
  - The **`show ip interface`** command becomes **`show ipv6 interface`**.
  - Testing commands, **`ping`** and **`trace`** maintain syntax and outcome consistency with their IPv4 counterparts.
- IPv6 is much more than a simple expansion in the address space.
  - It does away with broadcasts, which affects protocols such as DHCP
  - ARP does not exist.
  - Layer 2 addresses are gathered by hosts using the ICMPv6-based neighbor discovery process.
  - This process serves other purposes as well, including DAD or Duplicate Address Detection, stateless auto-configuration, and others.

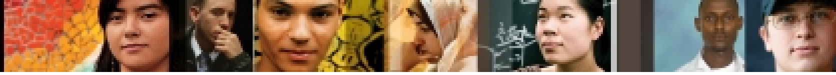

IPv6 Routing Overview – IPv6 vs IPv4 – Cont.

Some differences between IPv4 and IPv6

	IPv4	IPv6
Address Resolution Protocol	Used to find Layer 2 address mappings	Does not exist. ICMPv6 neighbor discovery is used instead
Secondary IP Addresses	Available, but the main IP address is used as packet source	Do not exist. Interfaces can have multiple concurrent IPv6 addresses of different types
Routing Protocols	Use interface IP address to exchange routing information	Use the link local address to create neighbor sessions and to assign as next-hop
Address Assignment	Static, or Dynamic (using DHCP)	Static, or Dynamic (using DHCP or stateless auto-configuration)

Troubleshooting IPv6 Issues - Cont.

- With IPv6 there are common configuration mistakes:
  - A misconfigured auto-configuration router that is not advertising network information to hosts will prevent IPv6 hosts from establishing full connectivity as they lack global unicast addresses.
- Other typical problem areas are related to IPv6 routing protocols and include:
  - Suboptimal routing due to improper summarization
  - Parameter mismatches on protocols such as OSPF that negotiate parameters.
- For tunnel scenarios, other components such as routing protocols, may need to change when the specific migration or tunneling method changes.
  - When using 6to4 tunnels, not all IGPs will function properly, due to the fact that when multicast addresses are used to establish adjacencies, those addresses are not properly mapped to a tunnel destination.

IPv6 Troubleshooting – `show` Commands

- **`show ipv6 interface:`**
  - Displays the usability status of interfaces configured for IPv6.
  - Validates the IPv6 status of an interface and its configured addresses.
  - If the interface’s hardware is usable, the interface is marked as up.
  - If the interface can provide two-way communication for IPv6, the line protocol is marked as up.
- **`show ipv6 routers:`**
  - IPv6-specific command (does not have an IPv4 counterpart).
  - Displays IPv6 router advertisement (RA) information received from onlink routers.
- **`show ipv6 route:`**
  - Displays the contents of the IPv6 routing table.
- **`show ipv6 protocols:`**
  - Displays the parameters and current state of the active IPv6 routing protocol processes.
  - The information displayed is useful in troubleshooting routing operations.

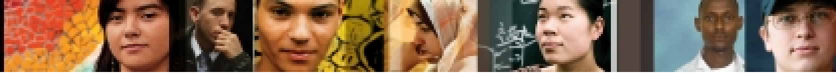

IPv6 Troubleshooting – debug Commands

- **`debug ipv6 routing:`**
  - Displays debugging messages for IPv6 routing table updates and route cache updates.
  - Displays messages whenever the routing table changes.
- **`debug ipv6 nd:`**
  - Displays debugging messages for IPv6 ICMP neighbor discovery (ND) transactions.
  - Can help determine whether the router is sending or receiving IPv6 ICMP ND messages.
- **`debug ipv6 packet:`**
  - Use this command to display debugging messages for IPv6 packets.
  - The debugging information includes packets received, generated, and forwarded.
  - Note that fast-switched packets do not generate messages.

IPv6 Troubleshooting Example 1: IPv6 Routing Problems

- Recent changes in the network have rendered router R3 isolated and with no connectivity outside the fast Ethernet segment connected to R1.
- This is verified by performing a ping from R3 to a remote destination (24::24:2), and it does not succeed.
- The changes were aimed at providing automatic IP address assignment and configuration to certain devices.
- After the change, R3 lost connectivity to the rest of the network.
- A bottom-up troubleshooting approach will be applied to resolve this problem, starting at R3.

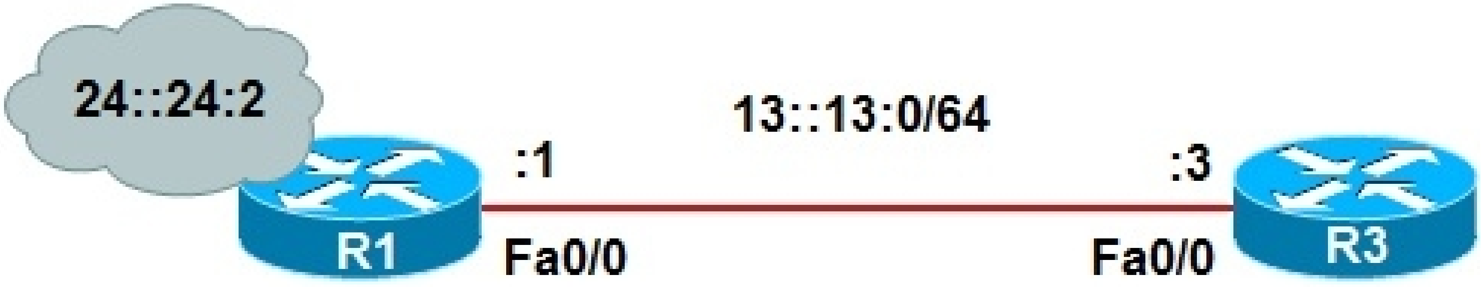

IPv6 Troubleshooting Example 1 – Cont.

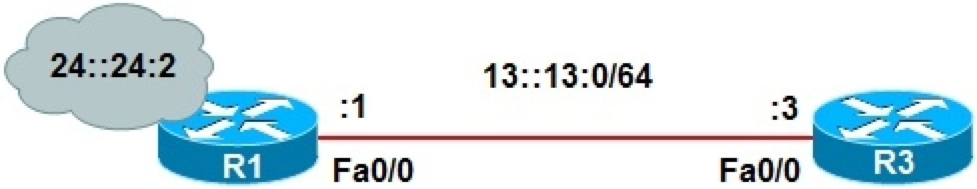

```
R3# show ipv6 interface Fa0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::219:55FF:FEF0:B7D0
  Global unicast address(es):
    13::13:3, subnet is 13::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF13:3
    FF02::1:FFF0:B7D0
  MTU is 1500 bytes
ICMP error messages limited to one every 100 milliseconds
ICMP redirects are enabled
ND DAD is enabled, number of DAD attempts: 1
ND reachable time is 30000 milliseconds
ND advertised reachable time is 0 milliseconds
ND advertised retransmit interval is 0 milliseconds
ND router advertisements are sent every 200 seconds
ND router advertisements live for 1800 seconds
Hosts use stateless autoconfig for addresses.
```

IPv6 Troubleshooting Example 1 – Cont.

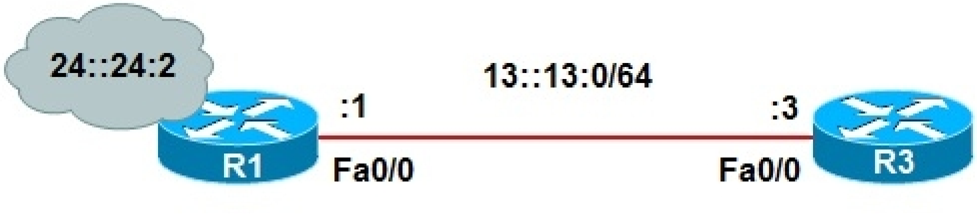

```
R3# show run int Fa0/0
Building configuration...
Current configuration : 111 bytes
!
Interface FastEthernet0/0
 no ip address
 duplex auto
 speed auto
 ipv6 address autoconfig
 ipv6 enable
end
```

IPv6 Troubleshooting Example 1 – Cont.

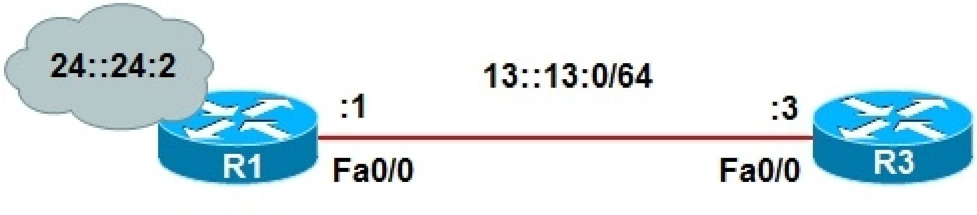

```
R3# show ipv6 route
IPv6 Routing Table - 6 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static Route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
C   13::/64 [0/0]
     via ::, Fast Ethernet0/0
L   13:219:55FF;FEF0:B7D0/128 [0/0]
     via ::, FastEthernet0/0
C   103::/64 [0/0]
     via ::, Loopback0
L   103::3/128 [0/0]
     via ::, Loopback0
L   FE80::/10 [0/0]
     via ::, Null0
L   FF00::/8 [0/0]
     via ::, Null0
```

IPv6 Troubleshooting Example 1 – Cont.

```
R3# debug ipv6 nd
ICMP Neighbor Discovery events debugging is on
R3#
R3#
R3# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# int Fa0/0
R3(config-if)# shutdown
R3(config-if)# no shutdown
R3(config-if)#
*Aug 23 21:44:18.491: ICMPv6-ND: Sending Final RA on FastEthernet0/0
*Aug 23 21:44:18.491: ICMPv6-ND: Address FE80::219:55FF:FEF0:B7D0/10 is down on
FastEthernet0/0
*Aug 23 21:44:20.491: %LINK-3-UPDOWN: Interface FastEthernet0/0, changed state to
up
*Aug 23 21:44:20.971: ICMPv6-ND: Sending NS for FE80::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:44:21.971: ICMPv6-ND: DAD: FE80::219:55FF:FEF0:B7D0 is unique
*Aug 23 21:44:21.971: ICMPv6-ND: Sending NA for FE80::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:44:21.971: ICMPv6-ND: Address FE80::219:55FF:FEF0:B7D0 is up on
FastEthernet0/0
*Aug 23 21:44:23.971: ICMPv6-ND: Sending RS on FastEthernet0/0
*Aug 23 21:44:27.971: ICMPv6-ND: Sending RS on FastEthernet0/0
*Aug 23 21:44:31.971: ICMPv6-ND: Sending RS on FastEthernet0/0
```

IPv6 Troubleshooting Example 1 – Cont.

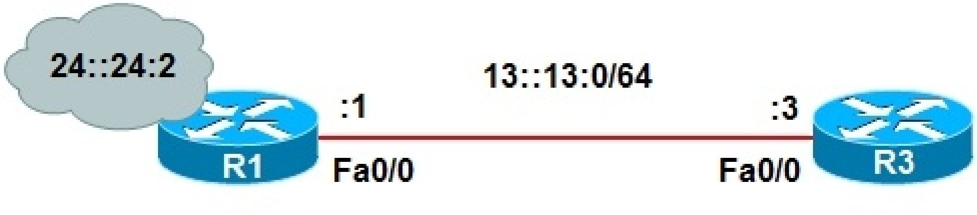

```
R1# show running-config interface f0/0
Building configuration...
Current configuration : 112 bytes
!
interface FastEthernet0/0
 no ip address
 duplex auto
 speed auto
 ipv6 address 13::13:1/64
 ipv6 enable
end
```

IPv6 Troubleshooting Example 1 – Cont.

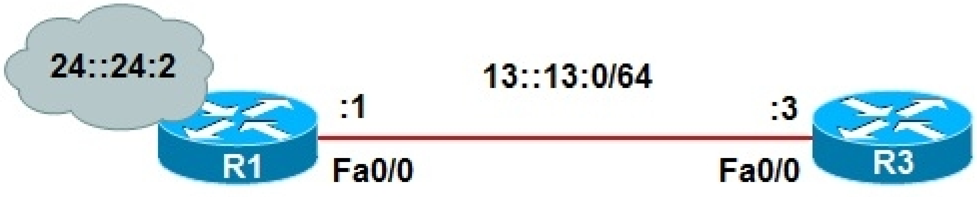

```
R1# show ipv6 interface f0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::219:56FF:FE2C:9856
  Global unicast address(es):
    13::13:1, subnet is 13::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF13:1
    FF02::1:FF2C:9856
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  Default router is FE80::219:55F:FEF0:B7D0 on FastEthernet0/0
```

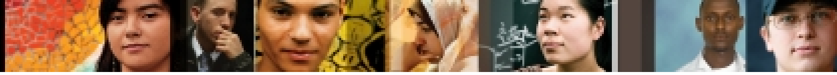

IPv6 Troubleshooting Example 1 – Cont.

- What else does R1 need to become a proper autoconfiguration router?
- How can it be that R3 has a working IPv6 address if the autoconfiguration process is not working?
- Why is it that R3 cannot access the rest of the network, even with a working IPv6 address and no noticeable physical issues?

- One requirement for autoconfiguration is explicitly enabling IPv6 unicast routing.

```
R1# show run | inc unicast-routing
R1#
R1# debug ipv6 nd
ICMP Neighbor Discovery events debugging is on
R1#
R1# conf t
Enter configuration commands, one per line. End with CNTL/A.
R1(config)# ipv6 unicast-routing
R1(config)# end
R1#
*Aug 23 22:01:45.175: ICMPv6-ND: Sending RA to FF02::1 on FastEthernet0/0
*Aug 23 22:01:45.175: ICMPv6-ND: MTU = 1500
*Aug 23 22:01:45.175: ICMPv6-ND: prefix = 13::/64 onlink autoconfig
*Aug 23 22:01:45.175: ICMPv6-ND: 2592000/604800 (valid/preferred)
```

IPv6 Troubleshooting Example 1 – Cont.

The **debug** output on R3 upon enabling IPv6 routing on R1

```
R3# debug ipv6 nd
ICMP Neighbor Dicovery events debugging is on
R3#
R3# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R3(config)# int Fa0/0
R3(config-if)# shutdown
R3(config-if)# no shutdown
R3(config-if)#
*Aug 23 21:57:47.547: ICMPv6-ND: Sending Final RA on FastEthernet0/0
*Aug 23 21:57:47.547: ICMPv6-ND: Address 13::219:55FF:FEF0:B7D0/64 is down on
FastEthernet0/0
*Aug 23 21:57:47.547: ICMPv6-ND: Address FE80::219:55FF:FEF0:B7D0/10 is down on FastEthernet0/0
R3#
*Aug 23 21:57:48.003: %SYS-5-CONFIG_I: Configured from console by console
*Aug 23 21:57:53.279: ICMPv6-ND: Sending NS for FE80::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:57:54.279: ICMPv6-ND: DAD: FE80::219:55FF:FEF0:B7D0 is unique
*Aug 23 21:57:54.279: ICMPv6-ND: Sending NA for FE80::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:57:56.279: ICMPv6-ND: Sending RS on FastEthernet0/0
*Aug 23 21:57:56.279: ICMPv6-ND: Received RA from FE80::219:56FF:FE2C:9856 on
FastEthernet0/0
*Aug 23 21:57:56.279: ICMPv6-ND: Sending NS for 13::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:57:56.279: ICMPv6-ND: Autoconfiguring 13::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:57:57.279: ICMPv6-ND: DAD: 13::219:55FF:FEF0:B7D0 is unique
*Aug 23 21:57:57.279: ICMPv6-ND: Sending NA for 13::219:55FF:FEF0:B7D0 on
FastEthernet0/0
*Aug 23 21:57:57.279: ICMPv6-ND: Address 13::219:55FF:FEF0:B7D0 is up on
FastEthernet0/0
```

IPv6 Troubleshooting Example 1 – Cont.

The debug output on R3 upon enabling IPv6 routing on R1

```
R1# sh ipv6 int f0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::219:55FF:FEF0:B7D0
  Global unicast address(es):
    13::219:55FF:FEF0:B7D0, subnet is 13::/64 [PRE]
      Valid lifetime 2591941 preferred lifetime 604741
  Joined group address(es):
   FF02::1
   FF02::2
   FF02::1:FFF0:B7D0
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 0 milliseconds
  ND advertised reachable time is 0 milliseconds
  ND advertised retransmit interval is 0 milliseconds
  ND router advertisements are sent every 200 seconds
  ND router advertisements live for 1800 seconds
  Hosts use stateless autoconfig for addresses
R1#

R3# ping 24::24:2
Type escape sequence to abort
Sending 5, 100-byte ICMP Echos to 24::24:2, timeout is 2 seconds:
!!!!!
Success reate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
R3#
```

IPv6 Troubleshooting Example 2: Redistribution Issue

- R1 and R2 have two connections: one through the 6to4 tunnel, and a main link over a Frame Relay network.
- A RIPng process is running between R3 and R1 and then across the tunnel.
- A second RIPng process is running between R4 and R2 and across the Frame Relay virtual circuit.
- A recent change ticket performed two-way redistribution between the two RIP processes on R2 and R4 has lost reachability to R3, specifically the loopback interface on R3.

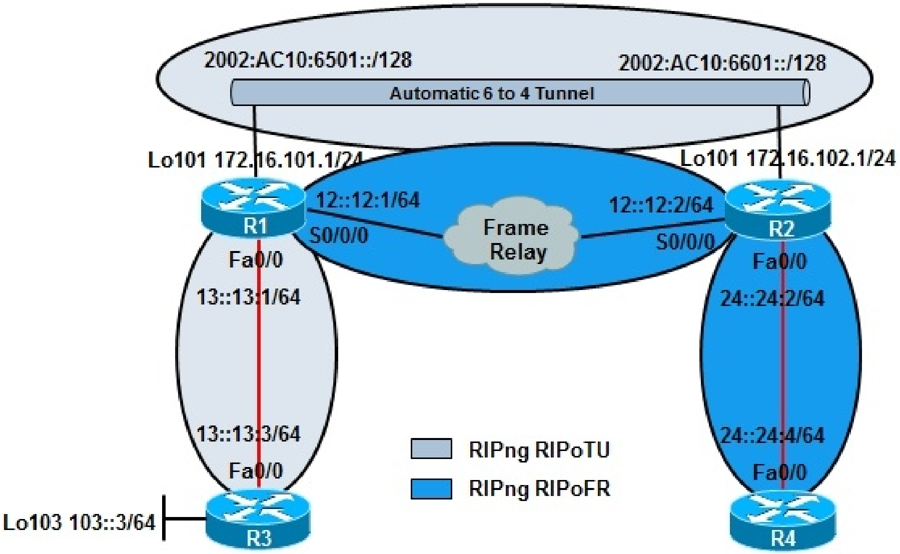

IPv6 Troubleshooting Example 2 – Cont.

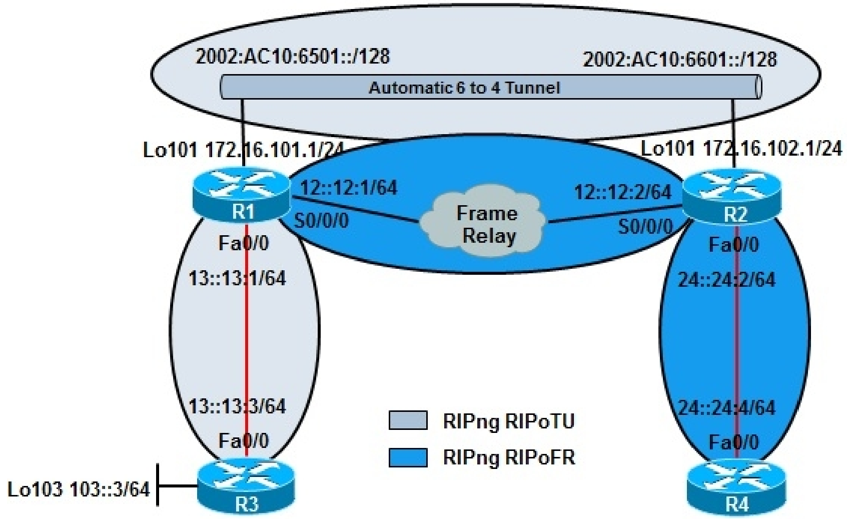

```
R4# show ipv6 route
IPv6 Routing Table - 6 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static Route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
R   12::/64 [120/2]
     via FE80::219:55FF:FE92:A442, Fast Ethernet0/0
R   103:::/64 [120/7]
     via FE80::219:55FF:FE92:A442, Fast Ethernet0/0

R4# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

IPv6 Troubleshooting Example 2 – Cont.

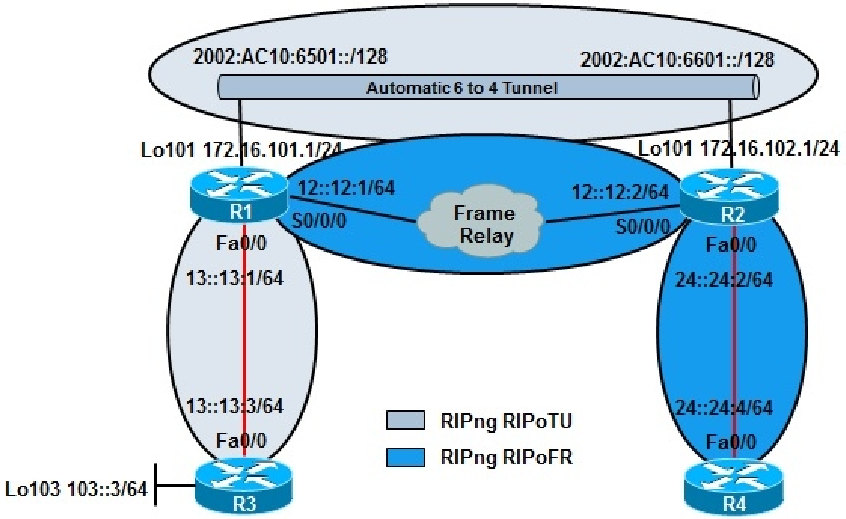

```
R2# show ipv6 protocols
IPv6 Routing Protocol is “connected”
IPv6 Routing Protocol is “static”
IPv6 Routing Protocol is “rip RIPoFR”
  Interfaces:
    FastEthernet0/0
    Serial0/0/0
  Redistribution:
    Redistributing protocol rip RIPoTU with metric 5
IPv6 Routing Protocol is “rip RIPoTU”
  Interfaces:
    Loopback101
    Tunnel0
Redistribution:
  Redistributing protocol rip RIPoFR with metric 15
```

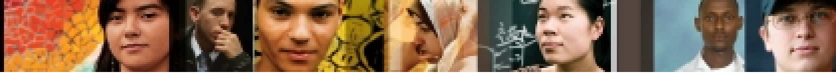

IPv6 Troubleshooting Example 2 – Cont.

R1’s IPv6 routing table is missing networks such as 24::24/64

```
R1# show ipv6 route
<output omitted>
S      ::/0 [1/0]
        via ::, Tunnel0
C      12::/64 [0/0]
        via ::, Serial0/0/0
L      12::12:1/128 [0/0]
        via ::, Serial0/0/0
C      13::/64 [0/0]
        via ::, FastEthernet0/0
L      13::13:1/128 [0/0]
        via ::, FastEthernet0/0
R      103::/64 {120/2]
        via FE80::219:55FF:FEF0:B7D0, FastEthernet0/0
LC     2002:AC10:6501::1/128 [0/0]
        Via ::, Tunnel0
L      FE80::/10 [0/0]
        via ::, Null0
L      FF00::/8 [0/0]
        via ::, Null0
```

IPv6 Troubleshooting Example 2 – Cont.

Correcting the redistribution metric on R2 from 15 to 10

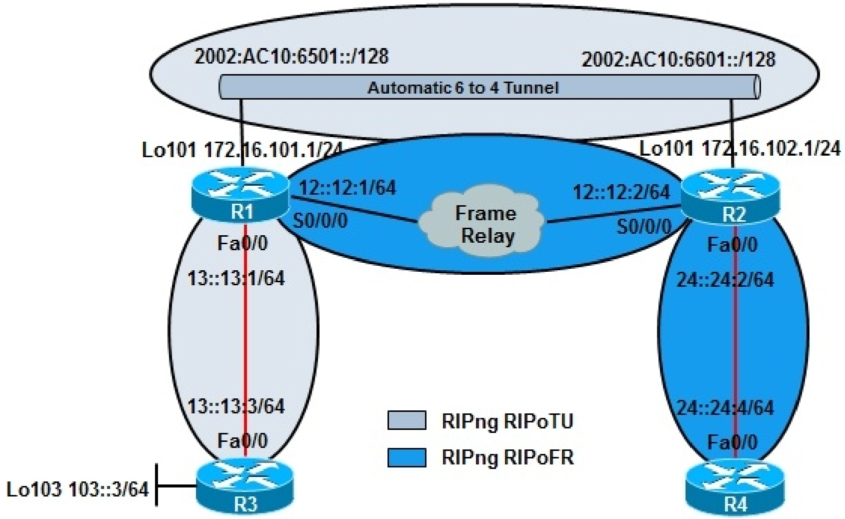

```
R2# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R2(config)# ipv6 router rip RIPoTU
R2(config-rtr)# redistribute rip RIPoFR metric 10
R2(config-rtr)#
R2(config-rtr)# end
R2#
```

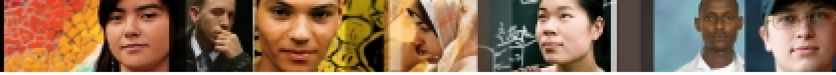

IPv6 Troubleshooting Example 2 – Cont.

The RIPoFR process has no interfaces enabled on R1

```
R1# show ipv6 rip
RIP process “RIPoTU”, port 521, multicast-grop FF02::9, pid 178
     Administrative distance is 120. Maximum paths is 16
     Updates every 30 seconds, expire after 180
     Holddown lasts 0 seconds, garbage collection after 120
     Split horizon is on; poison reverse is off
     Default routes are not generated
     Periodic updates 201, trigger updates 13
  Interfaces:
    Tunnel0
    FastEthernet0/0
  Redistribution:
    None
RIP process “RIPoFR”, port 521, multicast-grop FF02::9, pid 179
     Administrative distance is 120. Maximum paths is 16
     Updates every 30 seconds, expire after 180
     Holddown lasts 0 seconds, garbage collection after 120
     Split horizon is on; poison reverse is off
     Default routes are not generated
     Periodic updates 197, trigger updates 6
  Interfaces:
    None
  Redistribution:
    Redistributing protocol rip RIPoTU w
```

IPv6 Troubleshooting Example 2 – Cont.

The **debug ipv6 routing** results on R1 after enabling RIPoFR on S0/0/0

```
R1# debug ipv6 routing
IPv6 routing table events debugging is on
R1#
R1#
R1#
R1#conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# int s0/0/0
R1(config-if)# ipv6 rip RIPoFR enable
R1(config-if)#
R1(config-if)# end
R1#
R1#
*Aug 24 00:42:47.250: IPv6RTO: Event: 11::/64, Mod, owner connected, previous None
*Aug 24 00:42:47.250: IPv6RTO: Event: 12::/64, Mod, owner connected, previous None
*Aug 24 00:42:47.538: %SYS-5-CONFIG_I: Configured from console by console
*Aug 24 00:42:49.206: IPv6RTO: rip RIPoFR, Route add 24::/64 [new]
*Aug 24 00:42:49.206: IPv6RTO: rip RIPoFR, Add 24::/64 to table
*Aug 24 00:42:49.206: IPv6RTO: rip RIPoFR, Adding next-hop
FE80::219::55FF:FE92:A442 over Serial 0/0/0 for 24::/64, [120/2]
*Aug 24 00:42:49.206: IPv6RTO: rip RIPoFR, Added backup for 12::/64, distance 120
*Aug 24 00:42:49.206: IPv6RTO: rip RIPoFR, Added backup for 11::/64, distance 120
*Aug 24 00:42:49.206: IPv6RTO: Event: 24::/64, Add, owner rip, previous None
```

IPv6 Troubleshooting Example 2 – Cont.

The R1 IPv6 routing table after enabling RIPoFR on S0/0/0

```
R1# show ipv6 route
IPv6 Routing Table - 17 entries
<output omitted>
S    ::/0 [1/0]
      via ::, Tunnel0
C    11::/64 [0/0]
      via ::, Serial0/0/0
L    11::11:1/128 [0/0]
      via ::, Serial0/0/0
C    12::/64 [0/0]
      via ::, Serial0/0/0
L    12::12:1/128 [0/0]
      via ::, Serial0/0/0
C    13::/64 [0/0]
      via ::, FastEthernet0/0
L    13::13:1/128 [0/0]
      via ::, FastEthernet0/0
R    24::/64 [120/2]
      via FE80::219:55FF:FE92:A442, Serial0/0/0
R    103::/64 [120/2]
      via FE80::219:55FF:FEF0:B7D0, FastEthernet0/0
C    2001:AC10:6501::/64 [0/0]
      via ::, Tunnel0
L    2001:AC10:6501::1/128 [0/0]
      via ::, Tunnel0
R    2001:AC10:6601::/64 [120/2]
      via FE80::219:55FF:FE92:A442, Serial0/0/0
S    2002::/16 [1/0]
      via Tunnel0
LC   2002:AC10:6501::1/128 [0/0]
      via ::, Tunnel0
R    2002:AC10:6601::2/128 [120/2]
      via FE80::219:55FF:FE92:A442, Serial0/0/0
<output omitted>
```

IPv6 Troubleshooting Example 2 – Cont.

The **debug IPv6 packet** command output on R3

```
R3# show ipv6 access-list
IPv6 access list FILTER
Permit ipv6 any 24::/64 (5 matches) sequence 10
R3#
R3# debug ipv6 packet access-list FILTER
IPv6 packet debugging is on for access list FILTER

*Aug 24 11:25:11.748 IPv6: Sending on Loopback103
*Aug 24 11:25:11.748 IPv6: Sending on FastEthernet0/0
*Aug 24 11:25:39.812: IPv6: Sending on Loopback103
*Aug 24 11:25:39.812: IPv6: Sending on FastEthernet0/0

R3# ping 24::24:2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24::24:2, timeout is 2 seconds

*Aug 24 11:25:51.332: IPv6: SAS picked source 13::13:3 for 24::24:2
*Aug 24 11:25:51.332: IPv6: source 13::13:3 (local)
*Aug 24 11:25:51.332: dest 24::24:2
*Aug 24 11:25:51.332: traffic class 0, flow 0x0, len 100+0, prot 58, hops
64, Route not found
```

IPv6 Troubleshooting Example 2 – Cont.

Checking redistribution on R1

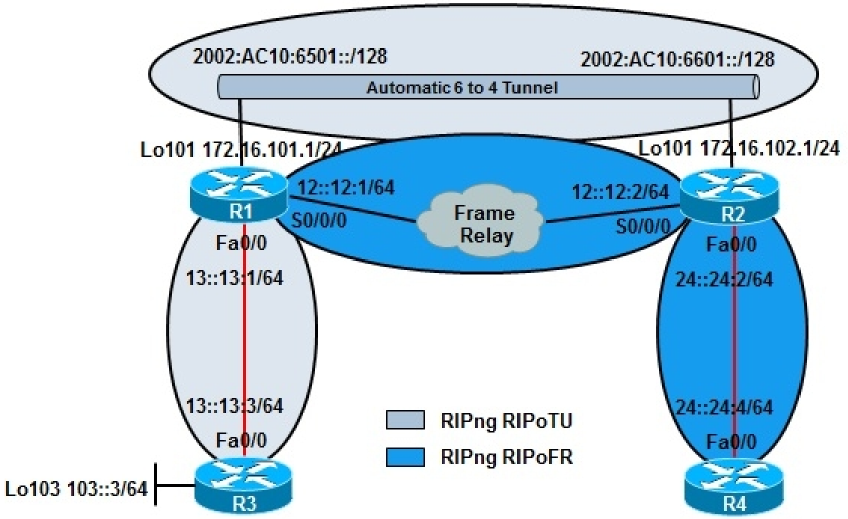

```
R1# show ipv6 protocols
IPv6 Routing Protocol is “connected”
IPv6 Routing Protocol is “static”
IPv6 Routing Protocol is “rip RIPoTU”
  Interfaces:
    Tunnel0
    FastEthernet0/0
  Redistribution:
    None
IPv6 Routing Protocol is “rip RIPoFR”
  Interfaces:
    Serial0/0/0
  Redistribution:
    Redistributing protocol rip RIPoTU with metric 5
```

IPv6 Troubleshooting Example 2 – Cont.

After redistributing RIPoFR into RIPoTU, R3 has an IPv6 path to R4.

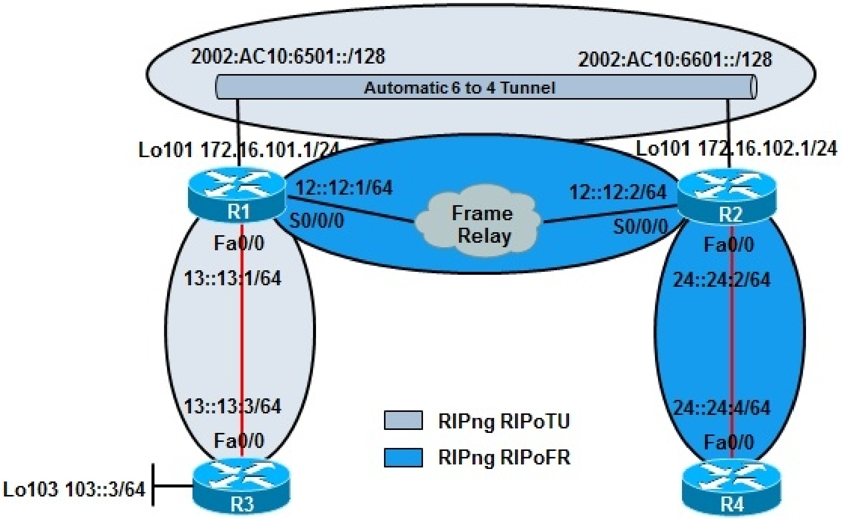

```
R1(config)# ipv6 router rip RIPoTU
R1(config-rtr)# redistribute rip RIPoFR include-connected metric 5
R1(config-rtr)#
R1(config)#

R3# ping 24::24:4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24::24:2, timeout is 2 seconds
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/4 ms
```

IPv6 Troubleshooting Example 3: OSPFv3 Configuration Errors

- OSPFv3 is running throughout the network.
- Backbone area 0 (routers R1 and R2) runs over the Frame Relay network.
- Stub area 1 (R1 and R3) and stub area 2 (R2 and R4) run over Fast Ethernet.
- A recent power outage might have caused some of the routers to restart.
- The configurations may not have been saved before the incident.
- End-to-end network connectivity has been lost since the power outage.

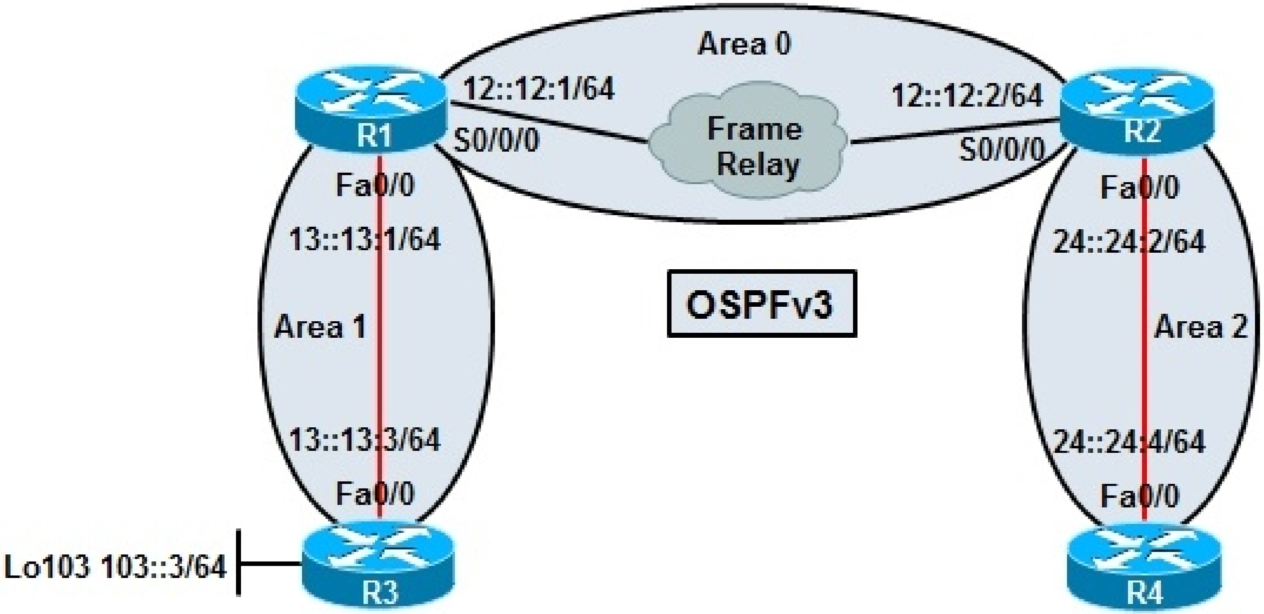

IPv6 Troubleshooting Example 3 - Cont.

Initial ping tests from R4 and R2 to R3’s Loopback interface fail

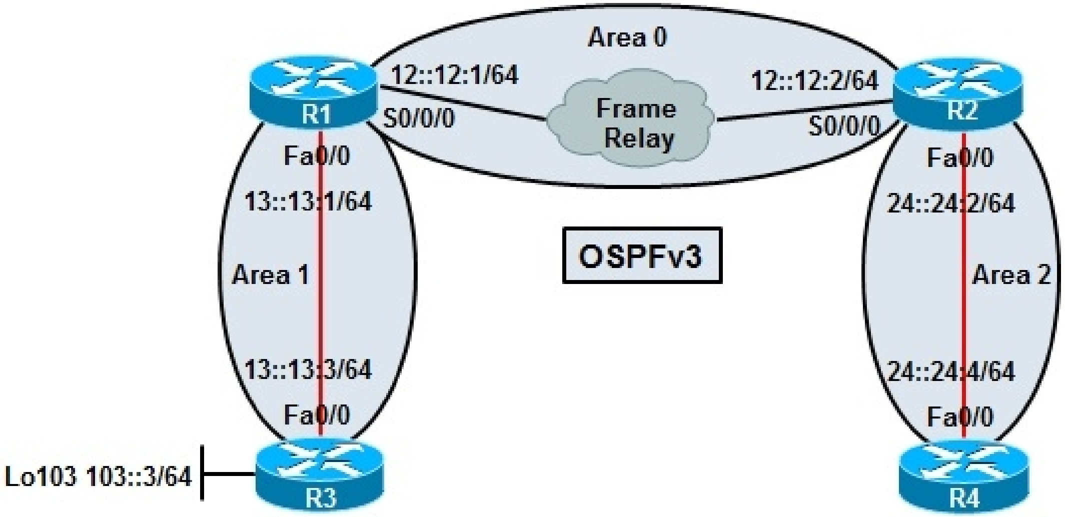

```
R4# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R2# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

IPv6 Troubleshooting Example 3 - Cont.

Ping tests from R1 to R3’s Loopback interface fail. The ping from R1 to R3 Fa0/0 succeeds.

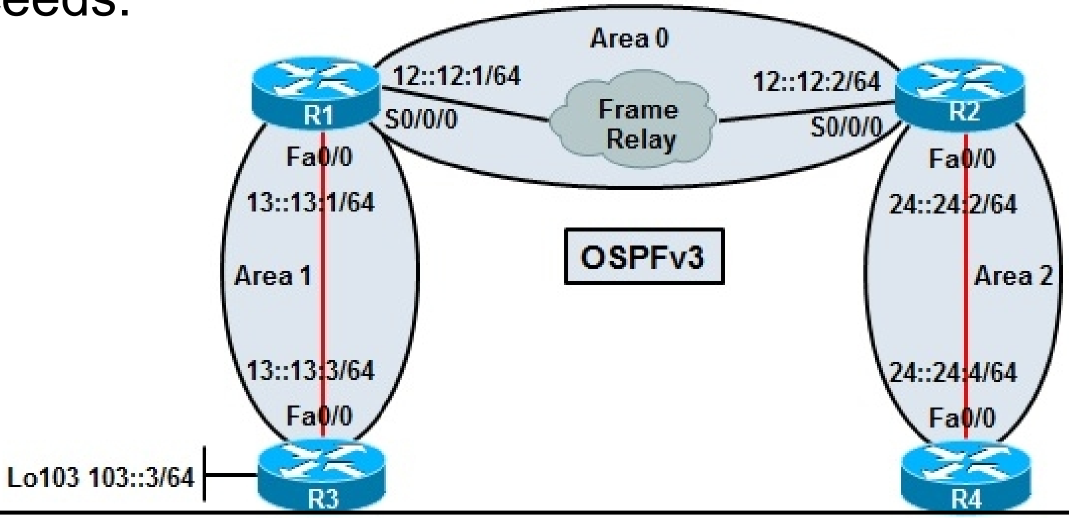

```
R1# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)

R1# ping 13::13:3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 13::13:3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/1 ms

R2# ping 13::13:3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 13::13:3, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
```

IPv6 Troubleshooting Example 3 - Cont.

`show` commands on R1 reveal that R1 has no neighbors.

```
R1# show ipv6 ospf interface
Serial0/0/0 is up, line protocol is up
  Link Local Address FE80::219:56FF:FE2C:9856, Interface ID 5
  Area 0, Process ID 1, Instance ID 0, Router ID 172.16.101.1
  Network Type POINT_TO_POINT, Cost: 64
  Transmit Delay is 1 sec, State POINT_TO_POINT
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Index 1/1/2, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 1, maximum is 2
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::219:56FF:FE2C:9856, Interface ID 3
  Area 1, Process ID 1, Instance ID 0, Router ID 172.16.101.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.101.1, local address FE80::219:56FF:FE2C:9856
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
    Hello due in 00:00:01
  Index 1/1/1, flood queue length 0
  Next 0x0(0)/0x0(0)/0x0(0)
  Last flood scan length is 0, maximum is 0
  Last flood scan time is 0 msec, maximum is 0 msec
  Neighbor Count is 0, Adjacent neighbor count is 0
  Suppress hello for 0 neighbor(s)

R1# show ipv6 ospf neighbor
R1#
```

IPv6 Troubleshooting Example 3 - Cont.

The **debug ipv6 ospf hello** shows mismatched stub/transit area.

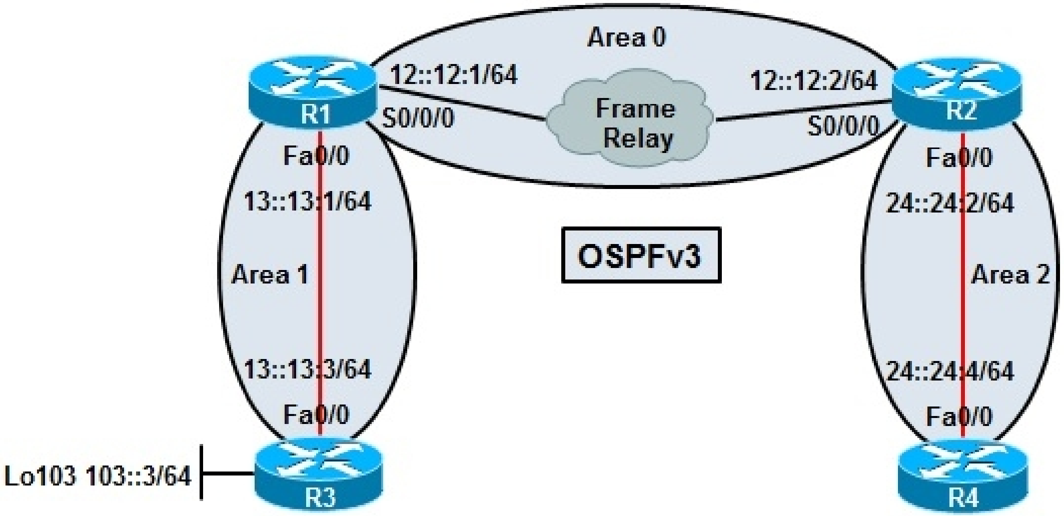

```
R1# debug ipv6 ospf hello
OSPFv3 hello events debugging is on
R1#
R1#
*Aug 24 13:19:49.751: OSPFv3: Rcv hello from 103.103.103.103 area 1 from
FastEthernet0/0 FE80::219:55FF:FEF0:B7D0 interface ID 3
*Aug 24 13:19:49.751: OSPFv3: Hello from FE80::219:55FF:FEF0:B7D0 with mismatched
Stub/Transit area option bit
*Aug 24 13:19:56.087: OSPFv3: Send hello to FF02::5 area 1 on FastEthernet0/0 from
FE80::219:56FF:FE2C:9856 interface ID 3
*Aug 24 13:19:59.751: OSPFv3: Rcv hello from 103.103.103.103 area 1 from
FastEthernet0/0 FE80::219:55FF:FEF0:B7D0 interface ID 3
*Aug 24 13:19:59.751: OSPFv3: Hello from FE80::219:55FF:FEF0:B7D0 with mismatched
Stub/Transit area option bit
```

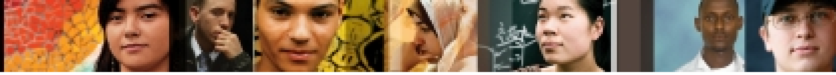

IPv6 Troubleshooting Example 3 - Cont.

The **show ipv6 ospf** command output on R1 indicates area 1 is totally stubby.

```
R1# show ipv6 ospf
Routing Process "ospfv3 1" with ID 172.16.101.1
It is an area border router
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Number of external LSA 0. Checksum Sum 0x000000
Number of areas in this router is 2. 1 normal 1 stub 0 nssa
Reference bandwidth unit is 100 mbps
  Area BACKBONE(0) (Inactive)
    Number of interfaces in this area is 1
    SPF algorithm executed 3 times
    Number of LSA 4. Checksum Sum 0x01F36B
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
  Area 1
    Number of interfaces in this area is 1
    It is a stub area, no summary LSA in this area
    Generates stub default route with cost 1
    SPF algorithm executed 5 times
    Number of LSA 4. Checksum Sum 0x016293
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

IPv6 Troubleshooting Example 3 - Cont.

The **`show ipv6 ospf`** command output on R3 indicates area 1 is a normal area.

```
R3# show ipv6 ospf
Routing Process “ospfv3 1” with ID 103.103.103.103
SPF schedule delay 5 secs, Hold time between two SPFs 10 secs
Minimum LSA interval 5 secs. Minimum LSA arrival 1 secs
LSA group pacing timer 240 secs
Interface flood pacing timer 33 msecs
Retransmission pacing timer 66 msecs
Number of external LSA 0. Checksum Sum 0x000000
Number of areas in this router is 1. 1 normal 0 stub 0 nssa
Reference bandwidth unit is 100 mbps
  Area 1
    Number of interfaces in this area is 1
    SPF algorithm executed 1 times
    Number of LSA 3. Checksum Sum 0x02286D
    Number of DCbitless LSA 0
    Number of indication LSA 0
    Number of DoNotAge LSA 0
    Flood list length 0
```

IPv6 Troubleshooting Example 3 - Cont.

Once area 1 is configured as stub on R3, the adjacency with R1 forms.

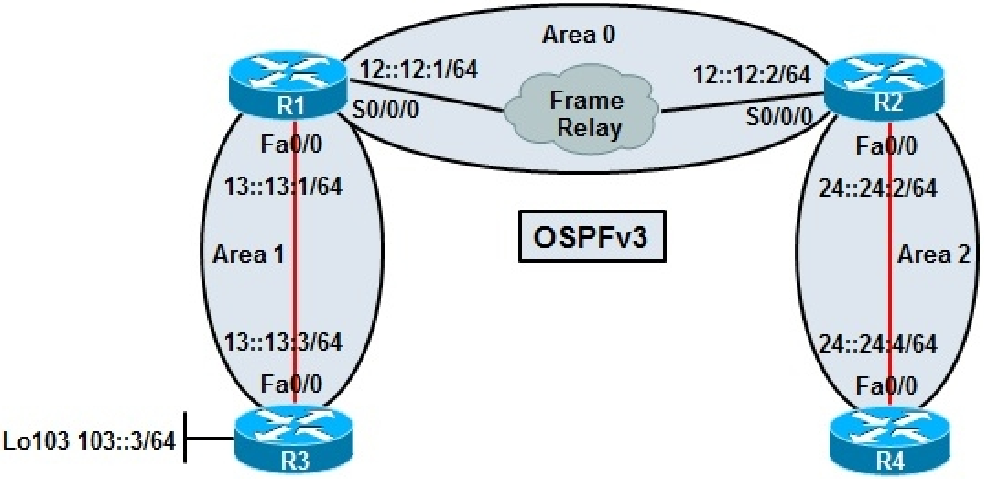

```
R3(config)# ipv6 router ospf 1
R3(config-rtr)# area 1 stub
R3(config-rtr)# end
R3#
Sep 27 20:17:38.367: %SYS-5-CONFIG_I: Configured from console by console
Sep 27 20:17:38.747: %OSPFv3-5-ADJCHG: Process 1, Nbr 172.16.101.1 on
FastEthernet0/0 from LOADING to FULL. Loading Done

R1# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 0/0/0 ms
```

IPv6 Troubleshooting Example 3 - Cont.

On R1, there is no DLCI mapping to the R2 S0/0/0 link local address, which prevents the exchange of hellos.

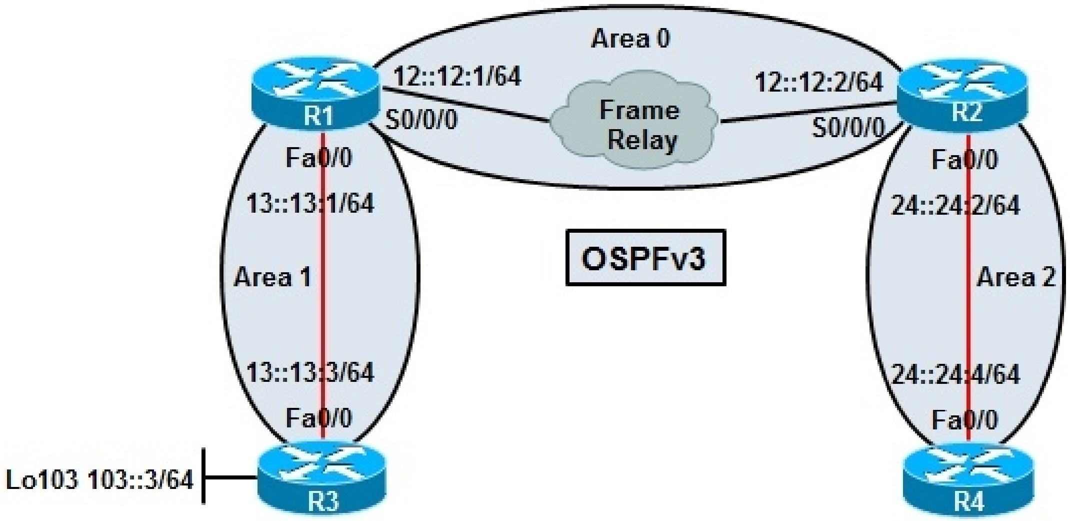

```
R1# show frame-relay map
Serial0/0/0 (up): ipv6 12::12:2 dlci 122(0x7A, 0x1CA0), static
              broadcast
              CISCO, status defined, active
```

IPv6 Troubleshooting Example 3 - Cont.

Configure DLCI-to-link local address mappings on R1 and R2 (R1 is shown here).

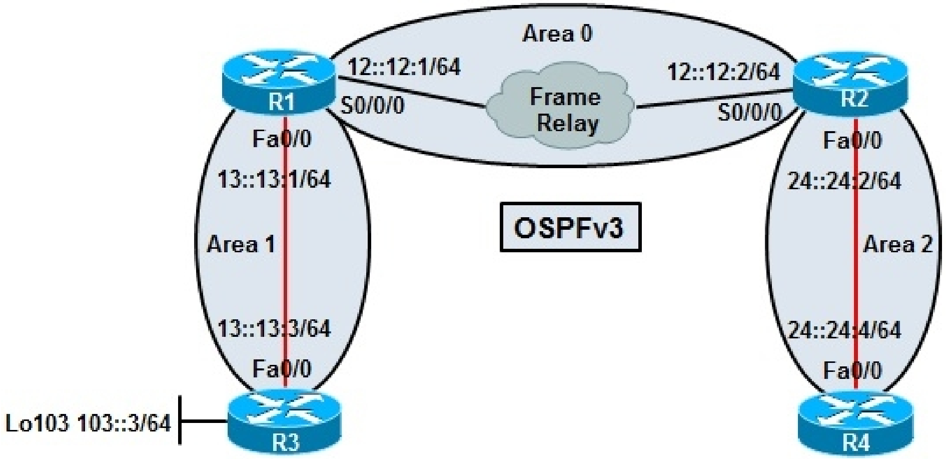

```
R2# show ipv6 interface brief
FastEthernet0/0                 [up/up]
    FE80::219:55FF:FE92:A442
    24::24:2
Serial0/0/0                     [up/up]
    FE80::219:55FF:FE92:A442
    11::1:2
    12::12:2

R1(config)# int s0/0/0
R1(config-if)# frame-relay map ipv6 FE80::219:55FF:FE92:A442 122 broadcast
R1(config-if)# end
```

IPv6 Troubleshooting Example 3 - Cont.

R2 Fa0/0 is defined in area 2.

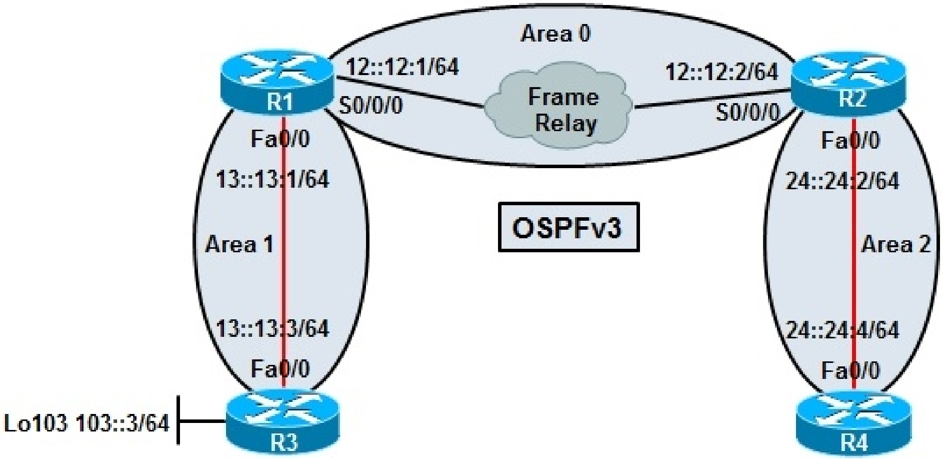

```
R2# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::219:55FF:FE92:A442, Interface ID 3
  Area 2, Process ID 1, Instance ID 0, Router ID 172.16.102.1
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State DR, Priority 1
  Designated Router (ID) 172.16.102.1, local address FE80::219:55FF:FE92:A442
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:04

<output omitted>
```

IPv6 Troubleshooting Example 3 - Cont.

R4 Fa0/0 is defined in area 0.

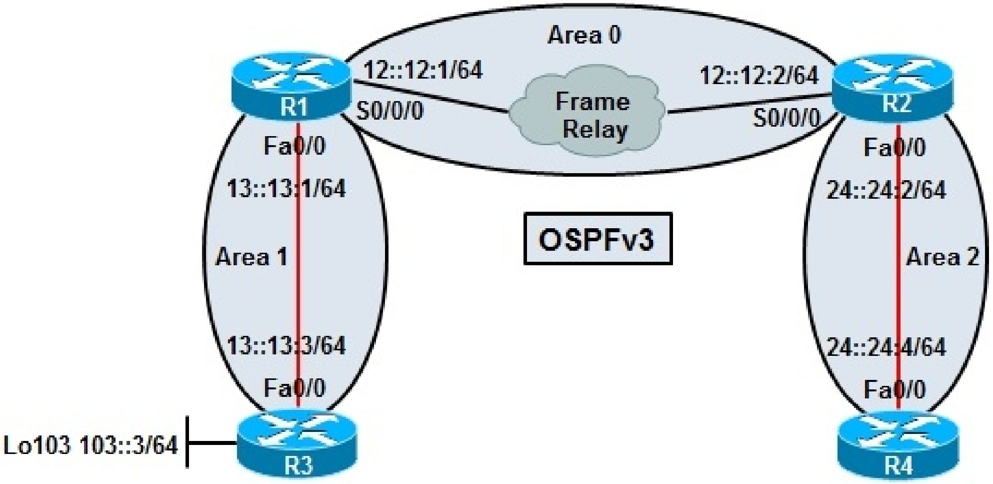

```
R4# show ipv6 ospf interface
FastEthernet0/0 is up, line protocol is up
  Link Local Address FE80::219:55FF:FEE0:F04, Interface ID 3
  Area 0, Process ID 1, Instance ID 0, Router ID 104.104.104.104
  Network Type BROADCAST, Cost: 1
  Transmit Delay is 1 sec, State WAITING, Priority 1
  No designated router on this network
  No backup designated router on this network
  Timer intervals configured, Hello 10, Dead 40, Wait 40, Retransmit 5
  Hello due in 00:00:03

<output omitted>
```

IPv6 Troubleshooting Example 3 - Cont.

Correct area number on R4 is configured and connectivity Is restored.

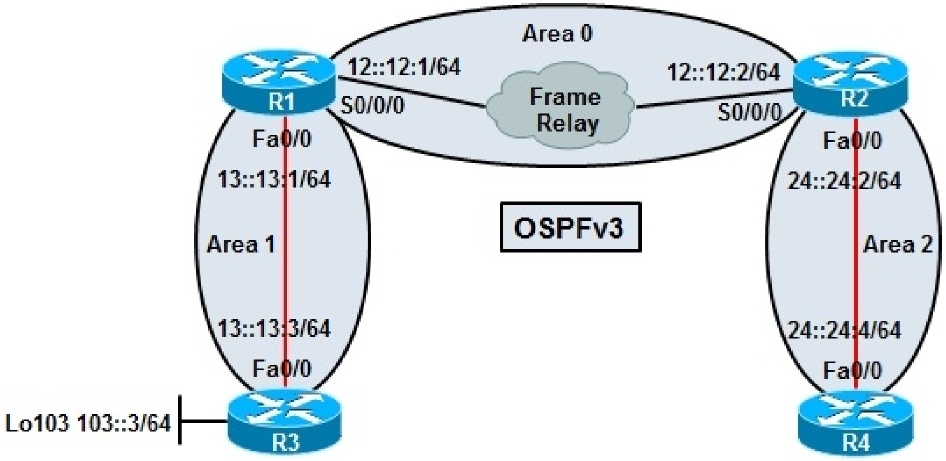

```
R4(config)# interface Fa0/0
R4(config-if)# ipv6 ospf 1 area 2
R4(config-if)# end
R4#

R4# ping 103::3
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 103::3, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 56/58/60
```

IPv6 Troubleshooting Example 4: OSPFv3 over 6to4 Tunnel

- Two islands of RIPng are connected through an OSPFv3 domain across the 6to4 tunnel.
- After a readdressing exercise, some traffic cannot traverse the tunnel.

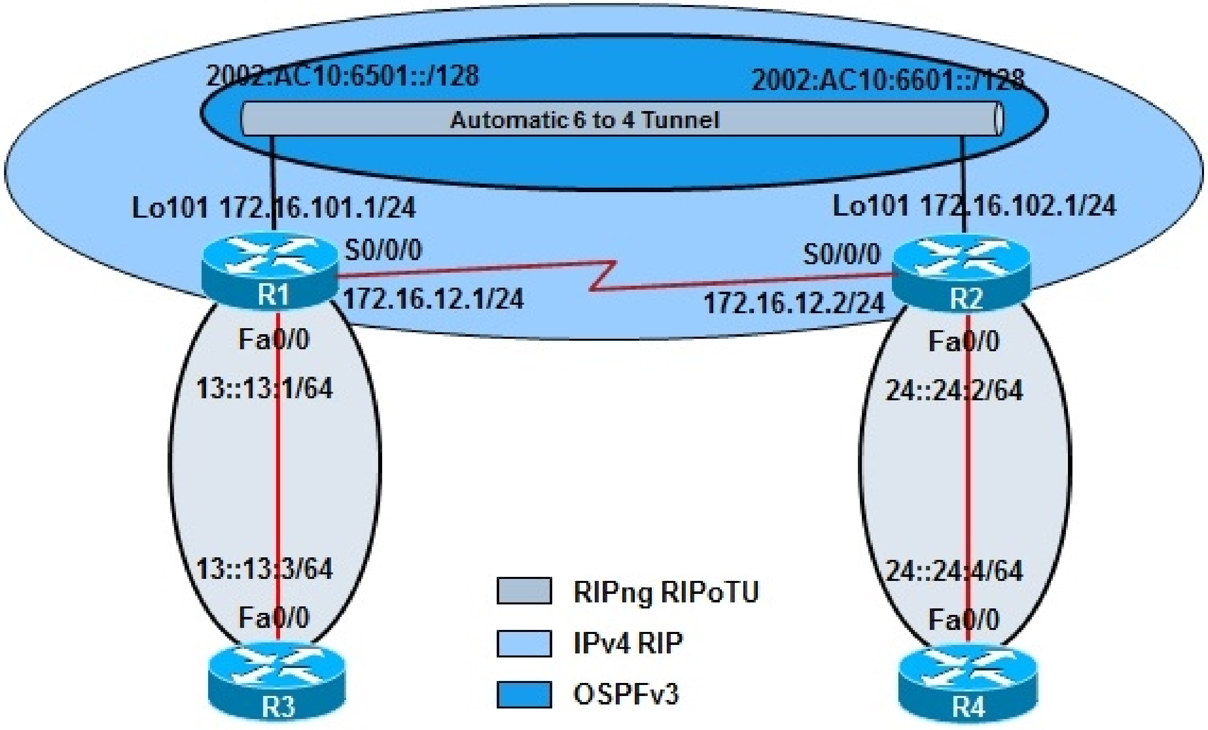

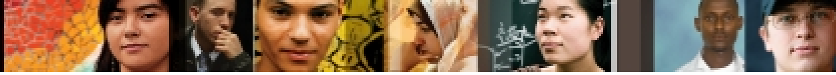

IPv6 Troubleshooting Example 4 – Cont.

The **ping** from R1 to R2 tunnel endpoint succeeds, and the output of the **debug tunnel** command shows the encapsulation and decapsulation.

```
R1# debug tunnel
Tunnel Interface debugging is on

R1# ping 2002:AC10:6601::2
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2002:AC10:6601::2, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 32/33/36 ms

R1# *Aug 24 14:39:52.863: Tunnel 0: count tx, adding 20 encap bytes
*Aug 24 14:39:52.895: Tunnel 0: IPv6/IP to classify 172.16.102.1 -> 172.16.101.1
(tbl=0, “default” len=120 ttl=254 tos=0x0)
*Aug 24 14:39:52.895: Tunnel 0: IPv6/IP (PS) to decaps 172.16.102.1 ->
172.16.101.1 (tbl=0, “default” len=120 ttl=254)
*Aug 24 14:39:52.895: Tunnel 0: decapsulated IPv6/IP packet (len 120)
*Aug 24 14:39:52.895: Tunnel 0 count tx, adding 20 encap bytes
*Aug 24 14:39:52.931: Tunnel 0: IPv6/IP to classify 172.16.102.1 -> 172.16.101.1
(tbl=0, “default” len=120 ttl=254 tos=0x0)
*Aug 24 14:39:52.931: Tunnel 0: IPv6/IP (PS) to decaps 172.16.102.1 ->
172.16.101.1 (tbl=0, “default” len=120 ttl=254)

<output omitted>
```

IPv6 Troubleshooting Example 4 – Cont.

R1 cannot ping R4’s Fa0/0 IPv6 address 24::24:2.

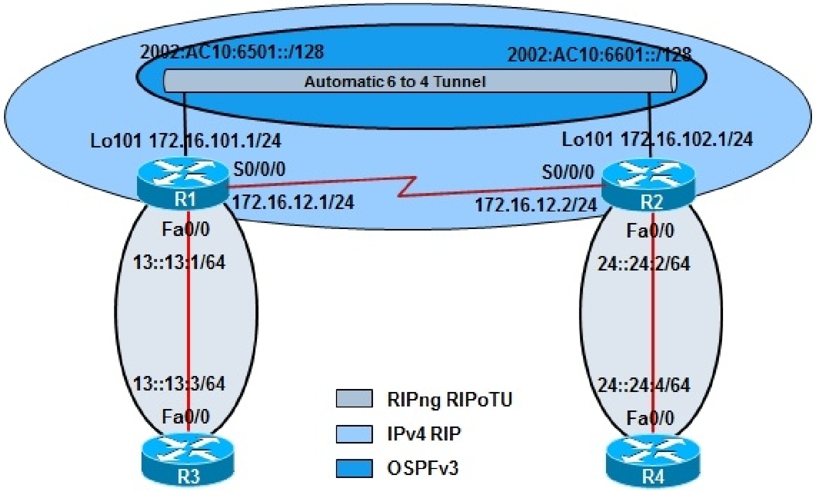

```
R1# ping 24::24:4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24::24:4, timeout is 2 seconds:
*Aug 24 14:41:12.927: Tunnel0 count tx, adding 20 encap bytes.
*Aug 24 14:41:14.899: Tunnel0 count tx, adding 20 encap bytes.
*Aug 24 14:41:14.927: Tunnel0 count tx, adding 20 encap bytes.
*Aug 24 14:41:16.927: Tunnel0 count tx, adding 20 encap bytes.
*Aug 24 14:41:18.927: Tunnel0 count tx, adding 20 encap bytes.

Success rate is 0 percent
```

IPv6 Troubleshooting Example 4 – Cont.

R1’s IPv6 routing table shows only connected and static routes.

```
R1# show ipv6 route
IPv6 Routing Table - 17 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static Route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea, IS - ISIS summary
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
       ON1 - OSPF NSSA ext 1, ON2 - OSPF NSSA ext 2
S  ::/0 [1/0]
   via ::, Tunnel0
C 11::/64 [0/0]
   via ::, Serial0/0/0
L 11::11:1/128 [0/0]
   via ::, Serial0/0/0
C 12::/64 [0/0]
   via ::, Serial0/0/0
L 12::12:1/128 [0/0]
   via ::, Serial0/0/0
C 13::/64 [0/0]
   via ::, FastEthernet0/0
L 13::13:1/128 [0/0]
   via ::, FastEthernet0/0
LC 2002:AC10:6501::1/128 [1/0]
   via ::, Tunnel0
S 2002:AC10:6610::1/128 [1/0]
   via Tunnel0
L FE80::/10 [0/0]
   via ::, Null0
L FF00::/8 [0/0]
   via ::, Null0
```

IPv6 Troubleshooting Example 4 – Cont.

The invalid IPv6 static route for the 6to4 tunnel is corrected but R1 still cannot ping R2’s Fa0/0 interface.

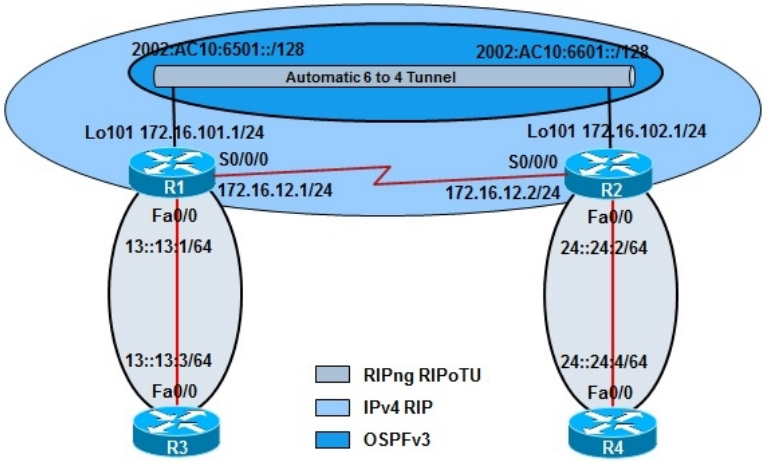

```
R1(config)# no ipv6 route 2002:AC10:6610::1/128 tun0
R1(config)#
R1(config)# ipv6 route 2002:AC10:6601::2/128 tun0
R1(config)#
R1(config)# end

R1# ping 24::24:4
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 24::24:24, timeout is 2 seconds:
.....
Success rate is 0 percent (0/5)
R1#
```

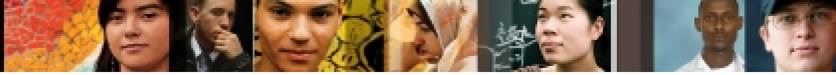

IPv6 Troubleshooting Example 4 – Cont.

Checking OSPFv3 configuration on R1 with **show** and **debug** commands.

```
R1# show ipv6 protocols
IPv6 Routing Protocol is “connected”
IPv6 Routing Protocol is “static”
IPv6 Routing Protocol is “rip RIPoTU”
  Interfaces:
    FastEthernet0/0
Redistribution:
  None
IPv6 Routing Protocol is “ospf 1”
  Interfaces (Area 0):
    Tunnel0
  Redistribution:
    None
R1#
R1#show ipv6 ospf neighbors

R1#
R1# debug ipv6 ospf hello
OSPFv3 hello events debugging is on
R1#
R1#
*Aug 24 14:59:53.247: OSPFv3: Send hello to FF02::5 area 0 on Tunnel0 from
FE80::AC10:6501 interface ID 13
```

IPv6 Troubleshooting Example 4 – Cont.

Checking OSPFv3 on R2 with **debug** command.

```
R2# debug ipv6 ospf hello
OSPFv3 hello events debugging is on
R2#
*Aug 24 15:07:04.635: OSPFv3: Send hello to FE80::219:56FF:FE2C:9856 area 0 on
Serial0/0/0 from FE80::219:55FF:FE92:A442 interface ID 5
*Aug 24 15:07:04.755: OSPFv3: Send hello to FE80::219:56FF:FE2C:9856 area 0 on
Serial0/0/0 from FE80::219:55FF:FE92:A442 interface ID 5
*Aug 24 15:07:04.875: OSPFv3: Send hello to FE80::219:56FF:FE2C:9856 area 0 on
Serial0/0/0 from FE80::219:55FF:FE92:A442 interface ID 5
*Aug 24 15:07:04.995: OSPFv3: send hello to FF02::5 area 0 on Tunnel0 from
FE80::AC10:6601 interface ID 13
*Aug 24 15:07:04.995: OSPFv3: Send hello to FE80::219:56FF:FE2C:9856 area 0 on
Serial0/0/0 from FE80::219:55FF:FE92:A442 interface ID 5
*Aug 24 15:07:05.115: OSPFv3: Send hello to FE80::219:56FF:FE2C:9856 area 0 on
Serial0/0/0 from FE80::219:55FF:FE92:A442 interface ID 5
```

IPv6 Troubleshooting Example 4 – Cont.

The OSPF **neighbor** command is not allowed on the tunnel interface.

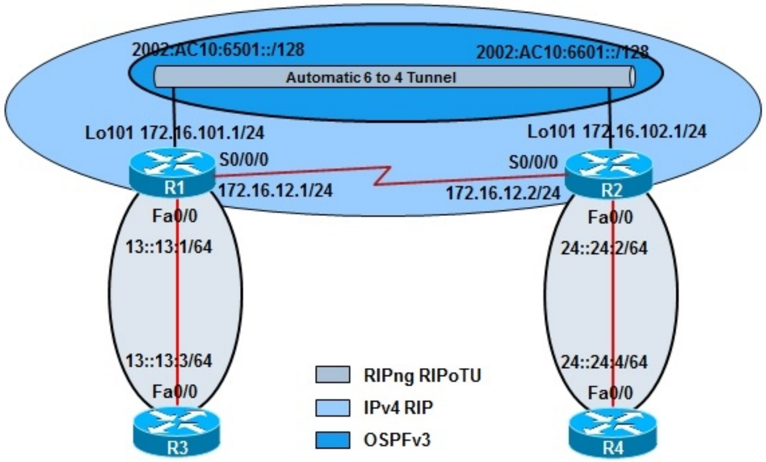

```
R1(config)# no ipv6 route 2002:AC10:6610::1/128 tun0
R1(config)# interface tunnel 0
R1(config-if)# ipv6 ospf neigh FE80::AC10:6601
R1(config-if)# end
R1#
R1#
*Aug 24 15:19:16.219: %OSPFv3-4-CFG_NBR_INVAL_NET_TYPE: Can not use configured
neighbor: neighbor command is allowed only on NBMA and point-to-multipoint
networks
*Aug 24 15:19:16.343: %SYS-5-CONFIG_I: Configured from console by console
```

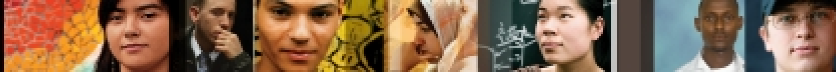

IPv6 Troubleshooting Example 4 - Summary

- In this example, the tunnel was configured correctly, and it was operational
- However, it was routing through the tunnel that was failing.
- In this example, OSPF was supposed to be running across the tunnel but it was not.
- If using a dynamic routing protocol between R1 and R2 is mandatory, the only option is BGP.
- BGP establishes peering based on TCP sessions that use unicasts.
- Those unicasts can be the 6to4 IPv6 address of the other side of the tunnel.
- If BGP is not used, the only way to make this work is with static routes.
- As seen earlier, R1 had a faulty static route, which was fixed, and a static default.
- Adding a proper static route on R2 will restore connectivity across the network.
- Only BGP would have worked in this scenario for dynamic routing across the 6to4 tunnel.

Cisco | Networking Academy®

Mind Wide Open™