Chapter 6 – Troubleshooting Addressing Services

Objectives

- Describe NAT & PAT operation & troubleshooting techniques.
- Describe DHCP operation & troubleshooting techniques.
- Describe the different methods of IPv6 address assignment.
- Explain the operation of OSPFv3 and RIPng.
- Describe typical IPv6 troubleshooting techniques.

NAT Benefits

- *Conserves* the legally registered addressing scheme

- Increases the *flexibility* of connections to the public network

- Provides *consistency* for internal network addressing schemes.

- Provides network *security.*

Configuring Static NAT

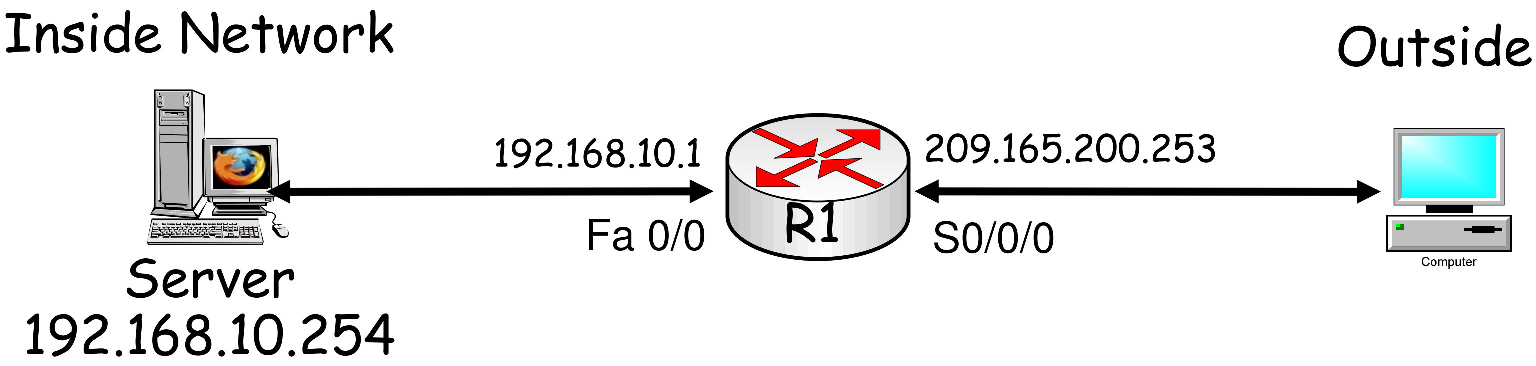

Inside Local Inside Global

```
R1(config)# ip nat inside source static 192.168.10.254 209.165.200.254
R1(config)#int s0/0/0
R1(config-if)#ip nat outside
R1(config-if)#int fa0/0
R1(config-if)#ip nat inside
```

- *Static NAT:* In this case, inside local (locally significant) and inside global (globally significant) addresses are mapped one to one.

- This mapping is particularly useful when an inside device must be accessible from the outside network, such as the case of web servers in an Internet data center.

Configuring Dynamic NAT

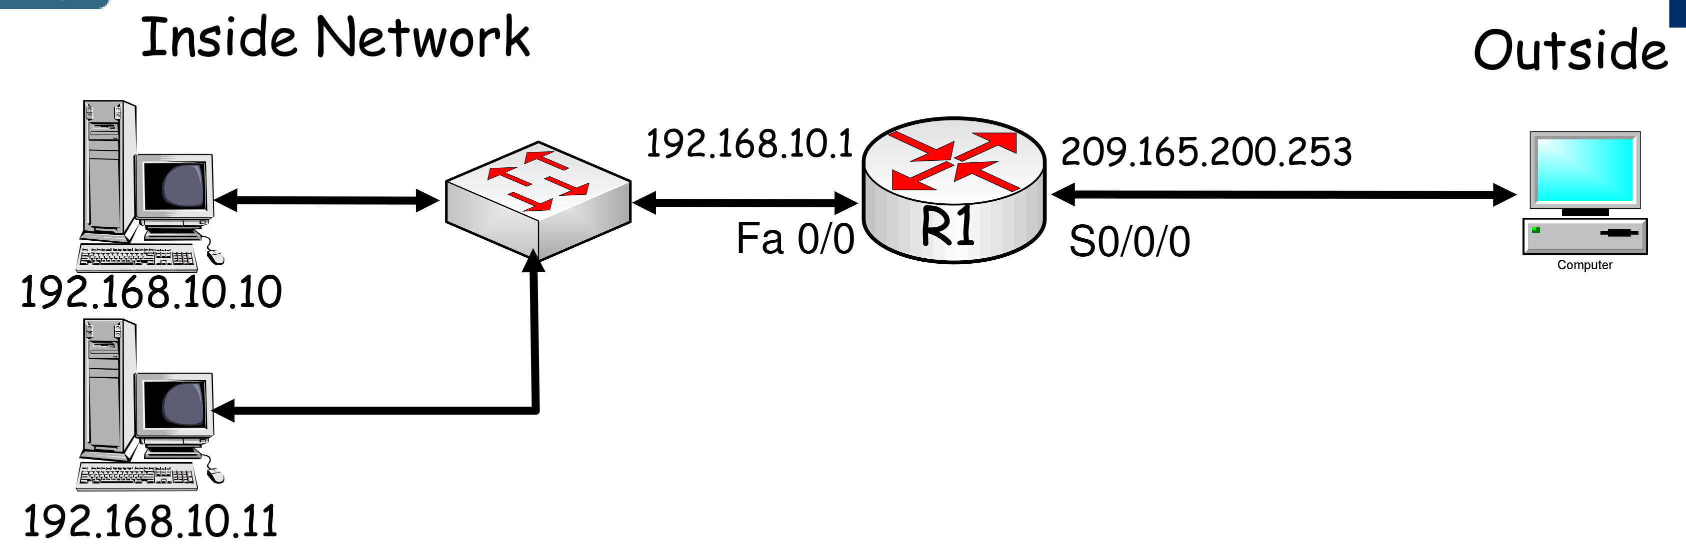

R1(config)# ip nat pool *POOL1* 209.165.200.226 209.165.200.240 netmask 255.255.255.224
R1(config) #access-list 1 permit 192.168.10.0 0.0.0.255
R1(config) #ip inside source list 1 pool POOL1
R1(config)#int s0/0/0
R1(config-if)#ip nat outside
R1(config-if)#int fa0/0
R1(config-if)#ip nat inside

- *Dynamic NAT:* translates addresses following the same underlying technology as static NAT; however, local addresses are translated to a *group* or *pool* of global addresses.
- Creates issues related to the size of that global pool, as there is a one-to-one translation once a global address has been selected.

Configuring NAT Overload (Single Address)

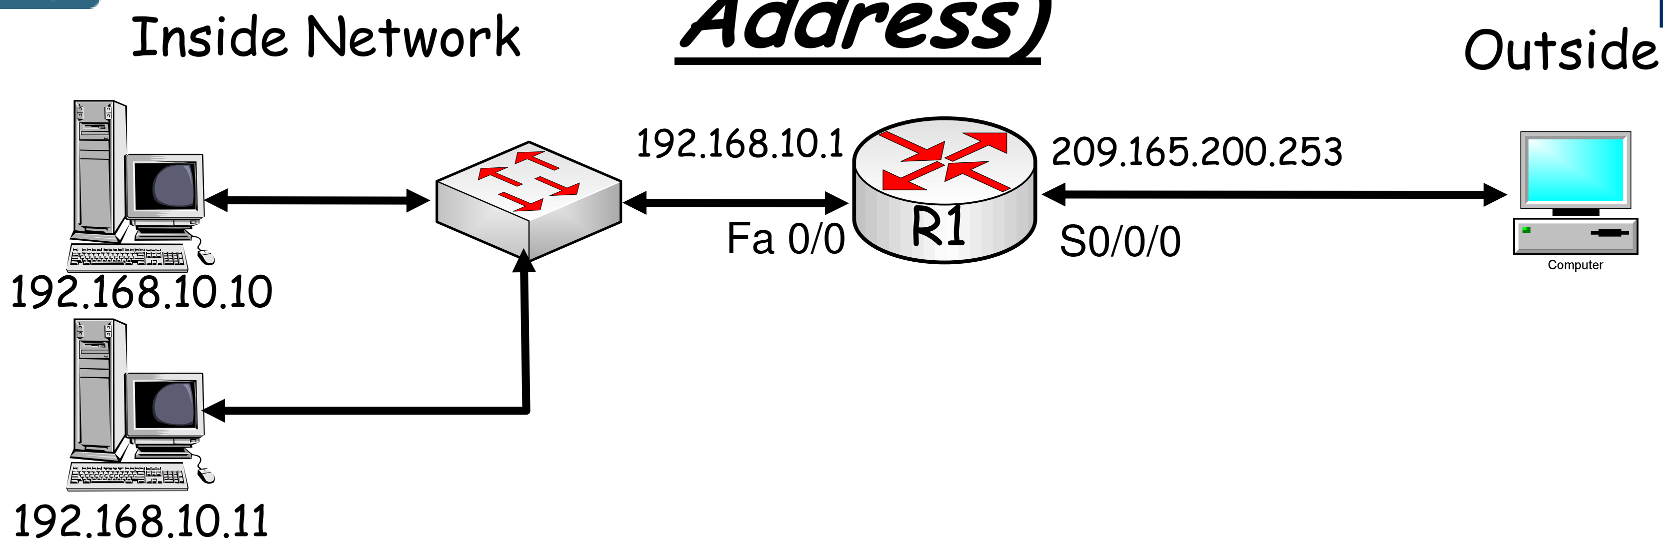

R1config) #access-list *1* permit 192.168.10.0 0.0.0.255
R1(config) #ip inside source list *1* interface serial 0/0/0 overload
R1(config)#int s0/0/0
R1(config-if)#ip nat outside
R1(config-if)#int fa0/0
R1(config-if)#ip nat inside

- With only one public IP address, the overload configuration typically assigns that public address to the outside interface that connects to the ISP.
- All inside addresses are translated to the single IP address when leaving the outside interface.

Configuring NAT Overload (Multiple Addresses)

Inside Network

Outside

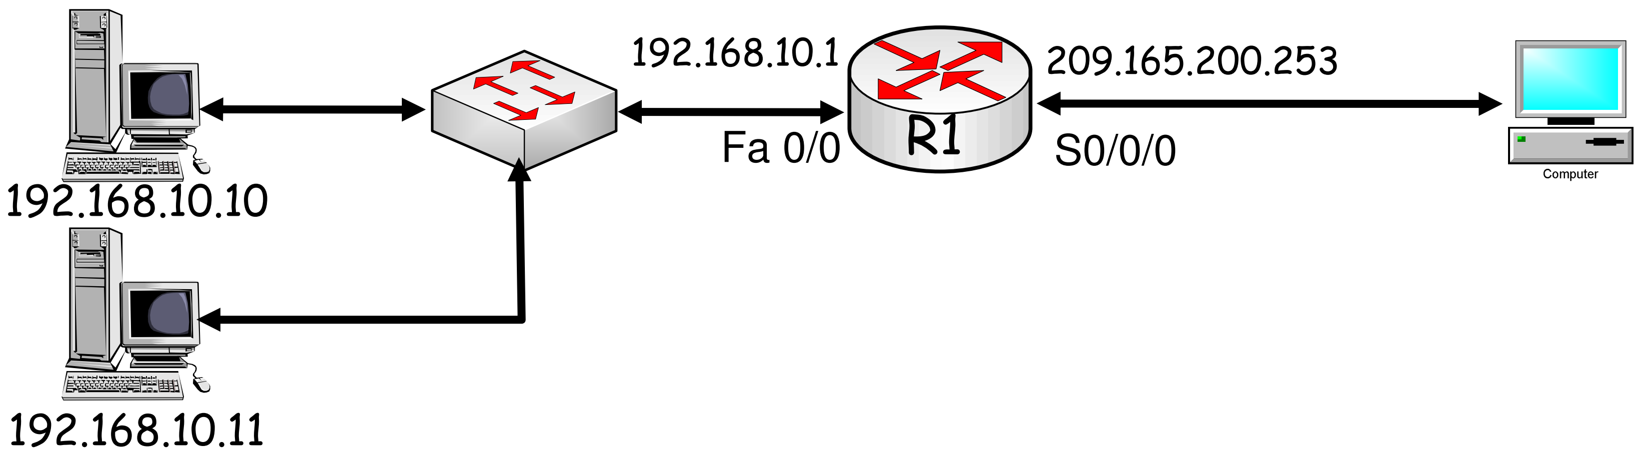

R1(config)# ip nat pool *POOL1* 209.165.200.226 209.165.200.240 netmask 255.255.255.224
R1config) #access-list *1* permit 192.168.10.0 0.0.0.255
R1(config) #ip inside source list *1* pool *POOL1* overload
R1(config)#int s0/0/0
R1(config-if)#ip nat outside
R1(config-if)#int fa0/0
R1(config-if)#ip nat inside

- In the scenario where the ISP has provided more than one public IP address, NAT overload is configured to use a pool.
- The primary difference between this configuration and the configuration for dynamic, one-to-one NAT is that the overload keyword is used.

NAT Drawbacks

- Some applications or protocols have direct conflict with NAT or PAT.
- VPNs encapsulates the original IP address, and doesn't provide access to UDP or TCP port numbers – *NAT Transparency* or *NAT Traversal* required.
- Multimedia applications negotiate ports at the moment of connection, or have IP addresses embedded in the payload of the packets, requiring NAT to be application-aware.
- Applications and protocols as such might be labeled as NAT-sensitive: Kerberos, X-windows, rsh, SIP, SNMP, FTP and DNS.

Router Interface Order of Operations

Inside Network

Outside

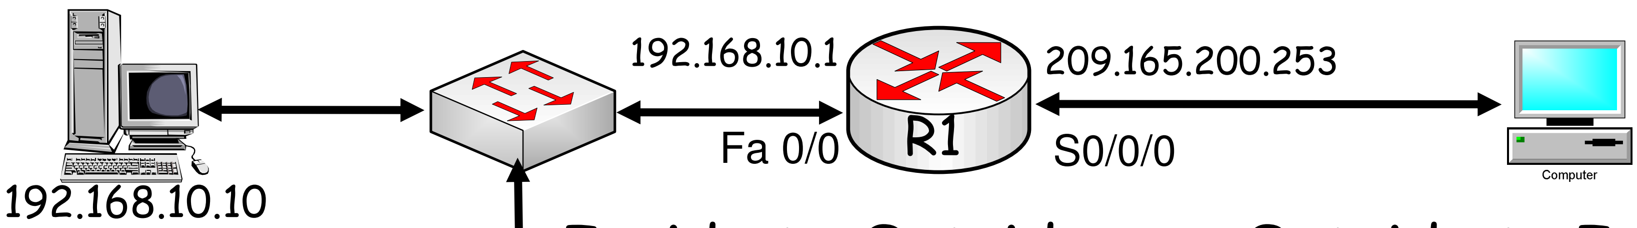

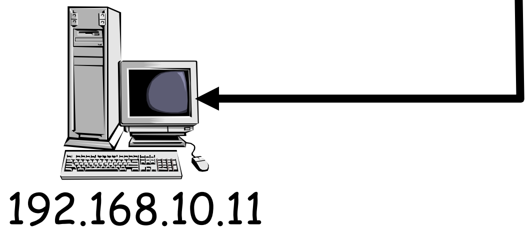

Inside to Outside

- IPsec decryption
- Input ACL
- Input rate limit
- Input accounting
- Policy routing
- IP routing
- Redirect to web cache
- NAT (inside to outside)
- Crypto map check
- Output ACL
- Firewall inspect
- TCP intercept
- Encryption

Outside to Inside

- IPsec decryption
- Input ACL
- Input rate limit
- Input accounting
- NAT (outside to inside)
- Policy routing
- IP routing
- Redirect to web cache
- Crypto map check
- Output ACL
- Firewall inspect
- TCP intercept
- Encryption

Common NAT Problems

- An ACL referenced by a NAT configuration is incorrect.
- Inside and outside interfaces are not correctly assigned.
- Incorrect IP addresses (or address ranges) are referenced by a NAT configuration.
- Applications are not NAT aware.
- A routing loop occurs as a result of a NAT address translation.

Verify NAT

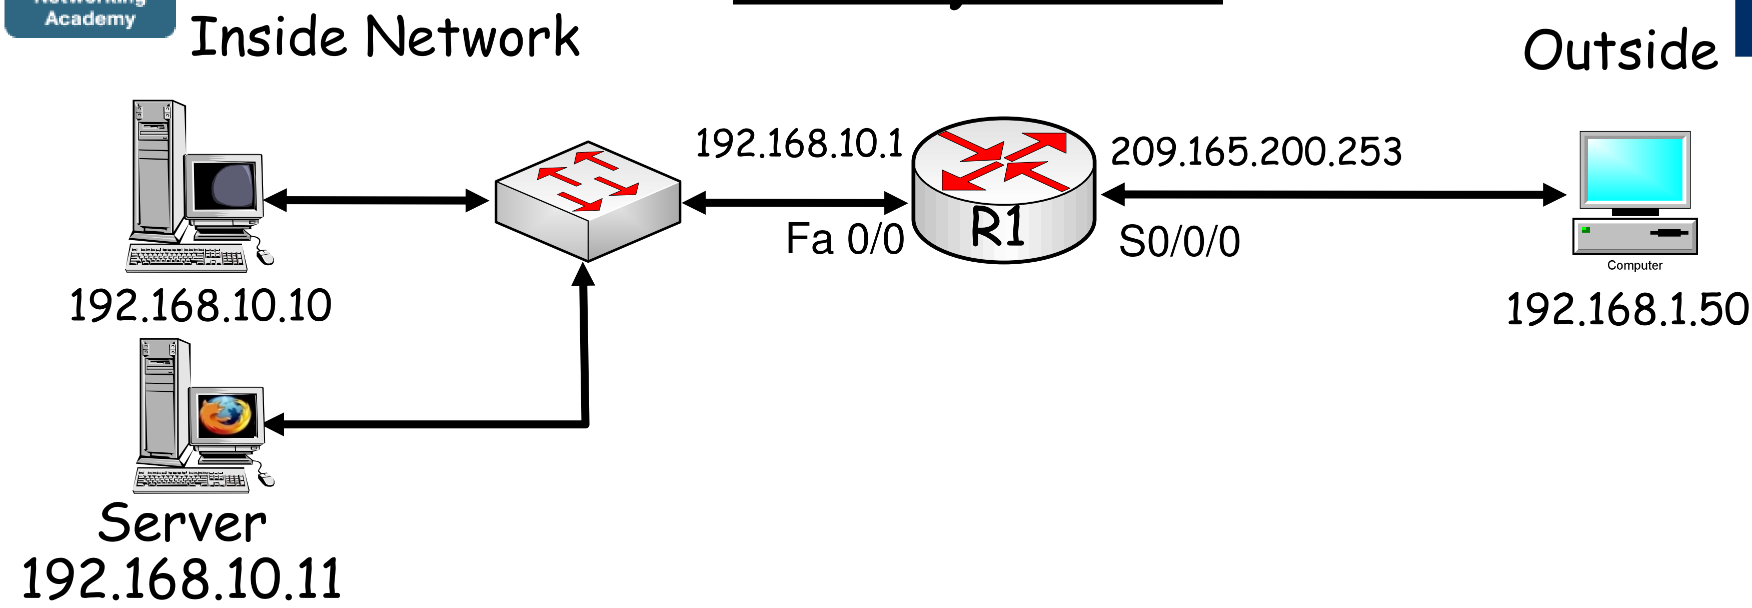

R1# **show ip nat translations**

Pro	Inside global	Inside local	Outside local	Outside global
---	209.165.200.254	192.168.10.11	---	---
tcp	209.165.200.254:23	192.168.10.10:23	192.168.1.50:1158	192.168.1.50:1158

R1 #clear ip nat translation*

R1# **show ip nat translations**

Pro	Inside global	Inside local	Outside local	Outside global
---	209.165.200.254	192.168.10.11	---	---

Verify NAT

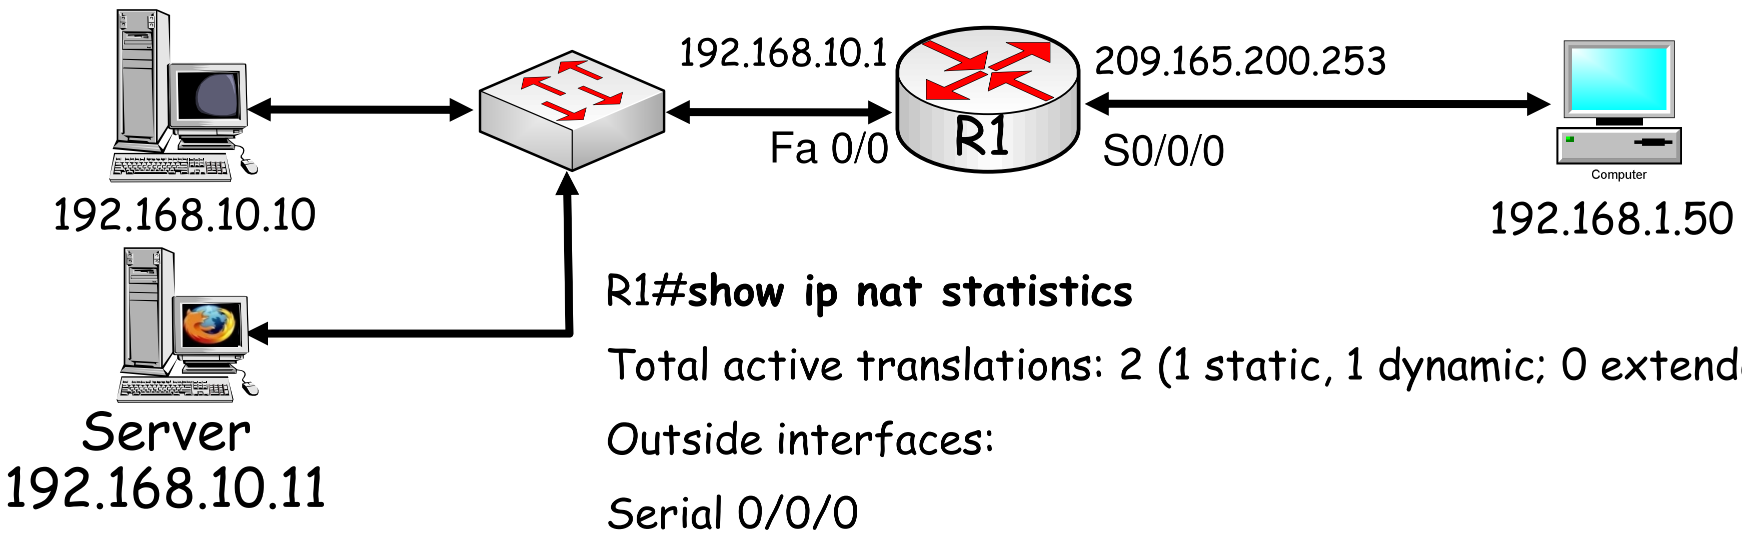

R1#**show ip nat statistics**

Total active translations: 2 (1 static, 1 dynamic; 0 extended)

Outside interfaces:

Serial 0/0/0

Inside interfaces:

FastEthernet0/0

Hits: 10 Misses: 0

CEF Translated packets: 5, CEF Punted packets: 0

Expired translations: 0

Dynamic mappings:

Appl doors: 0

Normal doors: 0

Queued Packets: 0

Verify NAT

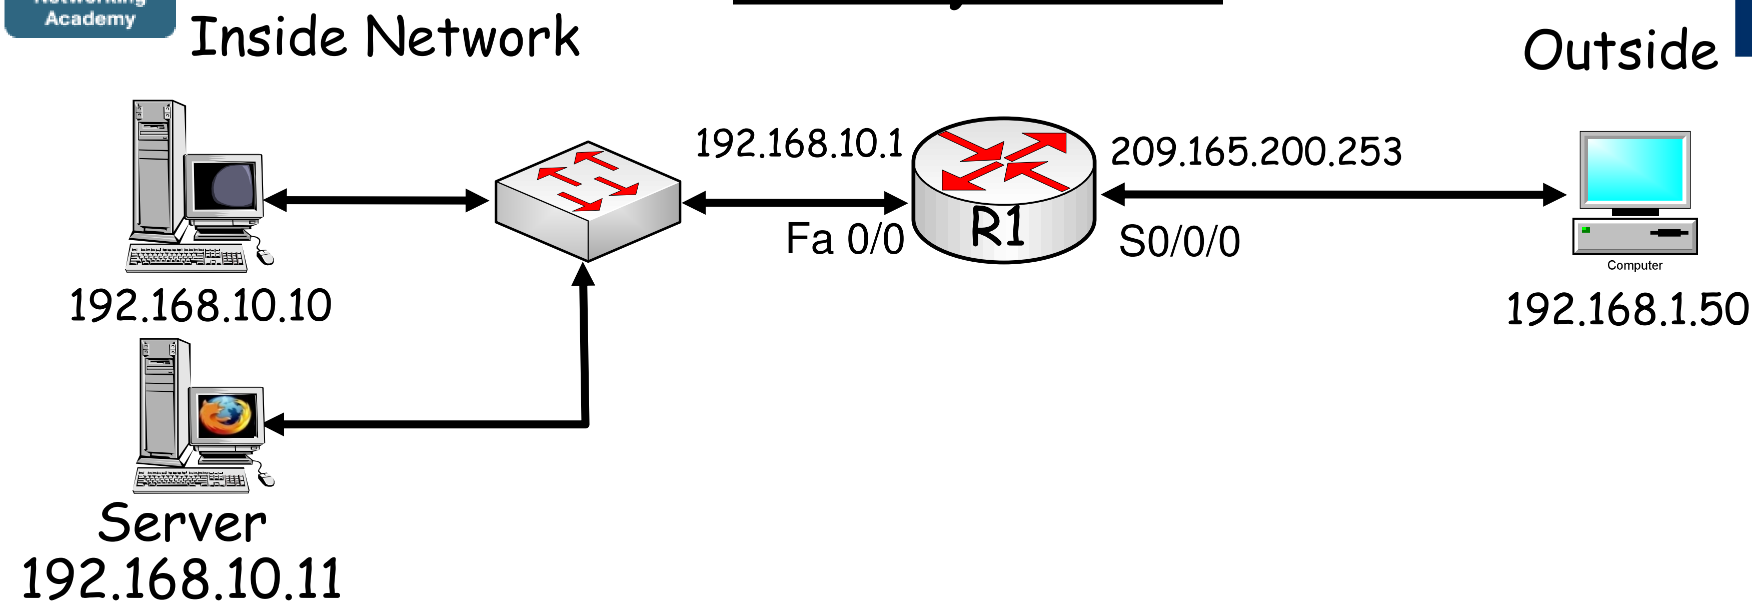

R1#**debug ip nat**

IP NAT debugging is on

Mar 3 13:01:28.162: NAT: s=192.168.1.50, d= 209.165.200.254 ->192.168.10.11 [10202]

*Mar 3 13:01:28.162: NAT: s=192.168.10.10->209.165.200.253, d=192.168.1.50 [210]

Mar 3 13:01:30.991: NAT: s=192.168.1.50, d= 209.165.200.254 ->192.168.10.11 [10370]

- The asterisk next to NAT indicates that the translation is occurring in the fast-switched path

Dynamic Host Configuration Protocol (DHCP)

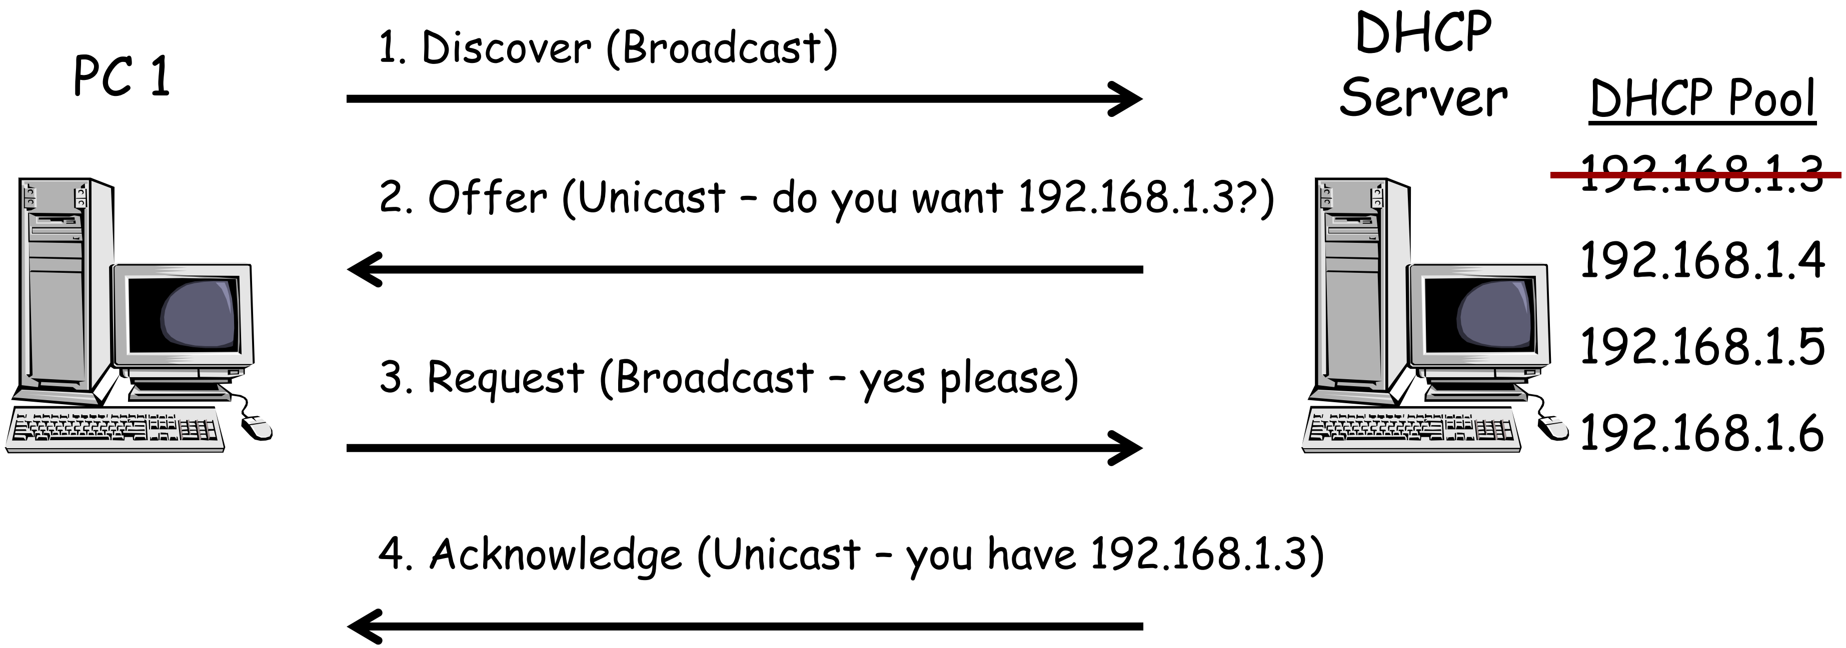

•*DHCP DECLINE:* Client-to-server communication, indicating that the IP address is already in use.
•*DHCP ACK :* Server-to-client communication. This is the server's response to a client REQUEST. This message includes all configuration parameters.
•*DHCP NACK:* Server-to-client communication. This is 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 Configuration using Importing

R1(config)#ip dhcp-excluded address 10.0.0.1 10.0.0.5
R1(config)#ip dhcp pool CENTRAL
R1(dhcp-config)#network 10.0.0.0 255.255.255.0
R1(dhcp-config)#default-router 10.0.0.1
R1(dhcp-config)#domain name central.com
R1(dhcp-config)#dns-server 10.0.0.2
R1(dhcp-config)#netbios-name-server 10.0.0.2
R1(config)#interface fastethernet0/0
R1(config-if)#ip address 10.0.0.1 255.255.255.0

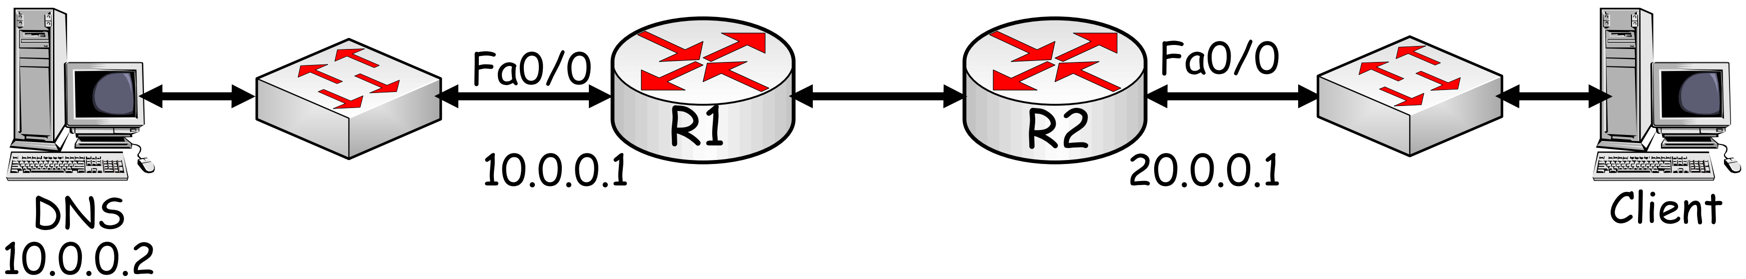

R2(config)#ip dhcp-excluded address 20.0.0.2
R2(dhcp-config)# ip dhcp pool CLIENT
R2(dhcp-config)# network 20.0.0.0 255.255.255.0
R2(dhcp-config)# default-router 20.0.0.1
R2(dhcp-config)# import all

DHCP Relay

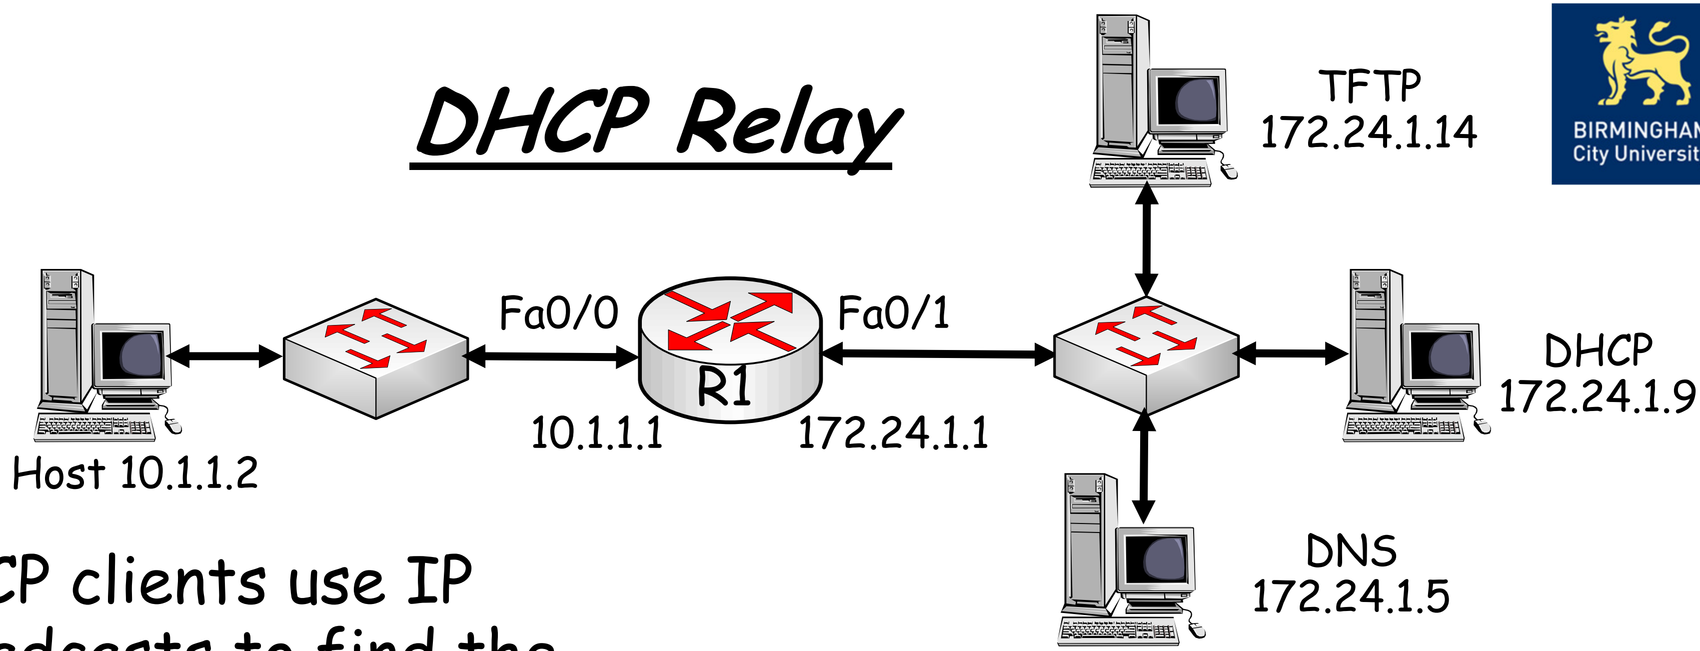

- DHCP clients use IP broadcasts to find the DHCP server on the segment - Routers do not forward these broadcasts.
- When possible, administrators should use the *ip helper-address* command to relay broadcast requests for key UDP services.

•Other protocols that are forwarded by a DHCP relay agent include the following:

•TFTP
•Domain Name System (DNS)
•Internet Time Service (ITS)
•NetBIOS name server
•NetBIOS datagram server
•BootP
•TACACS

Configuring IP helper addresses

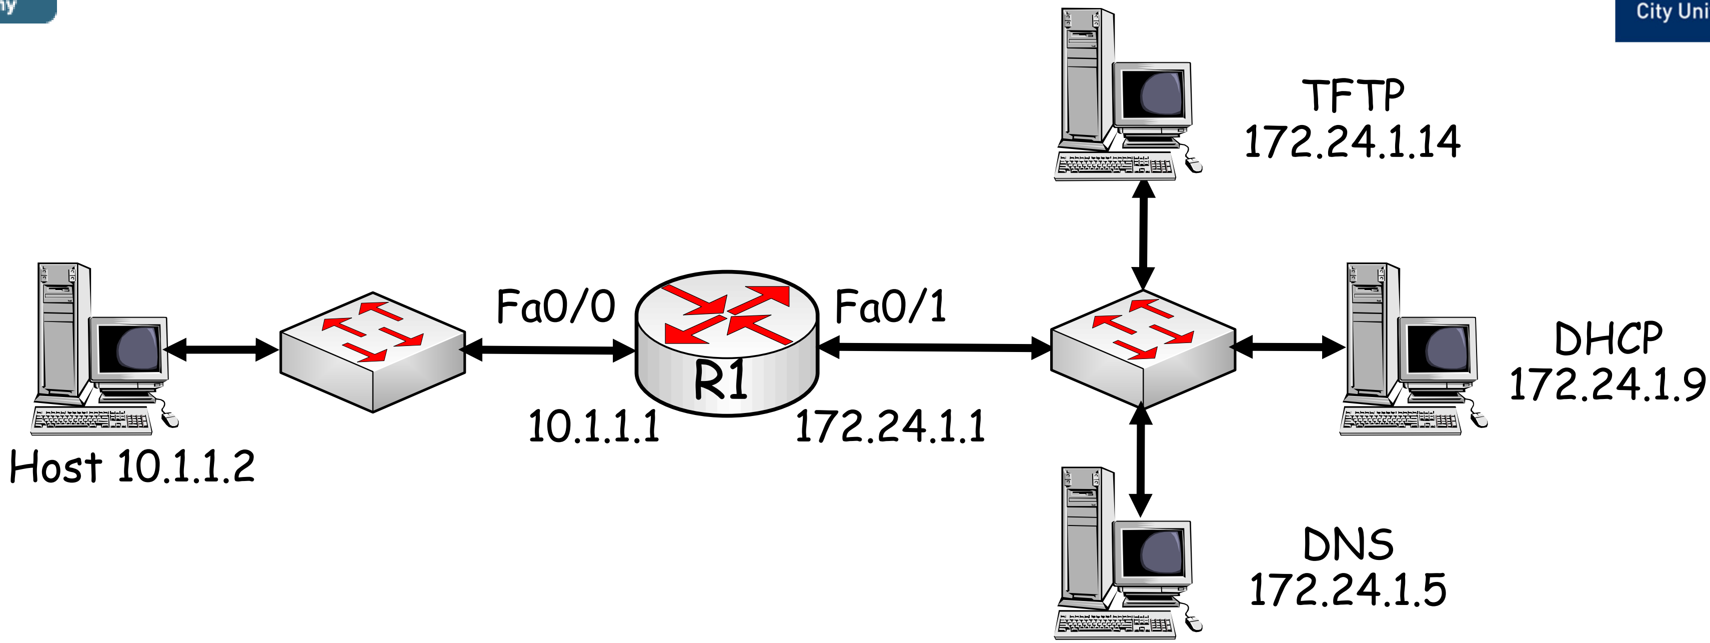

To configure R1 Fa0/0 (the interface that receives the Host broadcasts) to relay DHCP broadcasts as a unicast to the DHCP server, use the following commands:

```
R1(config)#interface Fa0/0
R1(config-if)#ip helper-address 172.24.1.9
```

Common DHCP Problems

- A router not forwarding broadcasts
- DHCP pool out of IP addresses
- Misconfiguration
- Duplicate IP addresses
- Redundant services not communicating
- The “pull” nature of DHCP

Verify DHCP

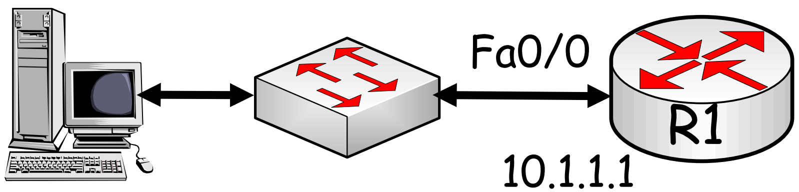

R1#sh ip dhcp binding

IP address	Client-ID/ Hardware address	Lease expiration	Type
10.1.1.2	0000.0C9B.9C83	Feb 11 2010 06:14 AM	Automatic

R1#clear ip dhcp binding*

R1#show ip dhcp conflict

IP address	Detection method	Detection time
10.1.1.2	Ping	Oct 15 2009 8:56 PM

R1# clear ip dhcp conflict *

Verifying DHCP

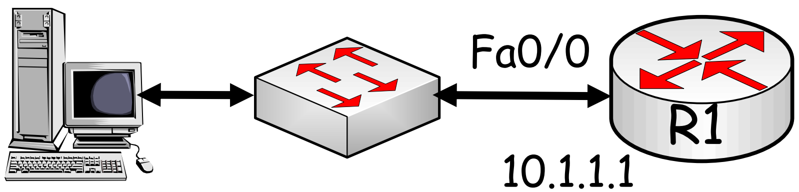

R1#sh ip dhcp server statistics

Memory usage	25307
Address pools	1
Database agents	0
Automatic bindings	1
Manual bindings	0
Expired bindings	0
Malformed messages	0
Secure arp entries	0

Message	Received
BOOTREQUEST	0
DHCPDISCOVER	8
DHCPREQUEST	1
DHCPDECLINE	0
DHCPRELEASE	0
DHCPINFORM	0

- To verify that messages are being received or sent by the router, use the ***show ip dhcp server statistics*** command.
- This command displays count information regarding the number of DHCP messages that have been sent and received.

Verifying DHCP

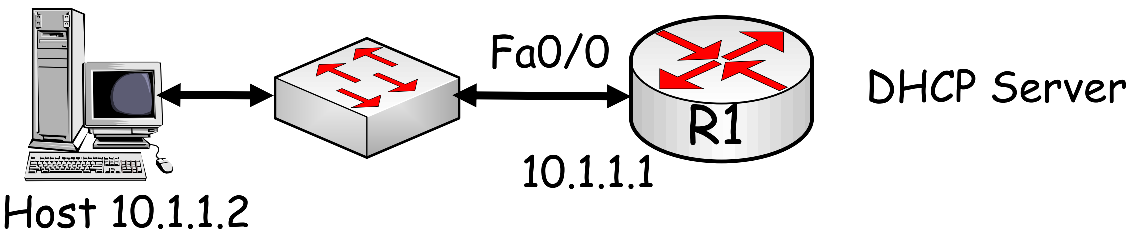

R1#sh ip dhcp pool

Pool POOL1:
Utilization mark (high/low) : 100 / 0
Subnet size (first/next) : 0 / 0
Total addresses : 254
Leased addresses : 1
Pending event : 1
1 subnet is currently in the pool:

Current index	IP address range	Leases addresses
10.1.1.3	10.1.1.2 - 10.1.1.254	1

Verifying DHCP

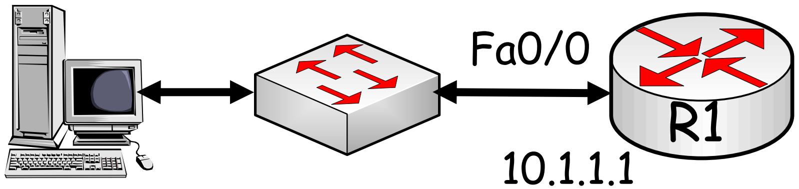

R1#**debug ip dhcp server packet**

*Mar 1 00:07:39.867: DHCPD: DHCPRELEASE message received from client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 (10.1.1.2).
*Mar 1 00:07:41.855: DHCPD: DHCPRELEASE message received from client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 (10.1.1.2).
*Mar 1 00:07:41.859: DHCPD: Finding a relay for client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 on interface FastEthernet0/1.
*Mar 1 00:07:54.775: DHCPD: DHCPDISCOVER received from client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 on interface FastEthernet0/0.
*Mar 1 00:07:54.779: DHCPD: Allocate an address without class information (10.1.1.0)
*Mar 1 00:07:56.783: DHCPD: Sending DHCPOFFER to client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 (10.1.1.2).
*Mar 1 00:07:56.787: DHCPD: broadcasting BOOTREPLY to client c001.0f1c.0000.
Mar 1 00:07:56.879: DHCPD: DHCPREQUEST received from client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30.
*Mar 1 00:07:56.887: DHCPD: No default domain to append - abort update
*Mar 1 00:07:56.887: DHCPD: Sending DHCPACK to client 0063.6973.636f.2d63.3030.312e.3066.3163.2e30.3030.302d.4661.302f.30 (10.1.1.2).
*Mar 1 00:07:56.891: DHCPD: broadcasting BOOTREPLY to client c001.0f1c.0000.

Troubleshooting IPv6 Issues

•Examining IPv6 issues reveals that there are many common configuration mistakes:

•Mis-configured auto-configuration on routers.

•IPv6 routing problems, such as suboptimal routing due to improper summarization, and parameter mismatches on protocols such as OSPF that negotiate parameters.

•For tunnel scenarios, due to the great variety of methods, there are often instances in which other components such as routing protocols need to change when the specific migration or tunneling method changes.

IPv6 Address Types

- Unicast
  - Address is for a single interface.
  - IPv6 has several types (for example, global and IPv4 mapped).
- Multicast
  - Broadcasts are replaced by multicast addresses. Multicast enables efficient network operation by using functionally specific multicast groups to send requests to a limited number of computers on the network.
  - A packet sent to a multicast address is delivered to all interfaces identified by that address.
- Anycast
  - IPv6 also defines a new type of address called anycast. An anycast address identifies a list of devices or nodes; therefore, an anycast address identifies multiple interfaces.
  - Routers decide on closest device to reach that destination.
  - Suitable for load balancing and content delivery services.

Multicast Address Structure

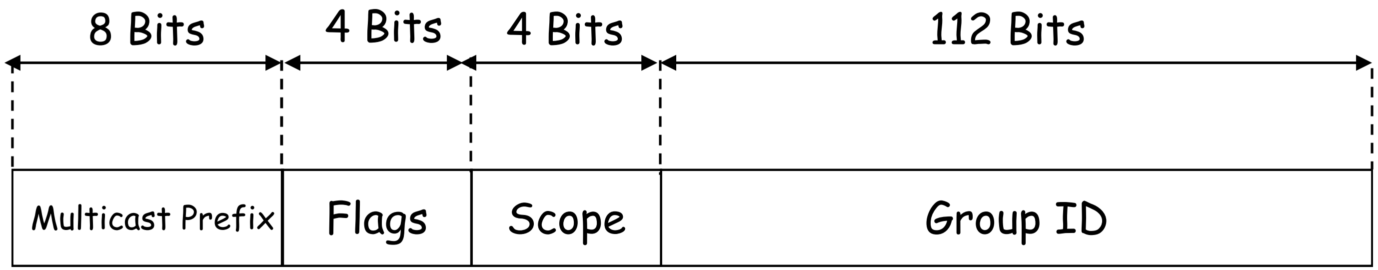

Address	Multicast Group
FF02::1	All Nodes
FF02::2	All Routers
FF02::5	OSPFv3 Routers
FF02::6	OSPFv3 DRs
FF02::9	RIPng Routers
FF02::A	EIGRP Routers
FF02::D	All PIM Routers

- A multicast address identifies not one device but a set of devices – a multicast *group*.
- A packet being sent to a multicast group is originated by a single device – a multicast packet has a *unicast* address as its source and a multicast address as its destination.
- Multicast is essential to the basic operation of IPv6, particularly some of its *plug-and-play* features such as neighbour discovery and autoconfiguration.

Assigning IPv6 Addresses

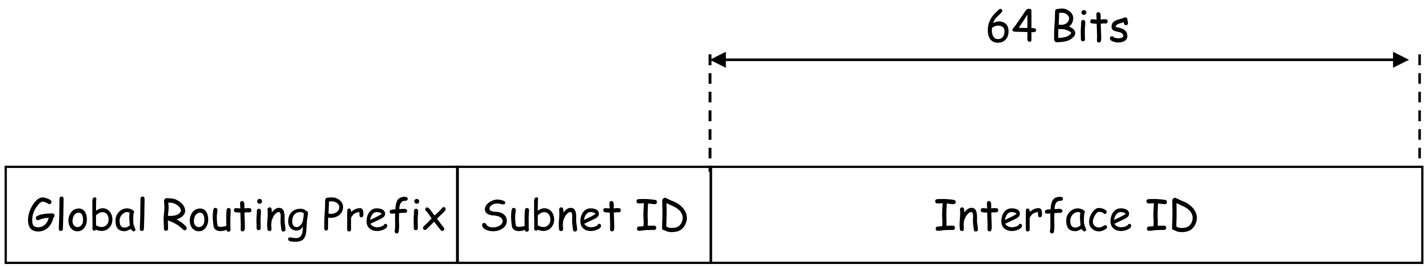

- IPv6 addresses use interface identifiers to identify *interfaces* on a link.
- Interface identifiers are required to be *unique* on a specific link.
- Interface identifiers are always 64 bits and can be dynamically derived from a *Layer 2* address (MAC).

- IPv6 address ID can be assigned *statically* or *dynamically*:

1. Static assignment using a manual interface ID
2. Static assignment using an Extended Universal Identifier 64 (EUI-64) interface ID
3. Stateless auto-configuration
4. DHCP for IPv6 (DHCPv6)

Stateless Address Autoconfiguration

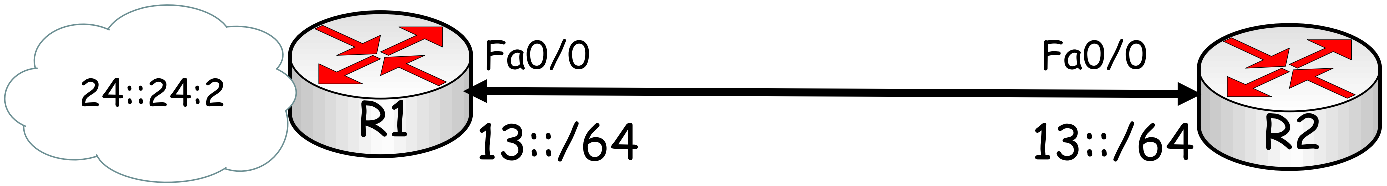

R2(config-if)# ipv6 address autoconfig

R2#**show ipv6 interface 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: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 Address Types

Address Type	MSB (Binary)	MSB (Hex)
Unspecified	00..0	::/128
Loopback	00..1	::1/128
Multicast	11111111	FF00::/8
Link-Local Unicast	1111111010	FE80::/10
Global Unicast	001	2xxx::/4 Or 3xxx::/4

Link-Local Address

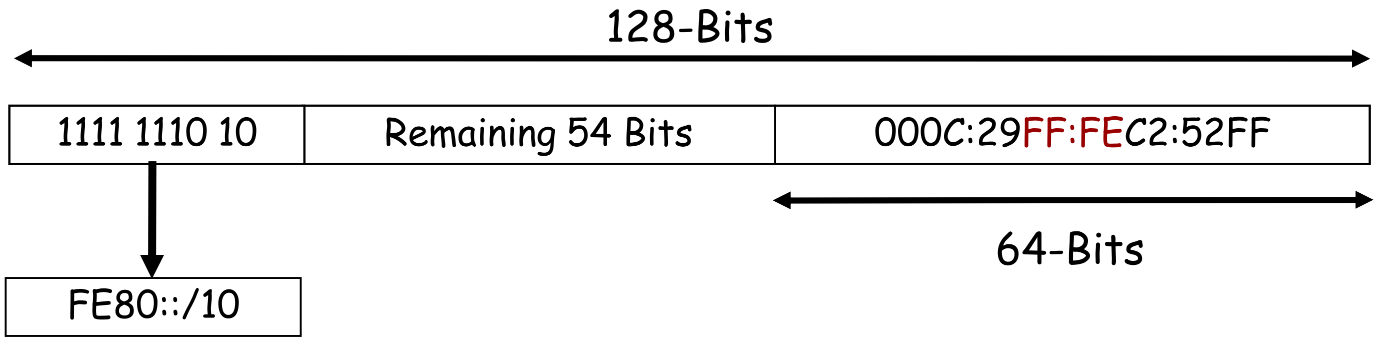

- Mandatory address for communication between two IPv6 devices (similar to ARP but at Layer 3).
- Automatically assigned by router as soon as IPv6 is enabled using stateless auto-configuration.
- Also used for neighbour relationships and next-hop calculation in routing protocols.

MAC Address: 00-0C-29-C2-52-FF
EUI-64 conversion: 000C:29FF:FEC2:52FF (locally administered)
EUI-64 conversion: 020C:29FF:FEC2:52FF (universal)

Phases of Stateless Auto-configuration

- **Phase 1:** The most common method to obtain a unique identifier on an Ethernet link is by using the EUI-48 MAC address and applying the modified IEEE *EUI-64* standard algorithm.
- **Phase 2:** The well-known link-local prefix *fe80::/64* is prepended to the 64-bit identifier from phase 1 to create the 128-bit link-local address. This address is associated with the interface and tagged *tentative*.
- **Phase3:** Before final association, it is necessary to verify the address's uniqueness on the link, called *duplicate address detection* (DAD).
- **Phase4:** This phase removes the tentative tag and formally assigns the address to the network interface. The system can now communicate with its neighbors on the link.

Neighbour Discovery Protocol (NDP)

- The most distinct characteristic of IPv6 after it's increased address space are it's plug-and-play features. NDP is the enable of these features, using the following functions:
  - Router Discovery
  - Prefix Discovery
  - Parameter Discovery
  - Address Auto-configuration
  - Address Resolution
  - Next-Hop Determination
  - Neighbour Unreachability Detection
  - Duplicate Address Detection
  - Redirect

NDP Messages

- *Router Advertisement (RA)* – Originated by routers to advertise their presence and link-specific parameters such as link prefixes, MTU and hop-limits. Sent periodically (every 200 seconds in Cisco routers), and in response to Router Solicitation messages.
- *Router Solicitation (RS)* –Originated by hosts to request that a router sends an RA.
- *Neighbour Solicitation (NS)* - Originated by nodes to request another node's link layer address and also for duplicate address detection (DAD) and neighbour reachability.
- *Neighbour Advertisement (NA)* – sent in response to NS messages.
- *Redirect* – Allows routers to advise clients of better exit gateways.

IPv6 Routing Protocols

- The following routing protocols have been developed to support IPv6:

1. *Routing Information Protocol next generation* (RIPng), is a distance vector routing protocol with a limit of 15 hops that uses split horizon and poison reverse to prevent routing loops.
2. *OSPFv3* is Based on OSPF version 2 (OSPFv2), with enhancements.
3. *IPv6 IS-IS* - with large address support facilitates the IPv6 address family
4. *EIGRP IPv6* runs over an IPv6 transport, communicates only with IPv6 peers, and advertises only IPv6 routes.
5. *Multiprotocol BGP* (MBGP)- RFC 2858 (which replaces the obsolete RFC 2283) defines multiprotocol extensions for BGP4

OSPFv3 Configuration

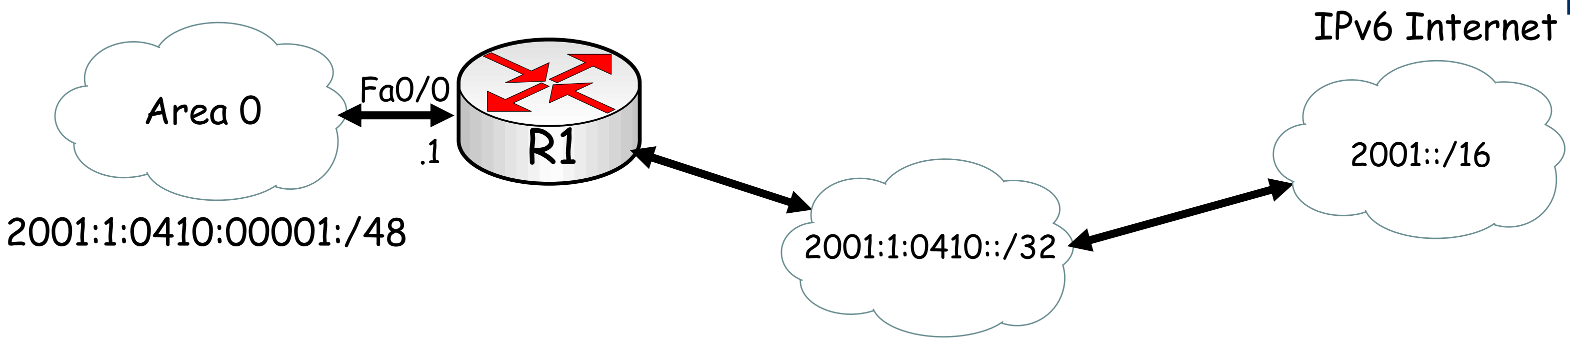

R1(config)#ipv6 unicast-routing
R1(config)#ipv6 cef
R1(config)#ipv6 router ospf 1
R1(config-router)# router-ID 1.1.1.1

R1(config)#interface fa0/0
R1(config-if)# ipv6 address 2001:4010:0001::1/48
R1(config-if)# ipv6 ospf 1 area 0
R1(config-if)#ipv6 ospf priority 20
R1(config-if)#ipv6 ospf cost 20

R1(config)ipv6 router ospf 1
R1(config-router)#area 0 range 2001:0410::/32

Verify OSPFv3

R1#**show ipv6 interface serial 0/0/0**

Serial0/0/0 is up, line protocol is up

IPv6 is enabled, link-local address is FE80::219:6FF:FE23:4380

No Virtual link-local address(es):

Global unicast address(es):

FEC0::12:1, subnet is FEC0::12:0/112

Joined group address(es):

FF02::1

FF02::2

FF02::1:FF12:1

FF02::1:FF23:4380

MTU is 1500 bytes

ICMP error messages limited to one every 100 milliseconds

ICMP redirects are enabled

ICMP unreachables are sent

ND DAD is enabled, number of DAD attempts: 1

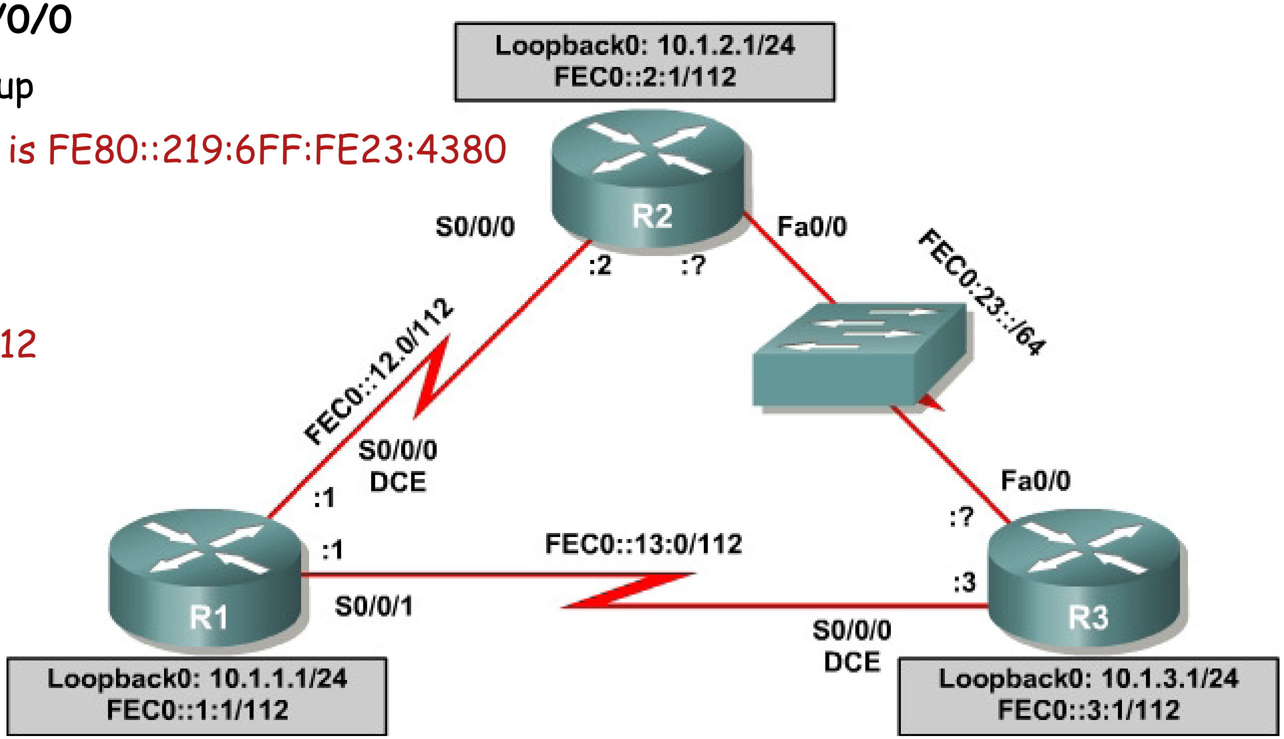

R1(config)# interface serial0/0/0
R1(config-if)# ipv6 address FE80::1 link-local

- When pinging link local addresses, specify an outgoing interface because the addresses are not routed and *not* in the routing table.

Verify OSPFv3

R1#**show ipv6 route**

IPv6 Routing Table - 11 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 D - EIGRP, EX - EIGRP external

```
L   FE80::/10 [0/0]                    <---- Link-Local bit bucket
     via ::, Null0
C   FEC0::1:0/112 [0/0]                <---- Subnet
     via ::, Loopback0
L   FEC0::1:1/128 [0/0]                <---- Interface
     via ::, Loopback0
O   FEC0::2:1/128 [110/64]
     via FE80::2, Serial0/0/0          <---- Link-Local next hop
O   FEC0::3:1/128 [110/64]
     via FE80::218:B9FF:FECD:BEF0, Serial0/0/1
C   FEC0::12:0/112 [0/0]
     via ::, Serial0/0/0
L   FEC0::12:1/128 [0/0]
     via ::, Serial0/0/0
C   FEC0::13:0/112 [0/0]
     via ::, Serial0/0/1
L   FEC0::13:1/128 [0/0]
     via ::, Serial0/0/1
O   FEC0:23::/64 [110/65]
     via FE80::2, Serial0/0/0
     via FE80::218:B9FF:FECD:BEF0, Serial0/0/1
L   FF00::/8 [0/0]                     <---- Multicast bit bucket
     via ::, Null0
```

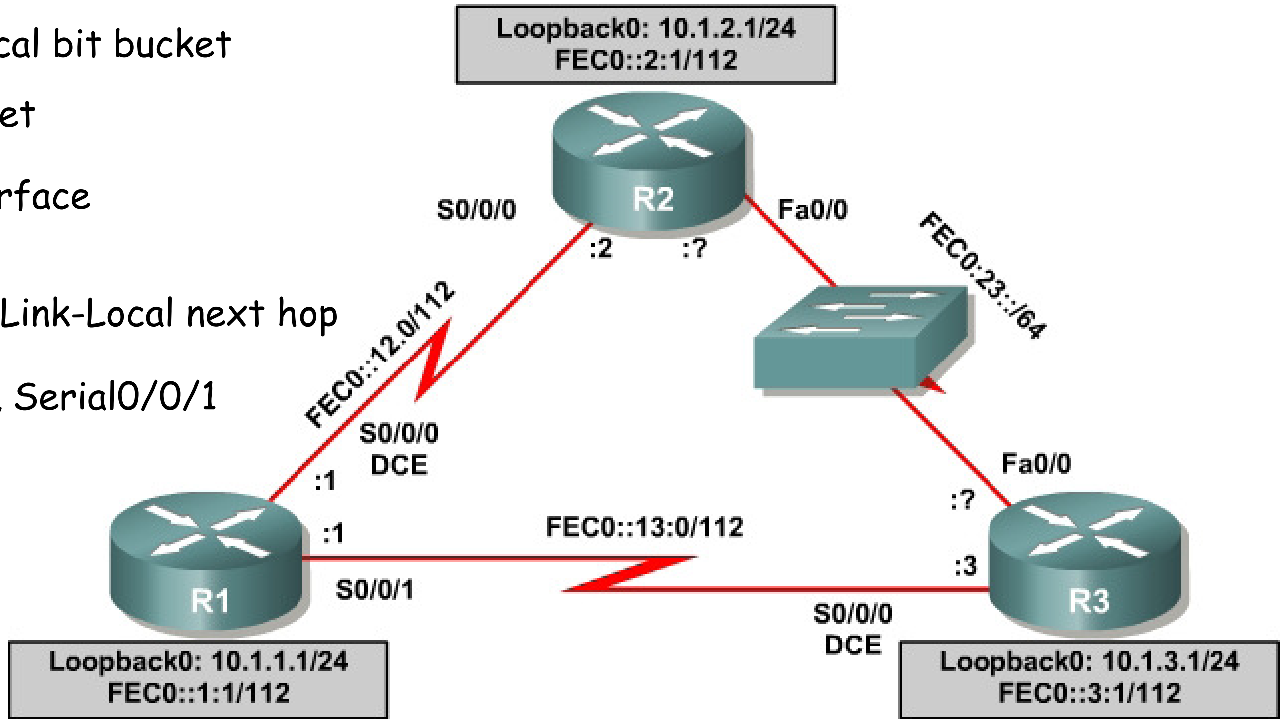

OSPFv3 Frame-Relay Configuration

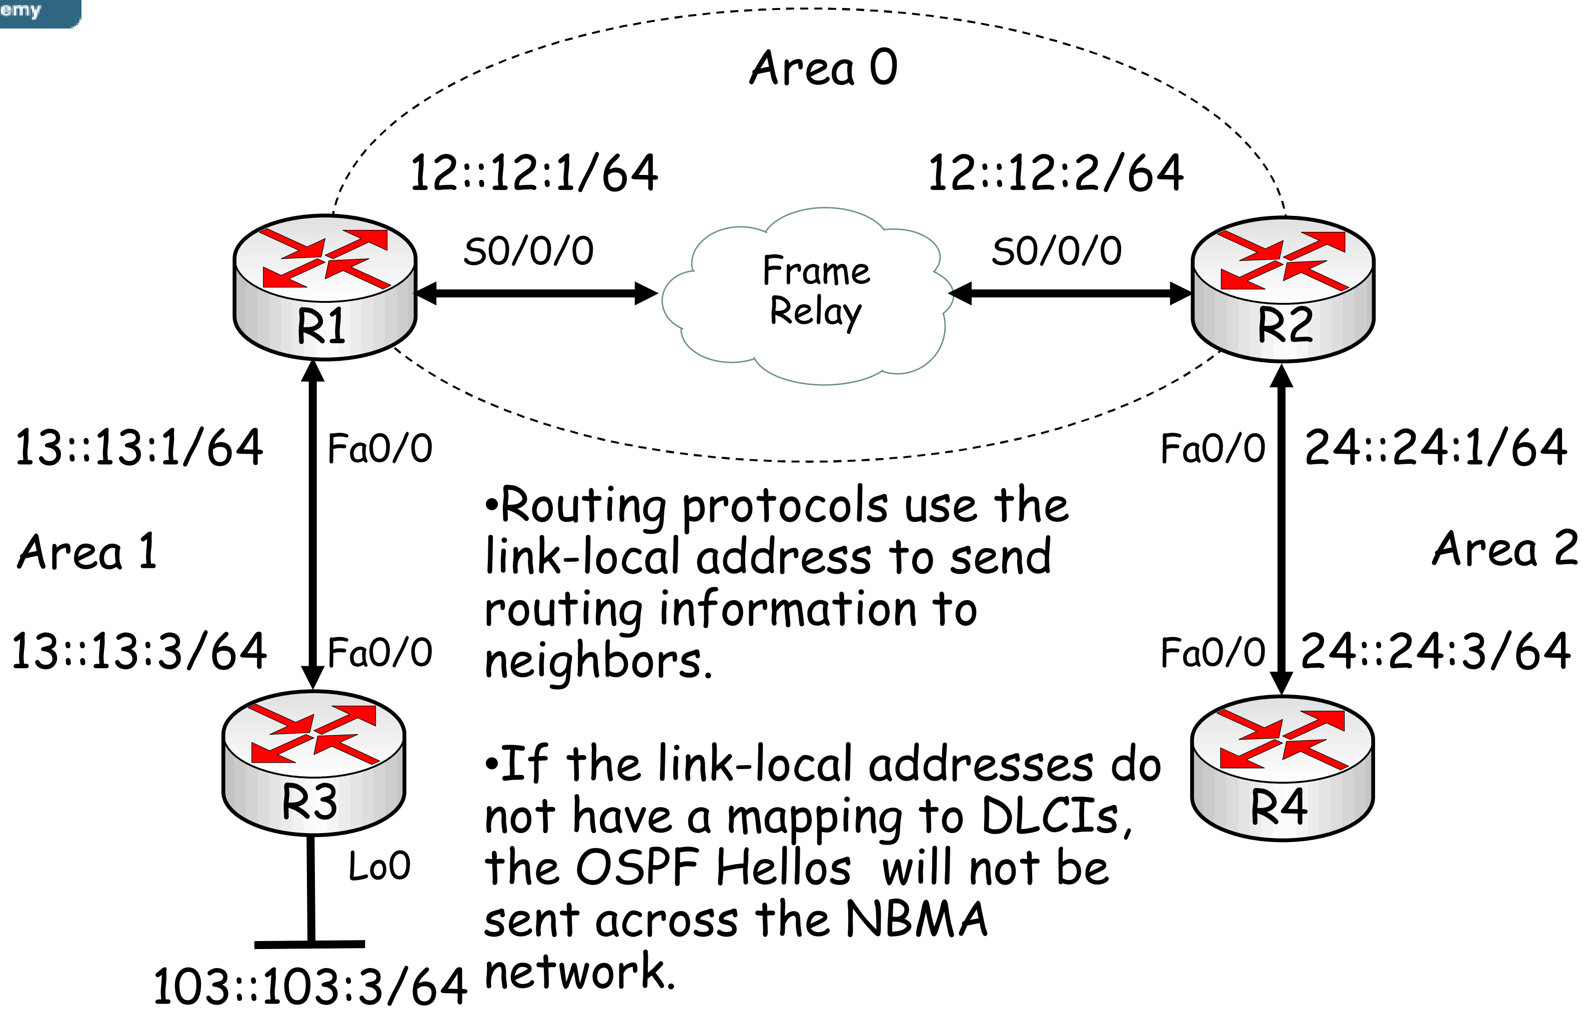

- Routing protocols use the link-local address to send routing information to neighbors.
- If the link-local addresses do not have a mapping to DLCIs, the OSPF Hellos will not be sent across the NBMA network.

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

RIPng Configuration

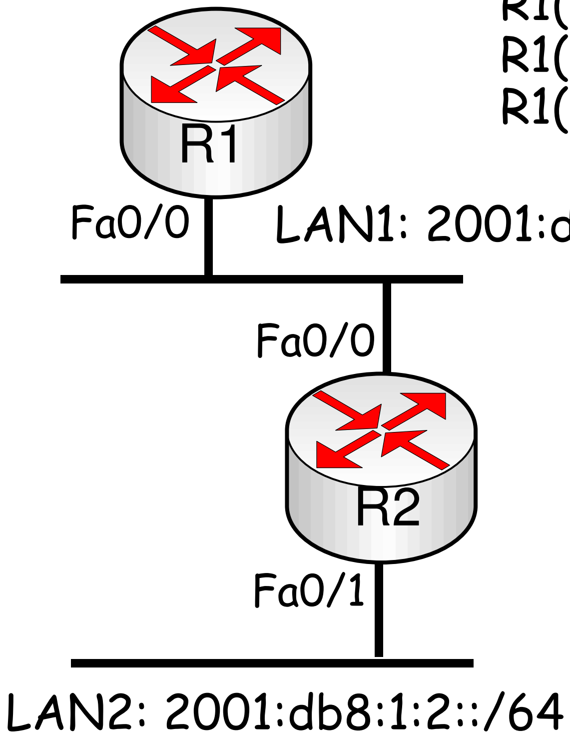

R1(config)#ipv6 unicast-routing
R1(config)#ipv6 cef
R1(config)#ipv6 router rip RT0

R1(config)#interface fa0/0
R1(config-if)# ipv6 address 2001:db8:1:1::/48 eui-64
R1(config-if)# ipv6 rip RT0 enable

R2(config)#ipv6 unicast-routing
R2(config)#ipv6 cef
R2(config)#ipv6 router rip RT0

R1(config)#interface fa0/0
R1(config-if)# ipv6 address 2001:db8:1:1::/48 eui-64
R1(config-if)# ipv6 rip RT0 enable
R1(config-if)#interface fa0/1
R1(config-if)# ipv6 address 2001:db8:1:2::/48 eui-64
R1(config-if)# ipv6 rip RT0 enable

IPv6 Redistribution

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

R2(config)#ipv6 router rip RIPoTU
R2(config-rtr)#redistribute rip RIPoFR metric 10

IPv6 Diagnostic Tools

- *debug ipv6 routing:* display debugging messages for IPv6 routing table updates and route cache updates.

- *debug ipv6 nd:* display debugging messages for IPv6 Internet Control Message Protocol (ICMP) ND transactions.

- *debug ipv6 packet:* display debugging messages for IPv6 packets. The debugging information includes packets received, generated, and forwarded. Note that fast-switched packets do not generate messages.

- *show ipv6 interface:* displays the usability status of interfaces configured for IPv6 or to validate the IPv6 status of an interface and its configured addresses.

- *show ipv6 routers:* This is an IPv6 specific command (doesn't have an IPv4 counterpart) you can use to display 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.

Chapter 6 – Troubleshooting Addressing Services
Objectives

- Describe NAT & PAT operation & troubleshooting techniques.
- Describe DHCP operation & troubleshooting techniques.
- Describe the different methods of IPv6 address assignment.
- Explain the operation of OSPFv3 and RIPng.
- Describe typical IPv6 troubleshooting techniques.

Any Questions?