Troubleshooting of Switching, HSRP and Addressing Services

CCNP TSHOOT: Maintaining and Troubleshooting IP Networks

Cisco | Networking Academy®
Mind Wide Open™

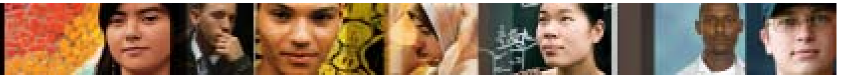

Lecture 1: Objectives

- LAN switch operation
- Troubleshooting of:
  - VLANs
  - STP and Etherchannel
  - Inter-VLAN routing
  - HSRP, VRRP, and GLBP
  - NAT/PAT
  - DHCP

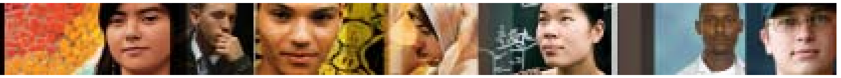

Review

- Before you start to troubleshoot, make sure you know the operation of the following protocols and functions:
  - LAN switch operation
  - VLANs
  - Spanning-Tree Protocol (STP)
  - Etherchannel
  - Inter-VLAN routing
  - First Hop Redundancy Protocols (HSRP, VRRP and GLBP)
  - Addressing Services (NAT/PAT and DHCP)

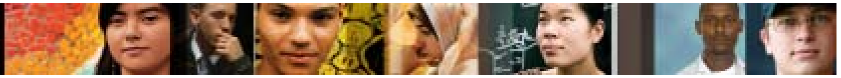

LAN Switch Operation

Issues that could cause the communication to fail:

- Physical problems
- Bad, missing, or miswired cables
- Bad ports
- Power failure
- Device problems
- Software bugs
- Performance problems
- Misconfiguration
- Missing or wrong VLANs
- Misconfigured VTP settings
- Wrong VLAN setting on access ports
- Missing or misconfigured trunks
- Native VLAN mismatch
- VLANs not allowed on trunk

Verifying Layer 2 Forwarding

Common findings when following the path of the frames through the switches:

- **Frames are not received on the correct VLAN**: This could point to VLAN or trunk misconfiguration as the cause of the problem.
- **Frames are received on a different port than you expected**: This could point to a physical problem, spanning tree issues, a native VLAN mismatch or duplicate MAC addresses.
- **The MAC address is not registered in the MAC address table:** This tells you that the problem is most likely upstream from this switch. Investigate between the last point where you know that frames were received and this switch.

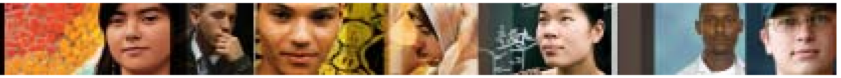

Verifying Layer 2 Forwarding – Cont.

Useful Layer 2 diagnostic commands:

- **`show mac-address-table:`** Shows learned MAC addresses and corresponding port and VLAN associations.
- **`show vlan`**: Verifies VLAN existence and port-to-VLAN associations.
- **`show interfaces trunk`**: Displays all interfaces configured as trunks, VLANs allowed and what the native VLAN is.
- **`show interfaces switchport`**: Provides a summary of all VLAN related information for interfaces.
- **`show platform forward`** *`interface`*: Used to determine how the hardware would forward a frame.
- **`traceroute mac`**: Provides a list of switch hops (layer 2 path) that a frame from a specified source MAC address to a destination MAC address passes through. CDP must be enabled on all switches in the network for this command to work.
- **`traceroute mac ip:`** Displays Layer 2 path taken between two IP hosts.

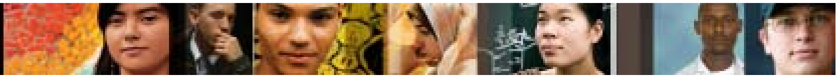

Spanning Tree Failures

- STP is a reliable but not an absolutely failproof protocol.
- If STP fails there are usually major negative consequences.
- With Spanning Tree, there are two different types of failures.
  - Type 1 - STP may erroneously block certain ports that should have gone to the forwarding state. You may lose connectivity to certain parts of the network, but the rest of the network is unaffected.
  - Type 2 - STP erroneously moves one or more ports to the Forwarding state. The failure is more disruptive as bridging loops and broadcast storms can occur.

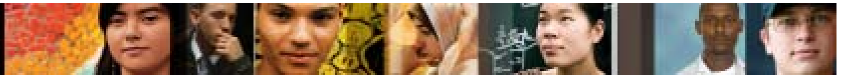

Spanning Tree Failures – Cont.

- Type 2 failures can cause these symptoms.
  - The load on all links in the switched LAN will quickly start increasing.
  - Layer 3 switches and routers report control plane failures such as continual HSRP, OSPF and EIGRP state changes or that they are running at a very high CPU utilization load.
  - Switches will experience very frequent MAC address table changes.
  - With high link loads and CPU utilization devices typically become unreachable, making it difficult to diagnose the problem while it is in progress.
- Eliminate topological loops and troubleshoot issues.
  - Physically disconnect links or shut down interfaces.
  - Diagnose potential problems.
  - A unidirectional link can cause STP problems. You may be able to identify and remove a faulty cable to correct the problem.

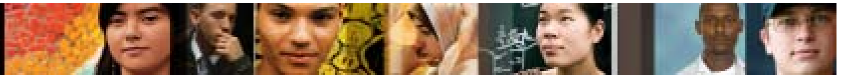

Spanning Tree Failures – Cont.

Using the **`show etherchannel 1 detail`** command

```
DSW2# show etherchannel 1 detail
Group state = L2
Ports: 2   Maxports = 8
Port-channels: 1 Max Port-channels = 1
Protocol:    -
Minimum Links: 0
Ports in the group:
-------------------
Port: Fa0/5
------------

Port state    = Up Cnt-bndl Suspend Not-in-Bndl
Channel group = 1           Mode = On              Gcchange = -
Port-channel  = null        GC   =   -             Pseudo port-channel = Po1
Port index    = 0           Load = 0x00            Protocol =    -

Age of the port in the current state: 0d:00h:25m:13s

Probable reason: vlan mask is different

<output omitted>
```

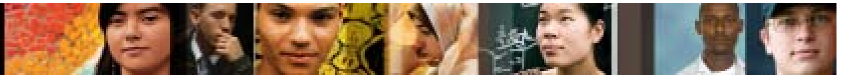

EtherChannel Problems

Three common EtherChannel problems:

1. Inconsistencies between the physical ports that are members of the channel
2. Inconsistencies between the ports on the opposite sides of the EtherChannel link
3. Uneven distribution of traffic between EtherChannel bundle members

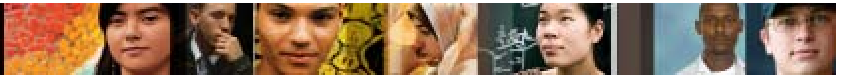

EtherChannel Diagnostic Commands

Using the **show etherchannel summary** command

```
DSW2# show etherchannel summary
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

Number of channel-groups in use: 2
Number of aggregators:           2

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-------------------------
1      Po1(SD)         -        Fa0/5(s)    Fa0/6(s)
2      Po2(SU)         -        Fa0/3(P)    Fa0/4(P)
```

EtherChannel Diagnostics

Using the **`show spanning-tree`** command to examine STP

```
ASW1# show spanning-tree vlan 17

MST0
  Spanning tree enabled protocol mstp
  Root ID    Priority    32768
             Address     001e.79a9.b580
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32768  (priority 32768 sys-id-ext 0)
             Address     001e.79a9.b580
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- ------------------------------
Fa0/7               Desg FWD 200000    128.9    P2p Edge
Po1                 Desg BLK 100000    128.56   P2p
Po2                 Desg BKN*100000    128.64   P2p Bound(PVST) *PVST_Inc
```

Troubleshooting Routers and Multi-Layer Switches

Sample Data Plane and Control Plane commands for routers and multi-layer switches

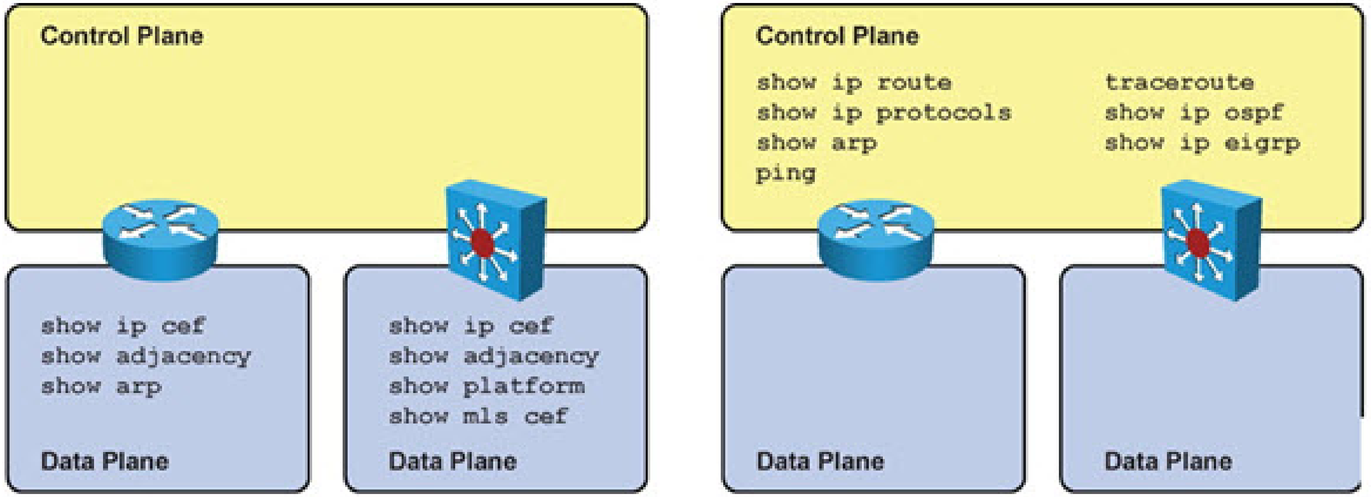

Control Plane

Data Plane

```
show ip cef
show adjacency
show arp
```

Data Plane

```
show ip cef
show adjacency
show platform
show mls cef
```

Control Plane

```
show ip route          traceroute
show ip protocols      show ip ospf
show arp               show ip eigrp
ping
```

Data Plane

Data Plane

Cisco 7206

Catalyst 6504

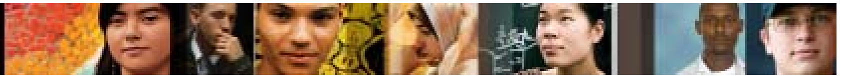

Troubleshooting Routers and Multi-Layer Switches – Cont.

Commands to check the CEF data structures for routers and multi-layer switches.

`show ip cef`

- Displays the content of the CEF FIB.
  - The FIB reflects the content of the routing table with all the recursive lookups resolved already and the output interface determined for each destination prefix.
  - The FIB also holds additional entries for directly connected hosts, the router's own IP addresses, and multicast and broadcast addresses.

`show adjacency`

- Displays the content of the CEF adjacency table.
  - This table contains preconstructed Layer 2 frame headers with all necessary fields already filled in. These frame headers are used to encapsulate the egress CEF-switched packets and deliver them to appropriate next hop devices..

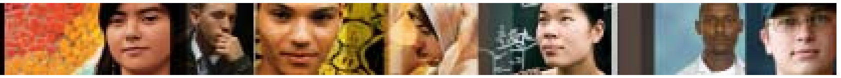

Troubleshooting Multi-layer Switches

Commands to check forwarding behavior of switches from the content of TCAM on Catalyst switches:

`show platform`

- On the Catalyst 3560, 3750 and 4500 platforms, the show platform family of commands can be used to obtain detailed information about the forwarding behavior of the hardware.

`show mls cef`

- On the Catalyst 6500 platform, the show mls cef family of commands can be used to obtain detailed information about the forwarding behavior of the hardware.

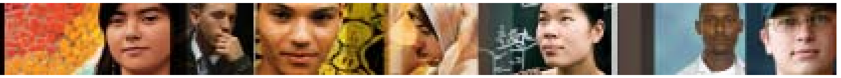

Checking SVI Status

Verifying the status of a VLAN and SVI

```
ASW1# show ip interfaces brief | exclude unassigned
Interface                   IP-Address      OK? Method Status            Protocol
Vlan128                     10.1.156.1      YES NVRAM  up                down


ASW1# show spanning-tree vlan 128
Spanning tree instance(s) for vlan 128 does not exist.


ASW1# show vlan id 128
VLAN id 128 not found in current VLAN database
```

Verifying HSRP Operation

Sample output from the **show standby brief** command

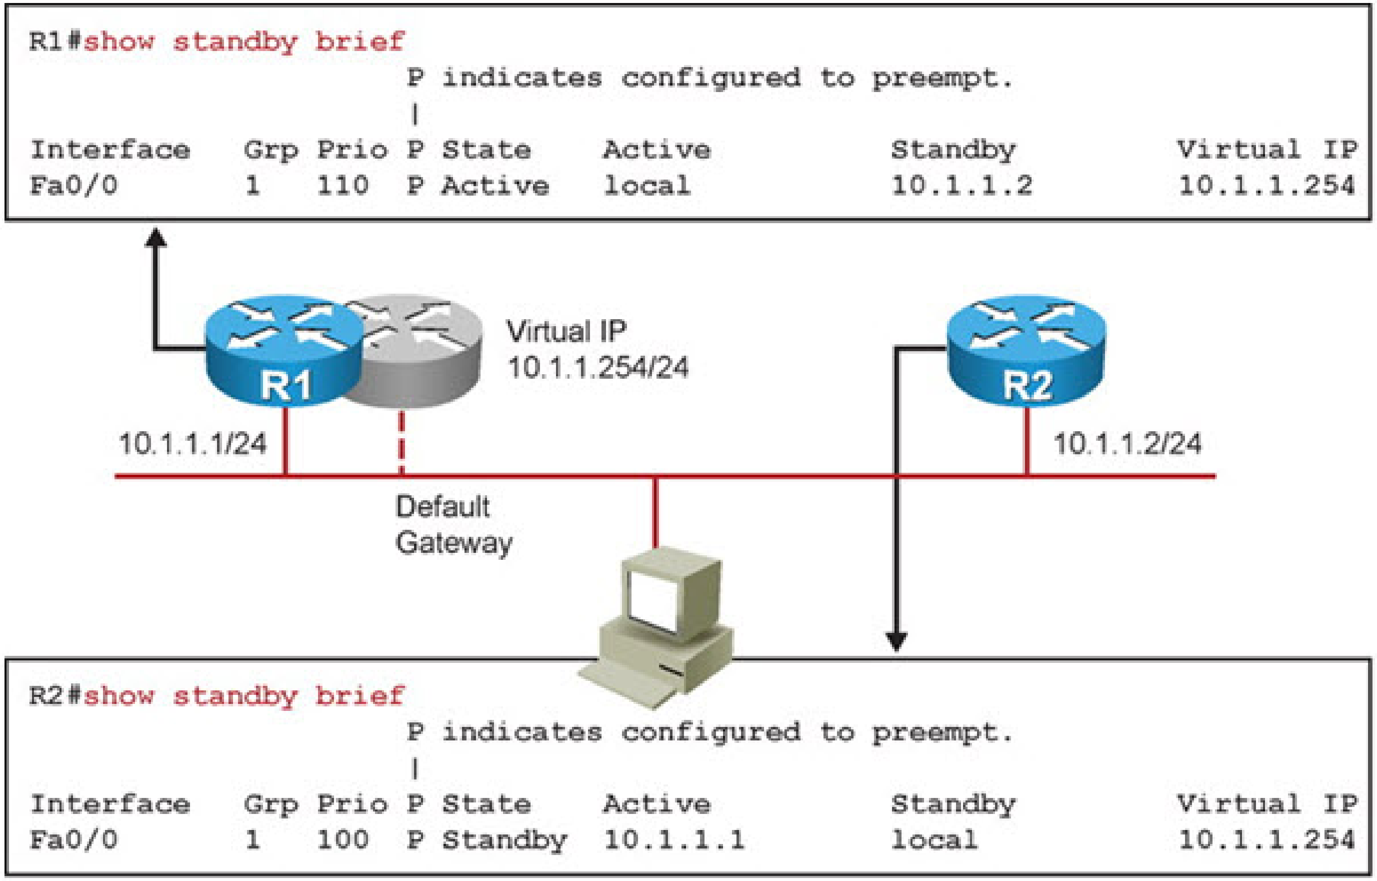

```
R1#show standby brief
                         P indicates configured to preempt.
                         |
Interface    Grp Prio P State    Active           Standby          Virtual IP
Fa0/0        1   110  P Active   local            10.1.1.2         10.1.1.254
```

```
R2#show standby brief
                         P indicates configured to preempt.
                         |
Interface    Grp Prio P State    Active           Standby          Virtual IP
Fa0/0        1   100  P Standby  10.1.1.1         local            10.1.1.254
```

Verifying HSRP Operation – Cont.

Sample output from the **`show standby`** ***`interface-id`*** command

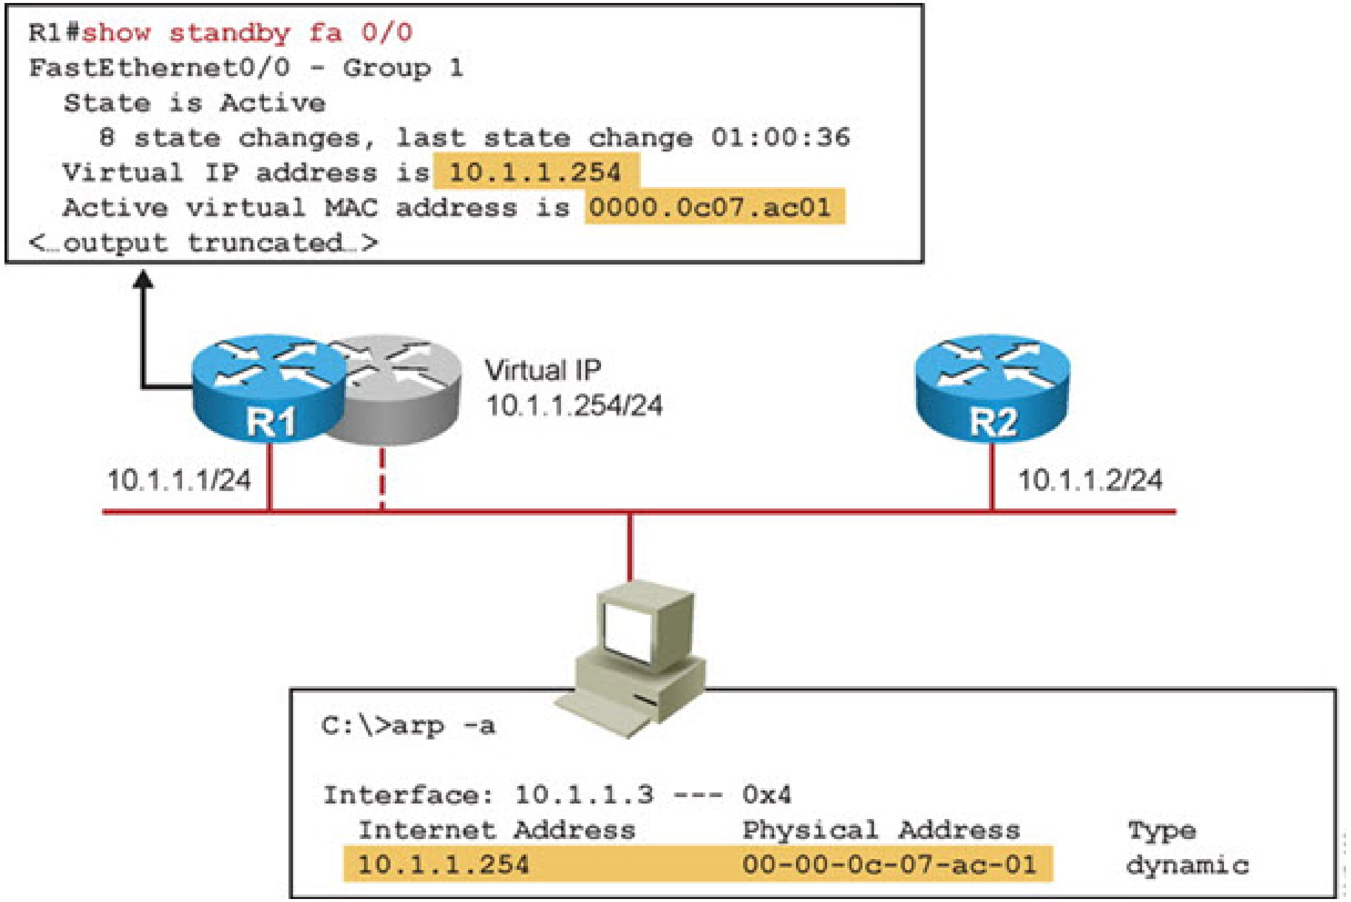

```
R1#show standby fa 0/0
FastEthernet0/0 - Group 1
  State is Active
    8 state changes, last state change 01:00:36
  Virtual IP address is 10.1.1.254
  Active virtual MAC address is 0000.0c07.ac01
<…output truncated…>
```

```
C:\>arp -a

Interface: 10.1.1.3 --- 0x4
  Internet Address      Physical Address      Type
  10.1.1.254            00-00-0c-07-ac-01     dynamic
```

Verifying HSRP Operation – Cont.

The interface of a router participating in HSRP is shutdown.

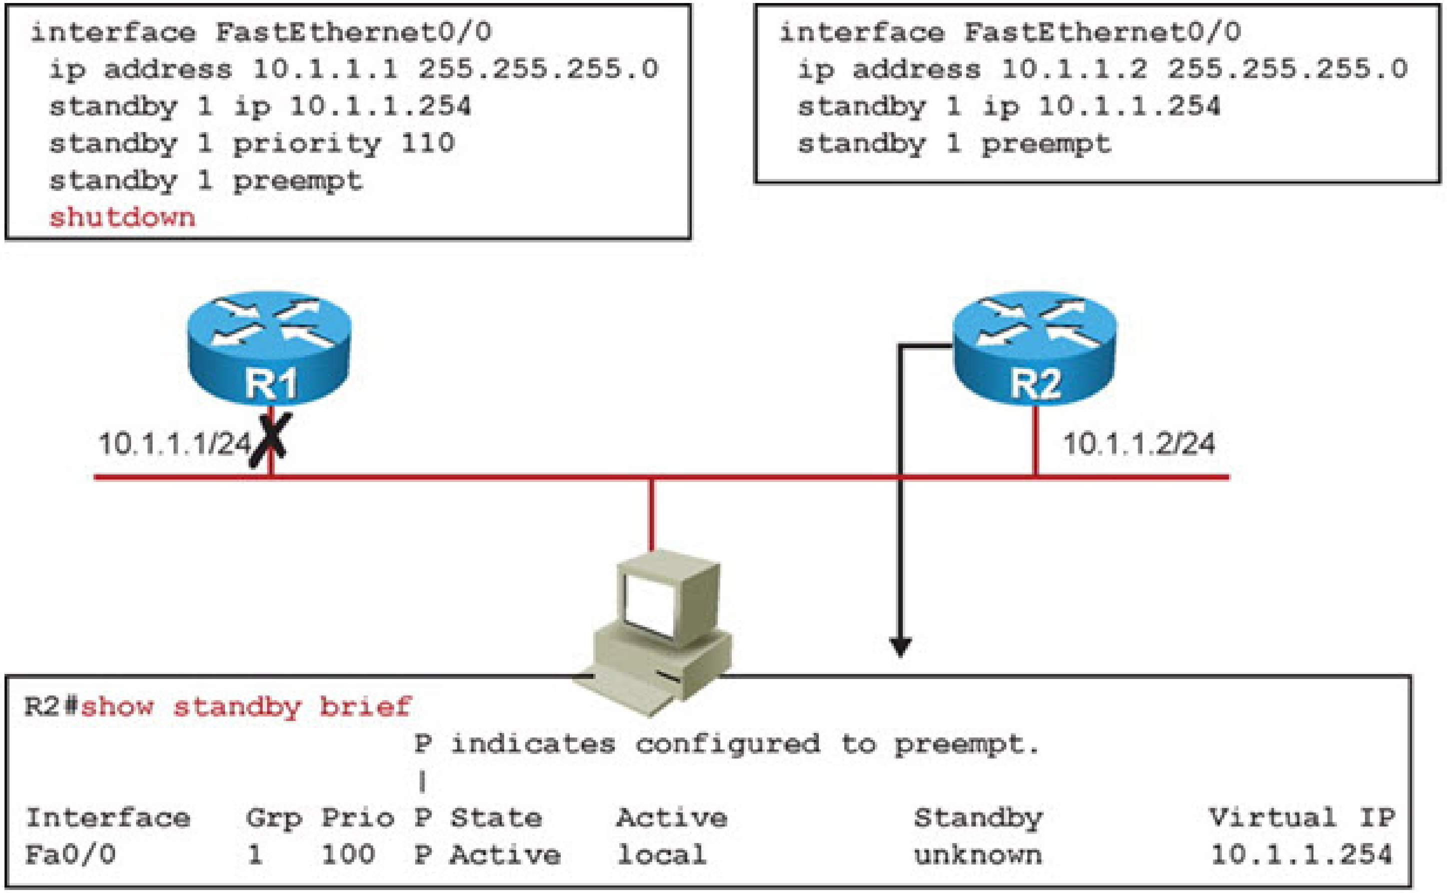

```
interface FastEthernet0/0
 ip address 10.1.1.1 255.255.255.0
 standby 1 ip 10.1.1.254
 standby 1 priority 110
 standby 1 preempt
 shutdown
```

```
interface FastEthernet0/0
 ip address 10.1.1.2 255.255.255.0
 standby 1 ip 10.1.1.254
 standby 1 preempt
```

```
R2#show standby brief
                         P indicates configured to preempt.
                         |
Interface    Grp Prio P State    Active          Standby         Virtual IP
Fa0/0        1   100  P Active   local           unknown         10.1.1.254
```

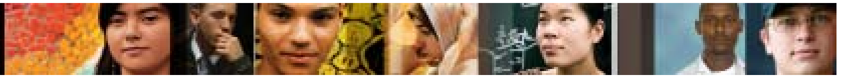

Verifying HSRP Operation – Cont.

While **debug standby terse** is enabled on R2, R1’s interface is enabled.

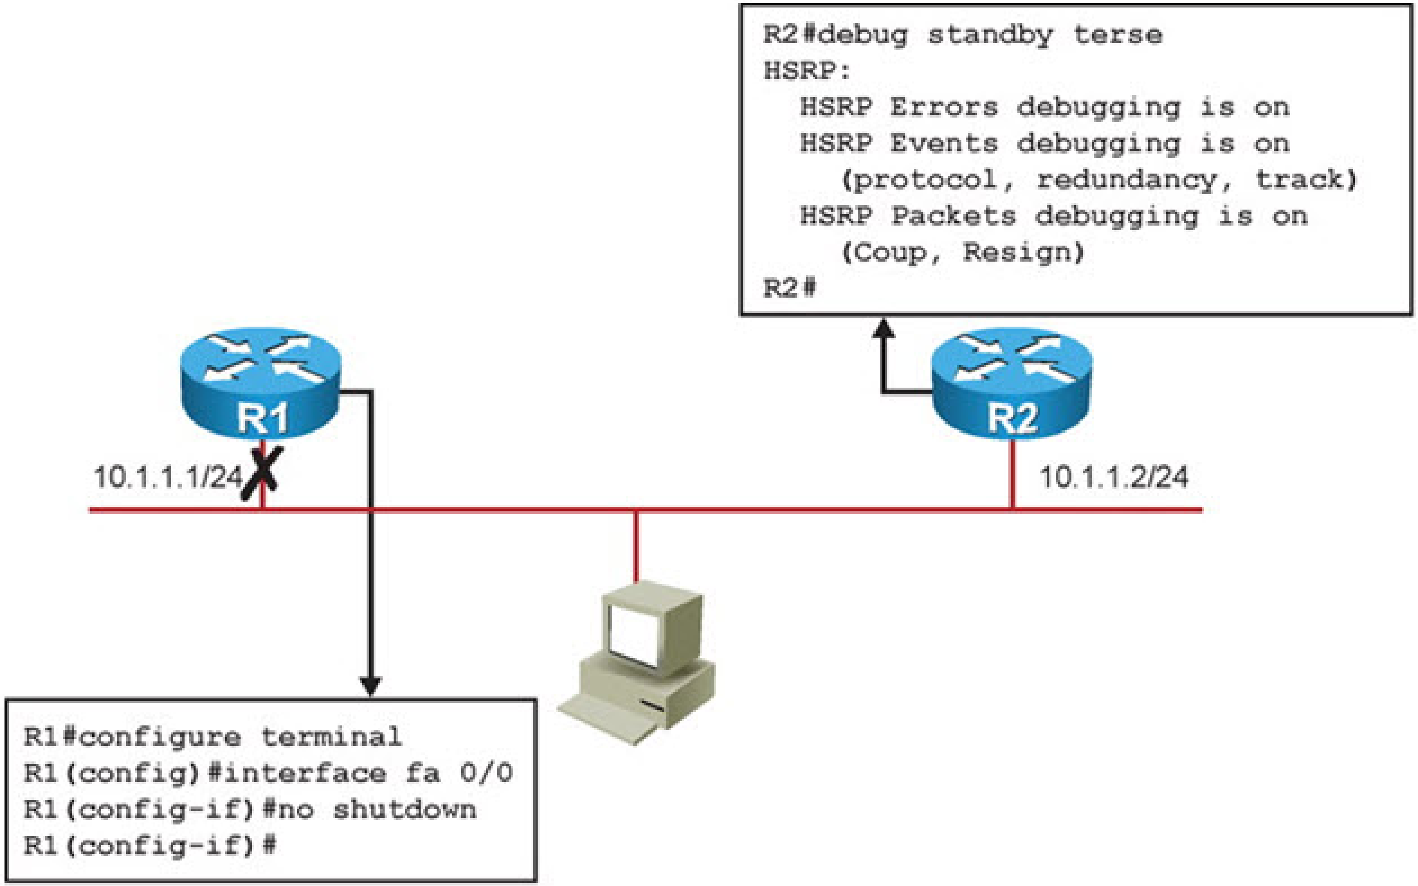

```
R2#debug standby terse
HSRP:
  HSRP Errors debugging is on
  HSRP Events debugging is on
    (protocol, redundancy, track)
  HSRP Packets debugging is on
    (Coup, Resign)
R2#
```

```
R1#configure terminal
R1(config)#interface fa 0/0
R1(config-if)#no shutdown
R1(config-if)#
```

Verifying HSRP Operation – Cont.

Output of **`debug standby terse`** on R2 as R1’s interface is enabled

```
R2#
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Coup   in  10.1.1.1 Listen  pri 110
vIP 10.1.1.254
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active: j/Coup rcvd from higher pri
router (110/10.1.1.1)
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active router is 10.1.1.1, was local
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Active -> Speak
*Mar  1 00:16:23.555: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Active
-> Speak
*Mar  1 00:16:23.555: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Active
-> Speak
*Mar  1 00:16:33.555: HSRP: Fa0/0 Grp 1 Speak: d/Standby timer expired
(unknown)
*Mar  1 00:16:33.555: HSRP: Fa0/0 Grp 1 Standby router is local
*Mar  1 00:16:33.555: HSRP: Fa0/0 Grp 1 Speak -> Standby
*Mar  1 00:16:33.555: %HSRP-5-STATECHANGE: FastEthernet0/0 Grp 1 state Speak -
> Standby
*Mar  1 00:16:33.559: HSRP: Fa0/0 Grp 1 Redundancy "hsrp-Fa0/0-1" state Speak
-> Standby
R2#
```

HSRP, VRRP, and GLBP Diagnostic Commands

Output of basic **show** commands for HSRP, VRRP, and GLBP

```
R1# show standby brief
                     P indicates configured to preempt.
                     |
Interface   Grp Prio P State    Active          Standby         Virtual IP
Fa0/0       1   110  P Active   local           10.1.1.2        10.1.1.254
…
R1# show vrrp brief
Interface         Grp Pri Time  Own Pre State   Master addr     Group addr
Fa0/0             1   110 3570       Y  Master  10.1.1.1        10.1.1.254
…
R1# show glbp brief
Interface   Grp  Fwd Pri State    Address         Active router   Standby
router
Fa0/0       1    -   110 Active   10.1.1.254      local           10.1.1.2
Fa0/0       1    1   -   Active   0007.b400.0101  local           -
Fa0/0       1    2   -   Listen   0007.b400.0102  10.1.1.2        -
```

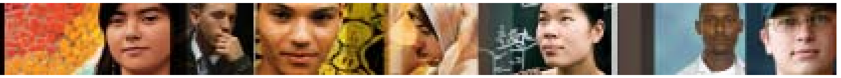

Troubleshooting NAT/PAT Issues

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 "to what those addresses translate", as packets go from IP NAT inside to IP NAT outside.
- 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.

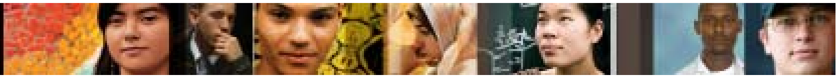

Troubleshooting NAT/PAT Issues – Cont.

The following commands can help determine if NAT is functioning correctly:

- **`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).

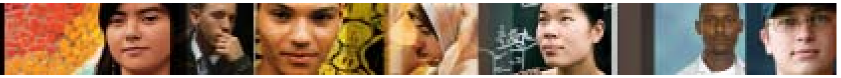

Troubleshooting NAT/PAT Issues – Cont.

Helpful NAT-related debug commands:

- **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.

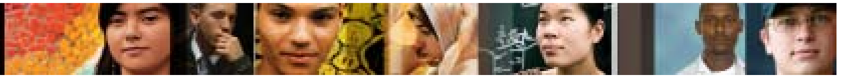

Troubleshooting NAT/PAT Issues – Cont.

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 the specified interface.
  - Will not generate debugging output for packets for a different interface.
  - First define the condition with the **`debug condition`** command. For example, define a condition of **`interface serial 0/0`**.
  - This definition means that all debug output will be limited to that particular interface.
  - 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.

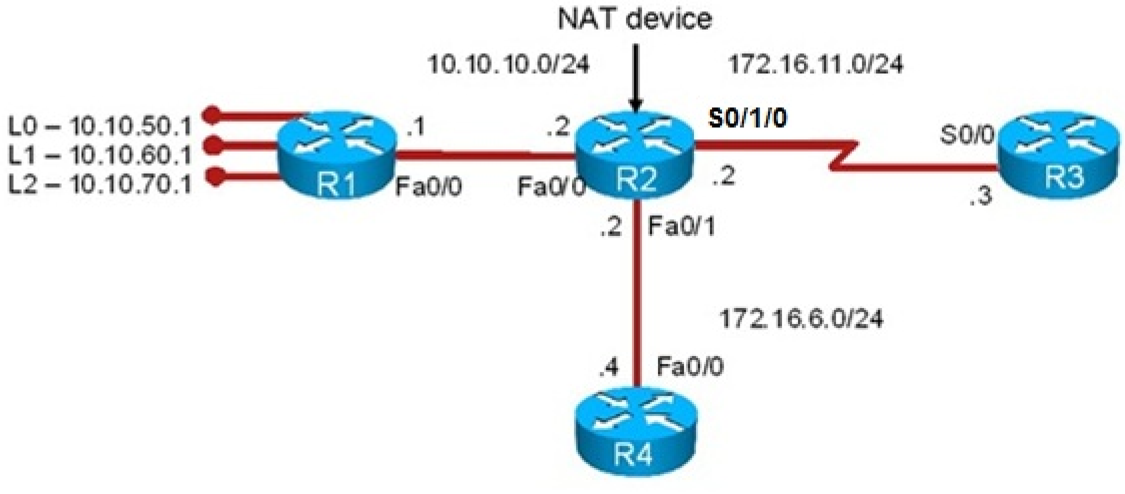

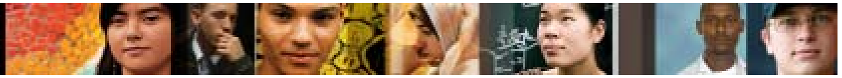

NAT/PAT Troubleshooting Example 1 – Cont.

```
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.

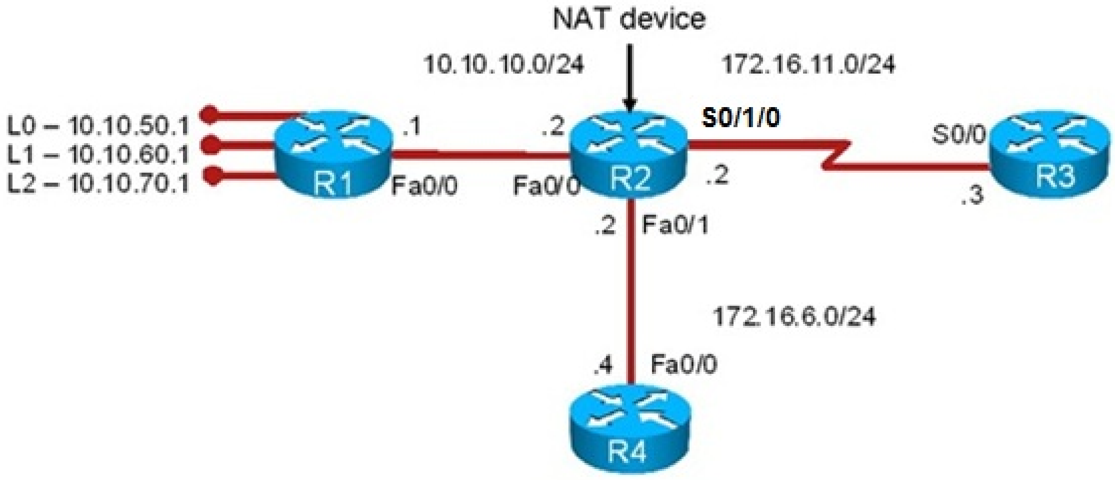

```
R2# sh ip nat translations
Pro  Inside global  Inside local     Outside local    Outside global
---  172.16.6.1     10.10.10.1         ---              ---
```

NAT/PAT Troubleshooting Example 1 – Cont.

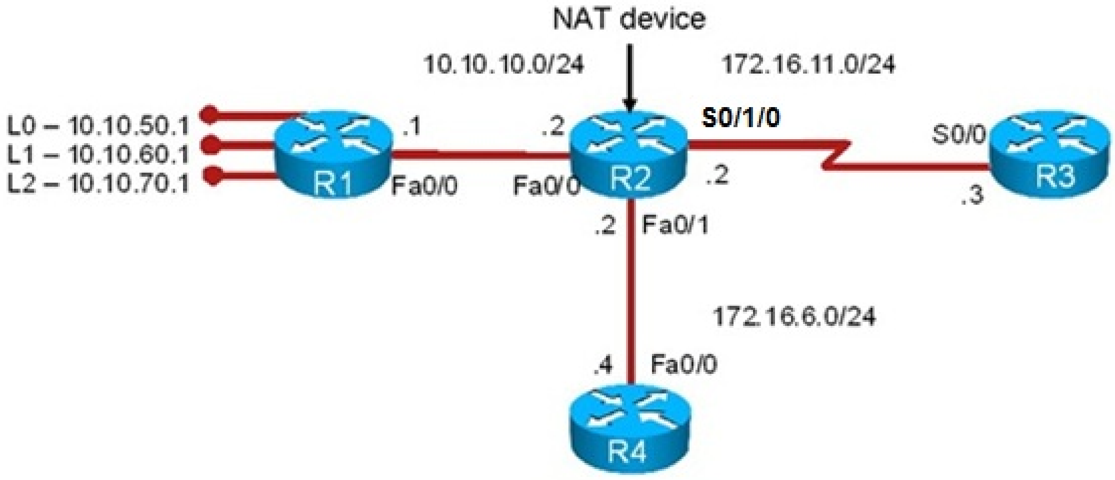

```
R3# debug ip icmp
ICMP packet debugging is on

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
```

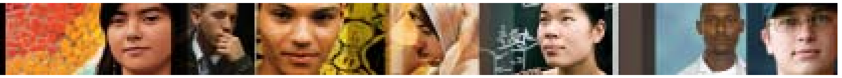

NAT/PAT Troubleshooting Example: – Cont.

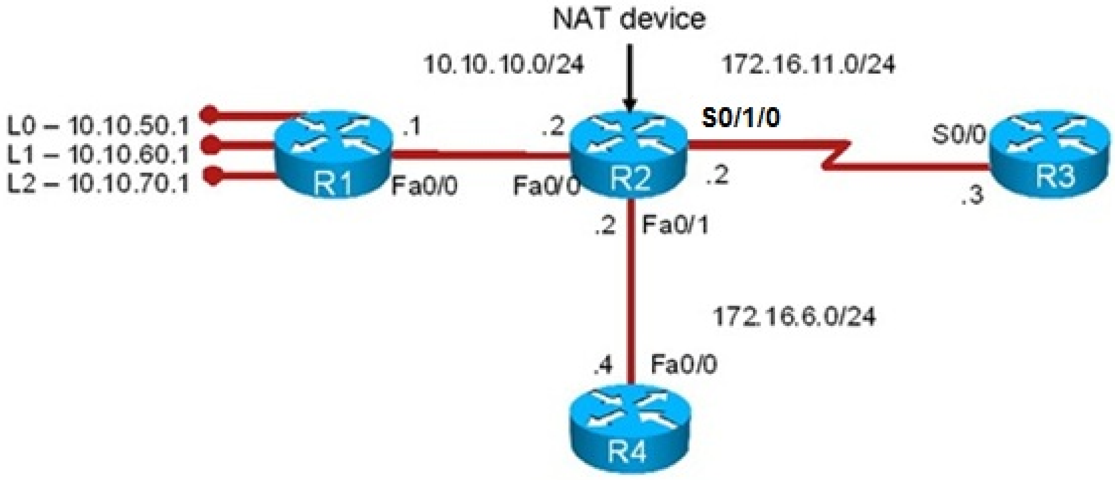

```
R3# show ip route 172.16.6.0 255.255.255.0
% Subnet not in table

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.
- They can accomplish connectivity from the R1 loopbacks.
- The risk management team recently performed an upgrade to router and firewall security policies.
- The routing protocol used is single-area OSPF.
- Goal to restore end-to-end connectivity and make sure SSH is operational to support management processes.

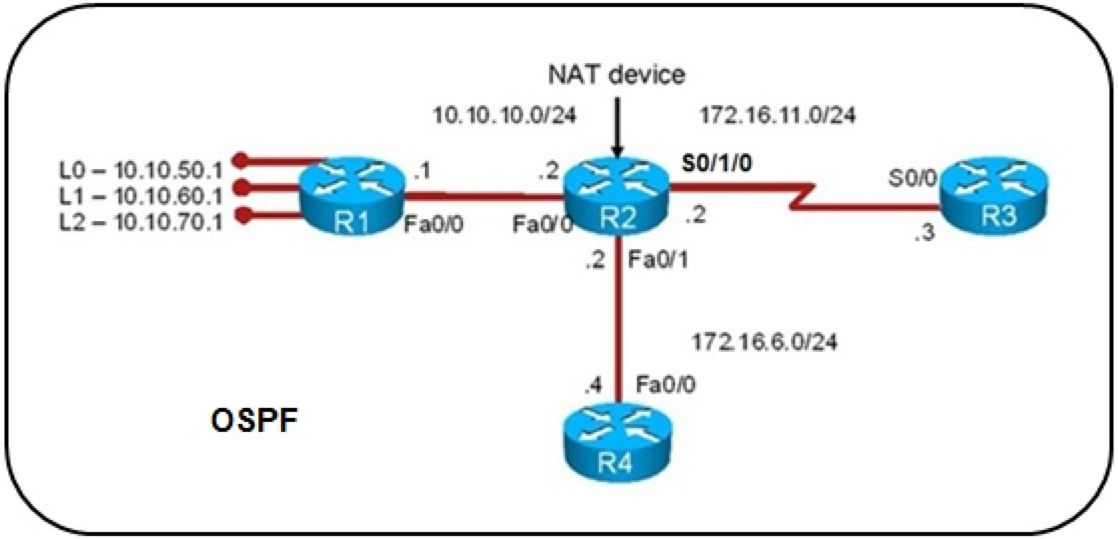

NAT/PAT Troubleshooting Example 2 – Cont.

- Extended ping and SSH results from R1 to R3

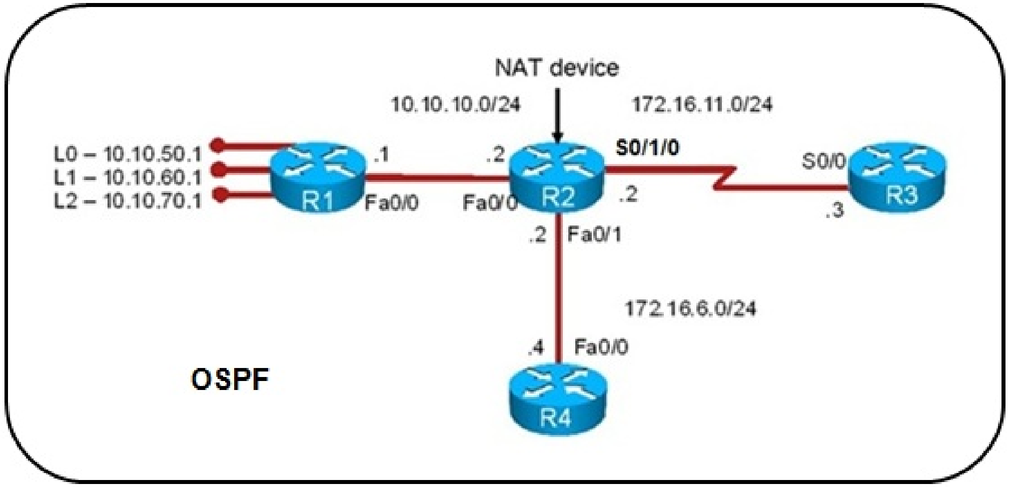

```
R1# ping 172.16.11.3 source 10.10.50.1
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 172.16.11.3, timeout is 2 seconds:
Packet sent with a source address of 10.10.50.1
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 1/1/4 ms

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
```

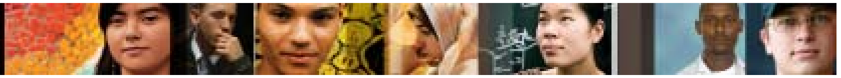

NAT/PAT Troubleshooting Example 2 – Cont.

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
```

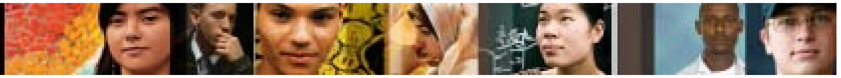

NAT/PAT Troubleshooting Example 2 – Cont.

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)
```

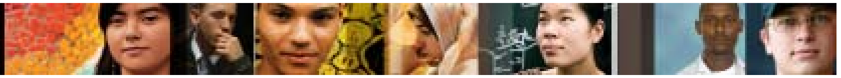

NAT/PAT Troubleshooting Example 2 – Cont.

Using **debug ip packet** while attempting SSH

```
R1# ssh -l user 172.16.11.3
% Connection refused by remote host
R1#
R3# debug ip packet
IP packet debugging is on
R3#
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
```

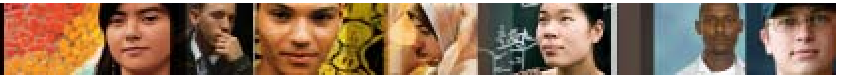

NAT/PAT Troubleshooting Example 2 – Cont.

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
```

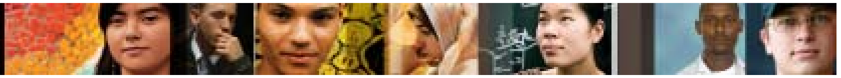

NAT/PAT Troubleshooting Example 2 – Cont.

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
```

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

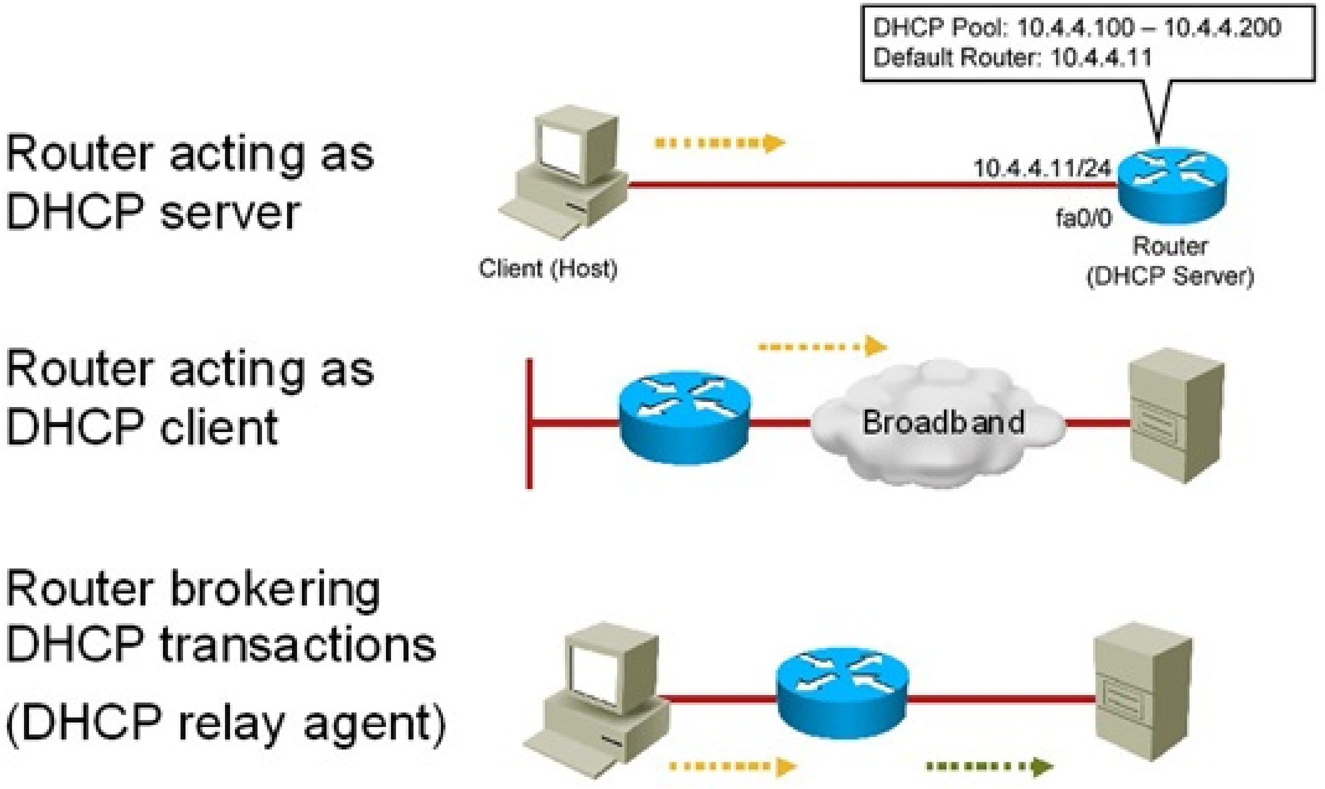

Router acting as DHCP server

Router acting as DHCP client

Router brokering DHCP transactions

(DHCP relay agent)

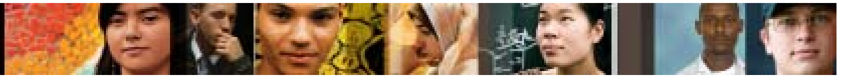

DHCP Troubleshooting Issues – Cont.

- 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.

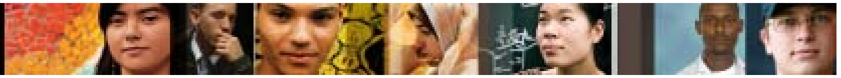

DHCP Troubleshooting Issues: DHCP Relay Agent

- The Cisco IOS command that makes a router a DHCP relay agent **is ip helper-address**.
- This is an interface configuration command that makes the router forward the BootP/DHCP requests from clients to the DHCP server.
- 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 makes the interface forward UDP broadcasts for six protocols (not just DHCP) to the IP address configured using the **`ip helper-address`** command.
  - 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 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.

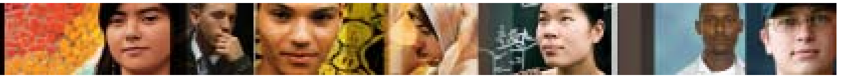

DHCP Troubleshooting Issues

- Troubleshooting can be related to DHCP security efforts.
- 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.

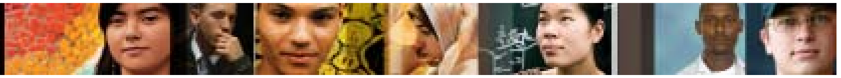

DHCP Troubleshooting Issues – Cont.

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 UDP port 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?

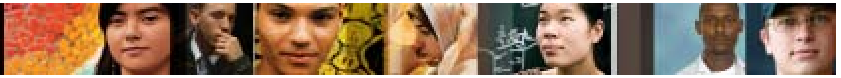

DHCP Troubleshooting Commands

- **`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.
- **`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

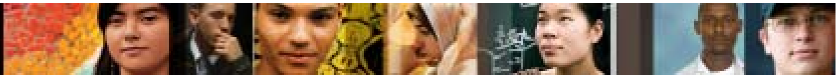

DHCP Troubleshooting Commands – Cont.

- **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.

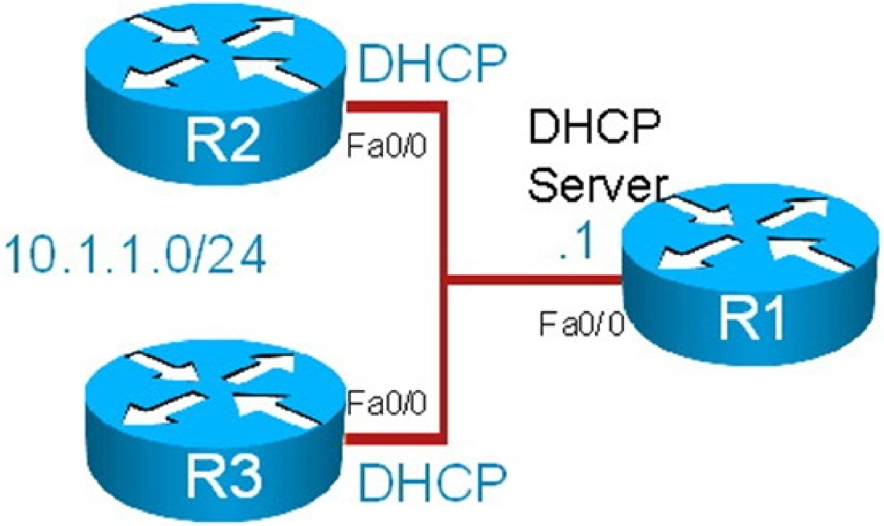

DHCP Troubleshooting Example 1 – Cont.

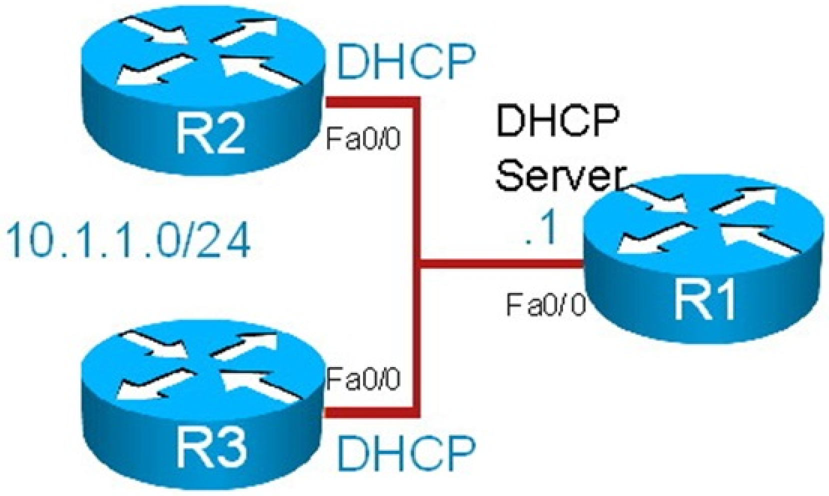

```
R2# show ip int brief
Interface          IP-Address   OK? Method Status                Protocol
FastEthernet0/0    unassigned   YES DHCP   up                    up
FastEthernet0/1    unassigned   YES NVRAM  administratively down down
Serial0/0/0        unassigned   YES NVRAM  administratively down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

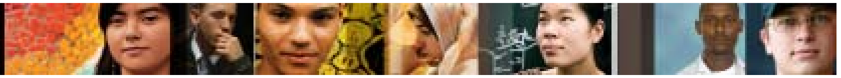

DHCP Troubleshooting Example 1 – Cont.

```
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.

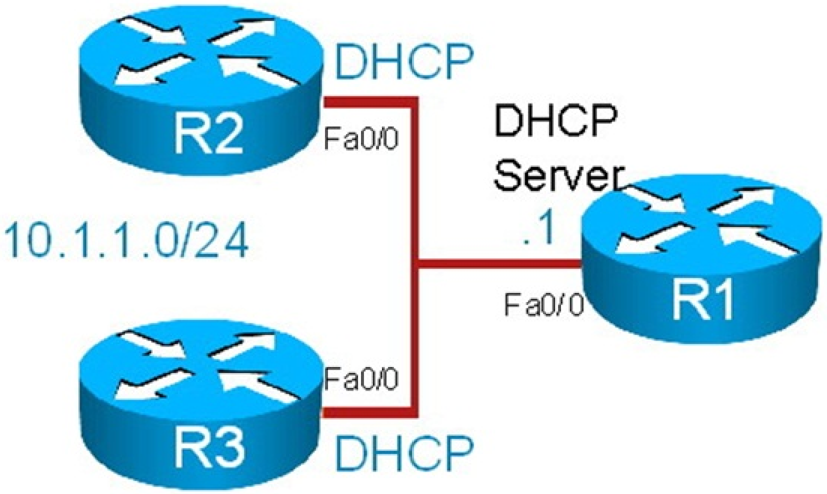

```
R1# show ip int brief
Interface         IP-Address   OK? Method Status                Protocol
FastEthernet0/0   10.1.1.1     YES manual up                    up
FastEthernet0/1   unassigned   YES NVRAM  administratively down down
Serial0/0/0       unassigned   YES NVRAM  administratively down down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

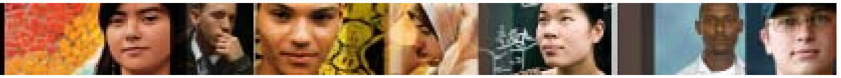

DHCP Troubleshooting Example 1 – Cont.

```
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
```

Chapter 4

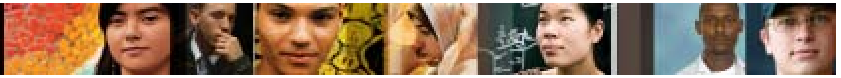

DHCP Troubleshooting Example 1 – Cont.

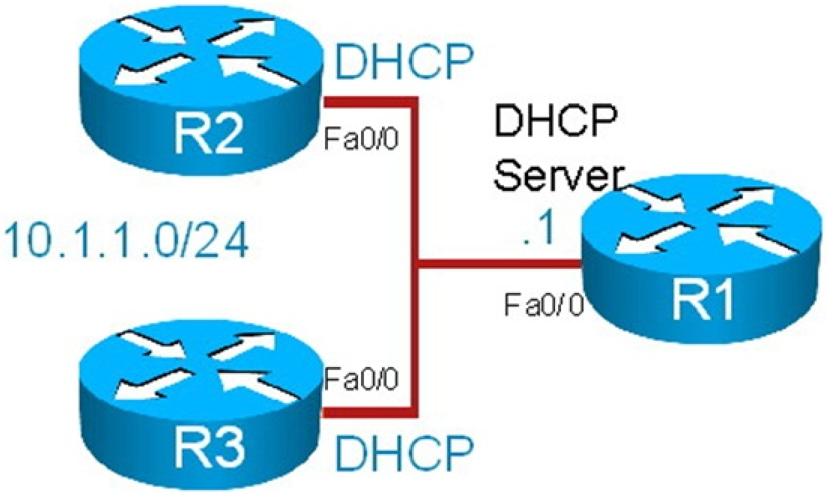

```
R1# show ip sockets
Proto  Remote    Port   Local     Port  In Out   Stat  TTY  OutputIF
88   --listen--     10.1.1.1      10  0    0      0    0
17   --listen--     10.1.1.1     161  0    0   1001    0
17   --listen--     10.1.1.1     162  0    0   1011    0
17   --listen--     10.1.1.1   57767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

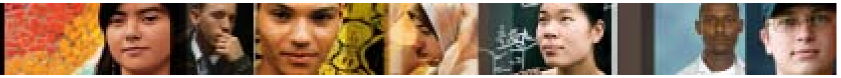

DHCP Troubleshooting Example 1 – Cont.

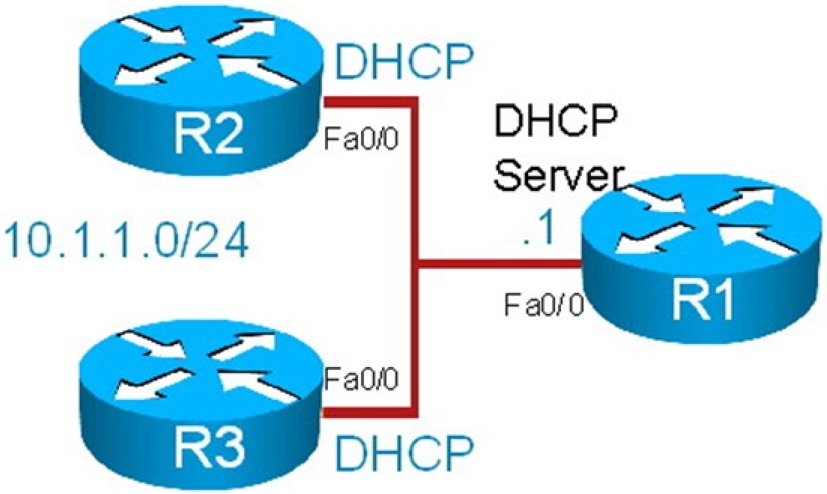

```
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.
R1(config)# service dhcp
R1(config)# end
R1#

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.

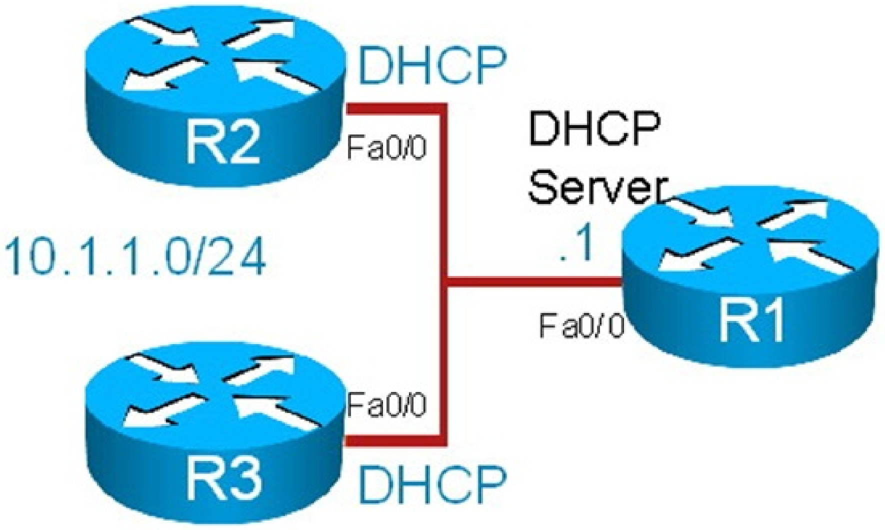

DHCP Troubleshooting Example 2 – Cont.

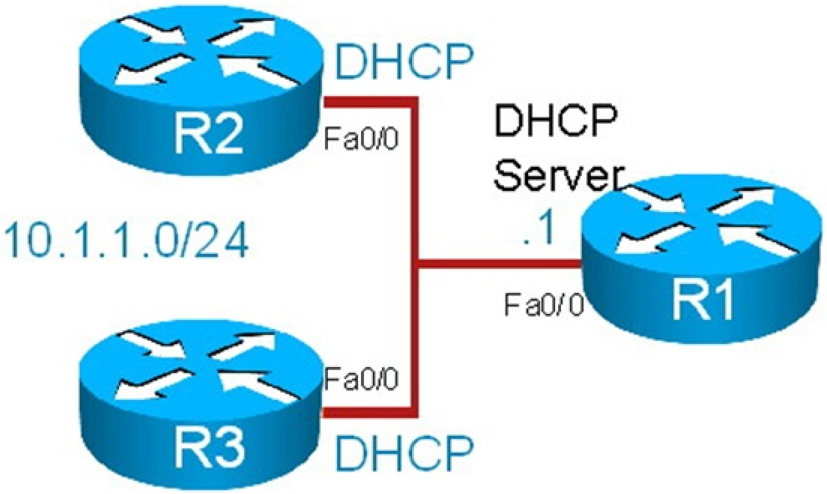

```
R1# show running-config | beg ip dhcp pool

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.

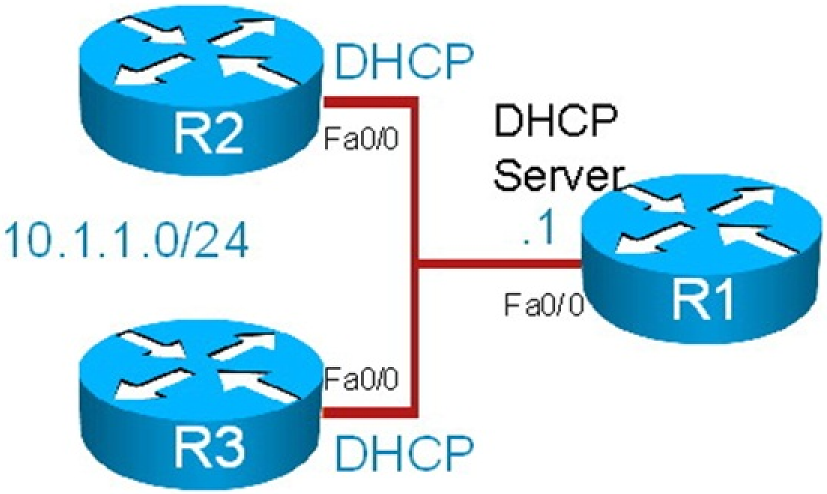

```
R1# show ip dhcp conflict
IP address          Detection method     Detection time              VRF
10.1.1.1            Gratuitous ARP       Aug 23 2009 06:28 PM
10.1.1.3            Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.3            Gratuitous ARP       Aug 23 2009 06:29 PM
10.1.1.4            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.

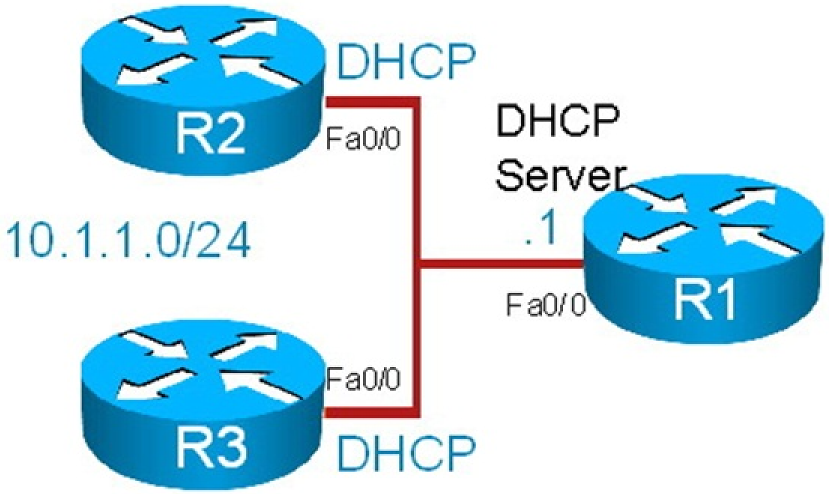

```
R1# sh run | inc excluded

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.

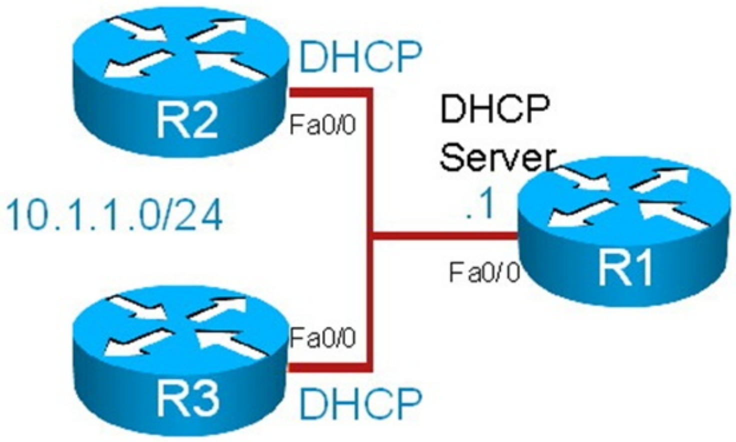

```
R1# conf t
Enter configuration commands, one per line. End with CNTL/Z.

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).

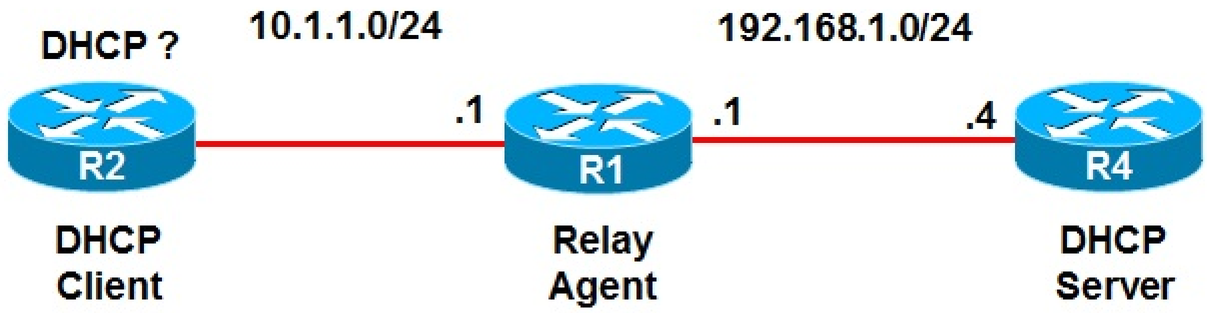

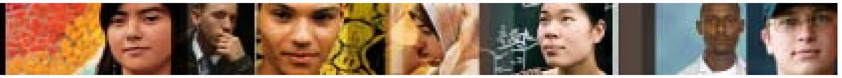

DHCP Troubleshooting Example 3 – Cont.

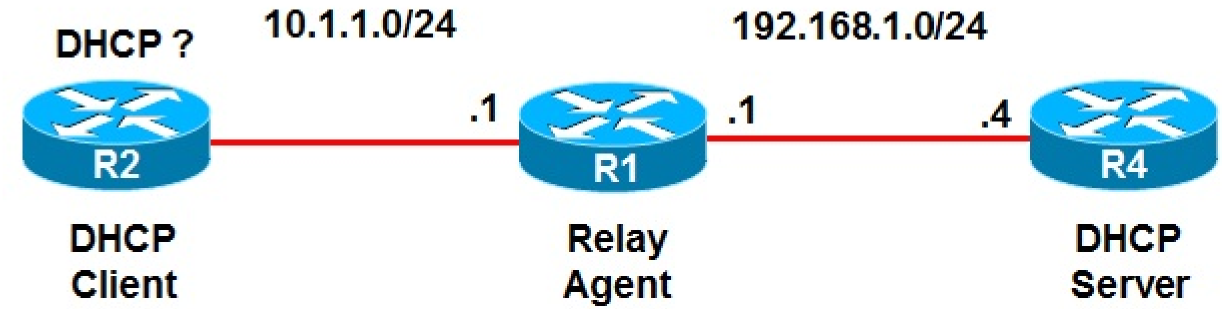

```
R1# debug ip udp
UDP packet debugging is on
R1#
R1#
*Aug 23 19:01:05.303: UDP: rcvd src-0.0.0.0(68), dst=255.255.255.255(67),
length=584
*Aug 23 19:01:05.303: UDP: broadcast packet dropped, src=0.0.0.0,
dst=192.168.1.255
*Aug 23 19:01:08.911: UDP: rcvd src-0.0.0.0(68), dst=255.255.255.255(67),
length=584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

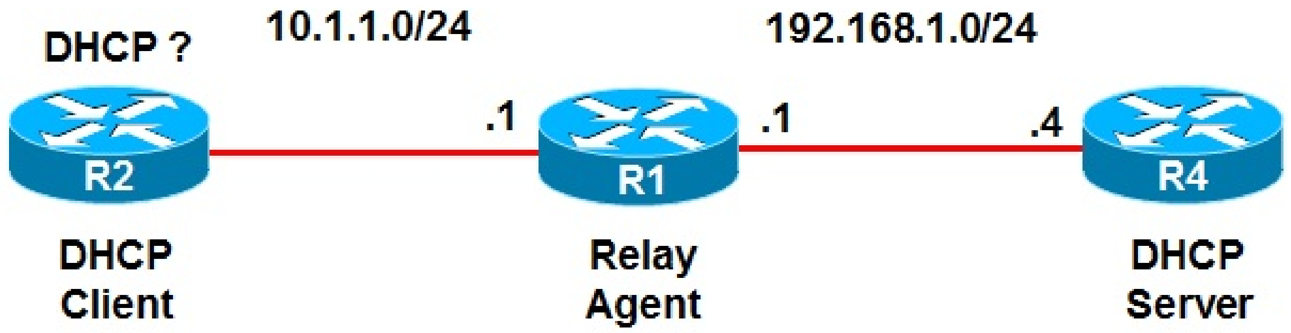

```
R1(config)# int fa0/0
R1(config-if)# ip helper-address 192.168.1.4
R1(config-if)# end
```

DHCP Troubleshooting Example 3 – Cont.

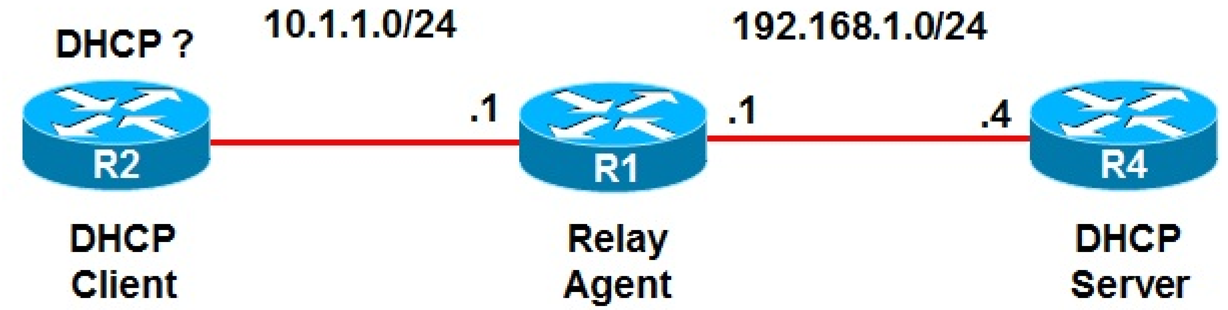

```
R4# debug ip udp
UDP packet debugging is on
R4#
*Aug 23 19:31:39.303: UDP: sent src=0.0.0.0(67), dst=255.255.255.255(68),length=308
*Aug 23 19:31:39.303: UDP: rcvd src=0.0.0.0(68), dst=255.255.255.255(67),length=584
*Aug 23 19:31:39.303: UDP: sent src=0.0.0.0(67), dst=255.255.255.255(68),length=308
*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
```

Cisco | Networking Academy®
Mind Wide Open™