IPv6 on router

Tsegereda Beyene and Benoit Lourdelet

tbeyene@cisco.com, blourdel@cisco.com

Agenda

- Router Basis
  - Addressing
  - Neighbor discovery
  - Renumbering
- Enabling IPv6 on Ethernet with EUI-64
- Enabling Static Routes
- Enabling RIPng
- Enabling OSPFv3
- Setup an IPv6/IPv4 tunnel : 6to4
- Enabling MP-BGP on the tunnel

IPv6 Router addresses

- **A router is required to recognize all addresses that a host is required to recognize, plus the following addresses as identifying itself:**
  - **The Subnet-Router Anycast Addresses for all interfaces for which it is configured to act as a router.**
  - **All other Anycast Addresses with which the router has been configured.**
  - **The All-Routers Multicast Addresses**

Enabling IPv6 on Cisco

- **To enable IPv6 on a Cisco router, you must:**
  - **Enable IPv6 traffic forwarding**
    - **ipv6 unicast-routing**
  - **Enable IPv6 on the interface(s) by configuring an IPv6 address on the interface**
    - **ipv6 address <ipv6addr>[/<prefix-length>]**
    - **ipv6 enable**
      - **Can be used, but only for link-local addresses**

IOS IPv6 Addressing Examples

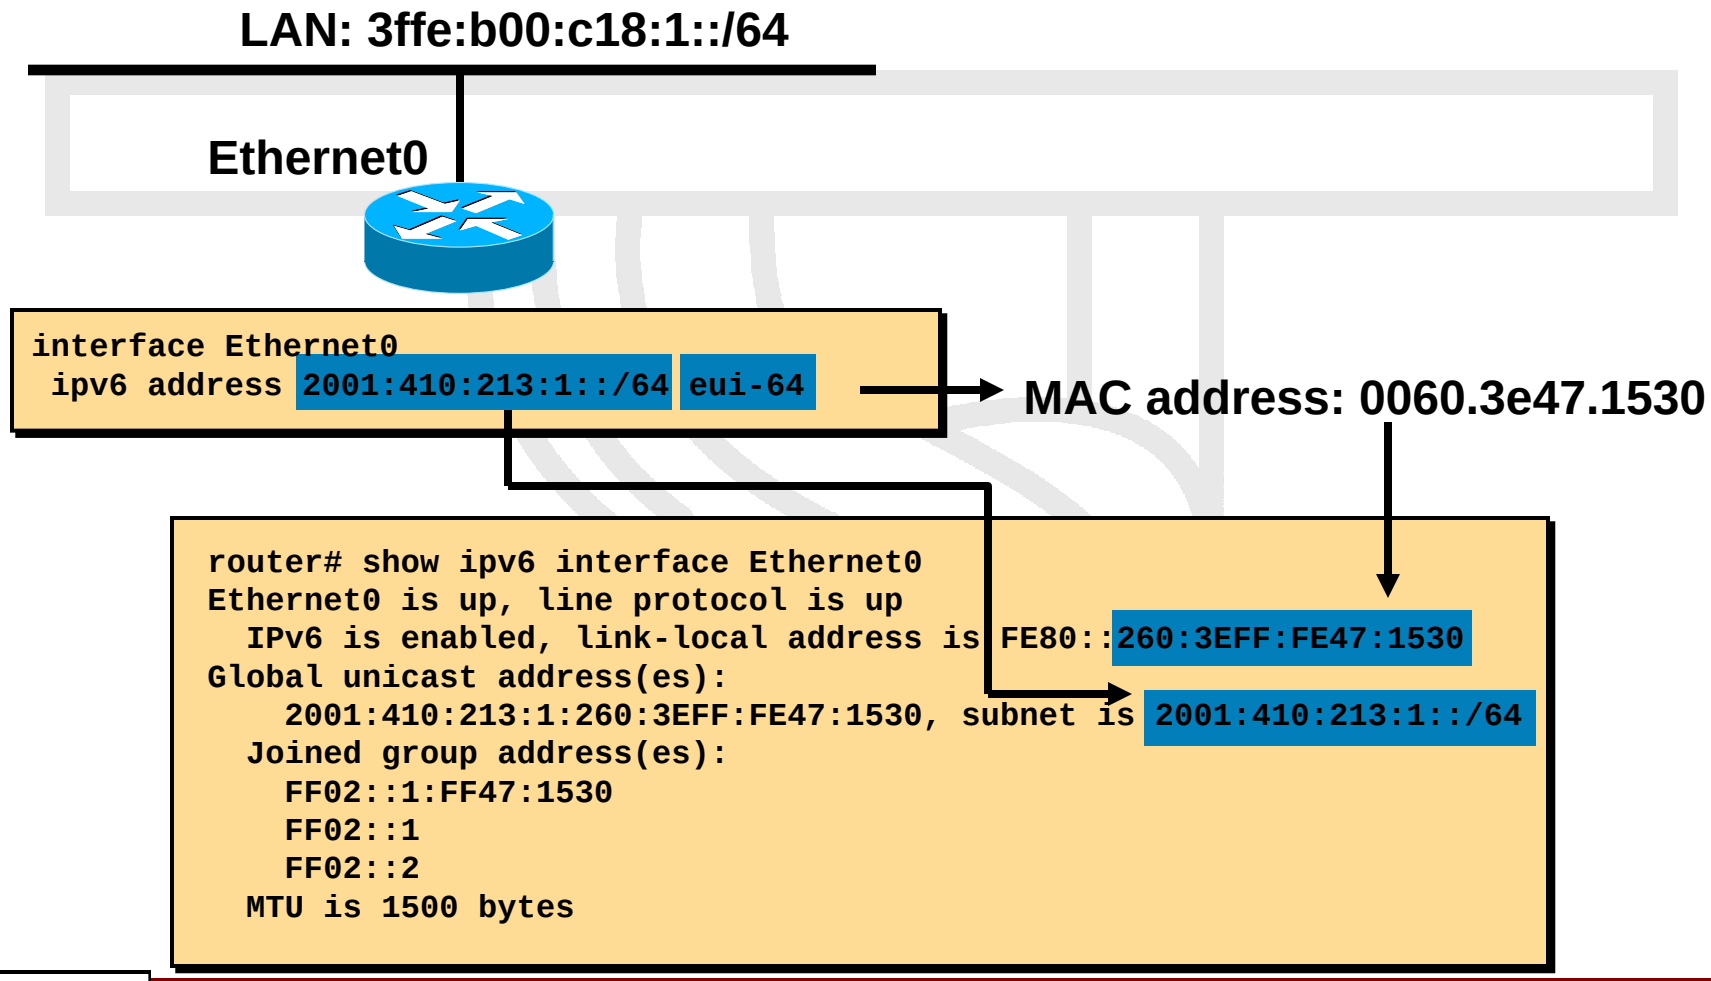

Configuring Interfaces Objectives

- Enable IPv6 on the router
- Configure an Ethernet interface for IPv6
- Identify which addresses are used on a router
- Test basic link-local connectivity

Configuring Interfaces on a Router

- **Configuring the IPv6 Ethernet Interface**
  - **Enable IPv6 on the router.**
  - **Enable IPv6 on the LAN interface**
- **Link-local Addresses**
  - **Find the link-local address of the Ethernet interface.**
  - **Multicast addresses ?**

Lab 1 – Configuring Interfaces Router

```
R1#sh ipv6 int e0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  No global unicast address is configured
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF3A:8B18
  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.
R1#
```

Configuring Interfaces Connectivity

```
R1#sh ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
FE80::210:A4FF:FE91:43B9                    0 0010.a491.43b9  REACH Ethernet0

R1#
```

```
R1#sh ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
FE80::210:A4FF:FE91:43B9                    0 0010.a491.43b9  STALE Ethernet0

R1#
```

- **Verifying the IPv6 Link-local Connectivity**

  From the PC (ping6)

  or from the router using the link-local address

- **Verify the list of IPv6 neighbors.**

Neighbor Discovery

ICMPv6

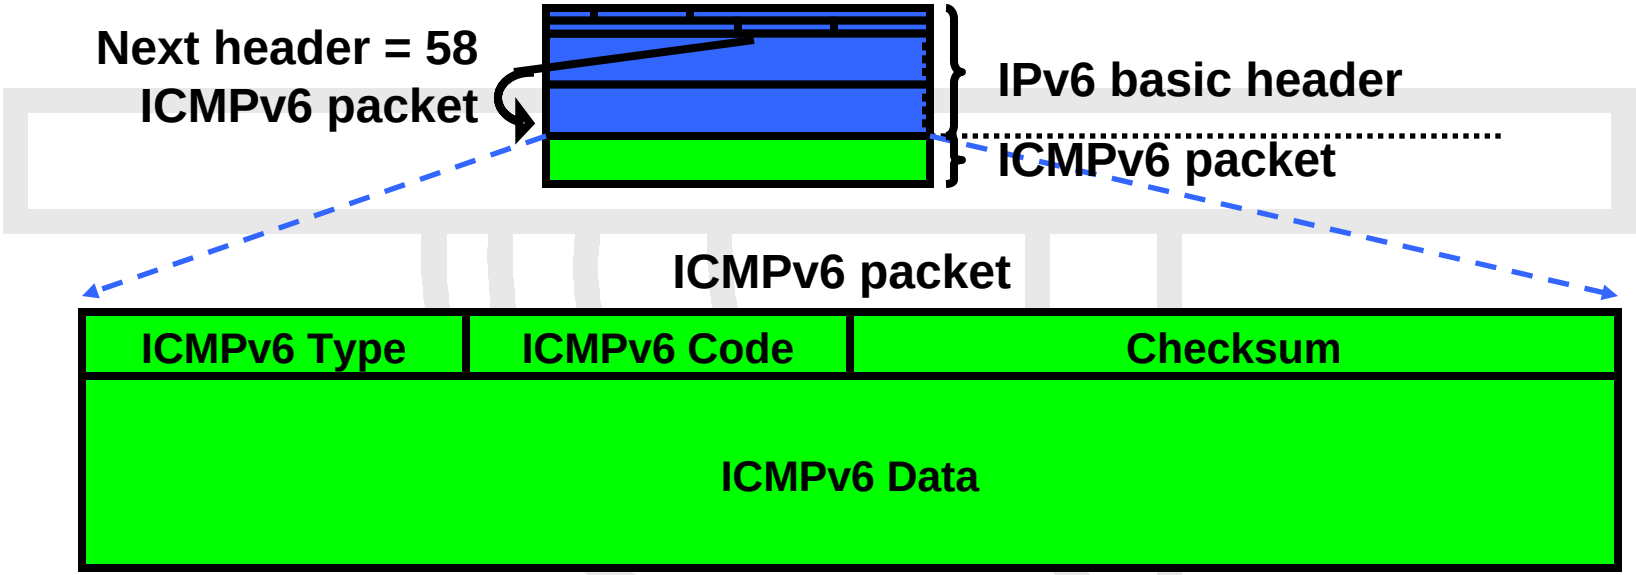

- **ICMPv6 is similar to IPv4:**
  - **Provides diagnostic and error messages**
  - **Is used for path MTU discovery**

ICMPv6

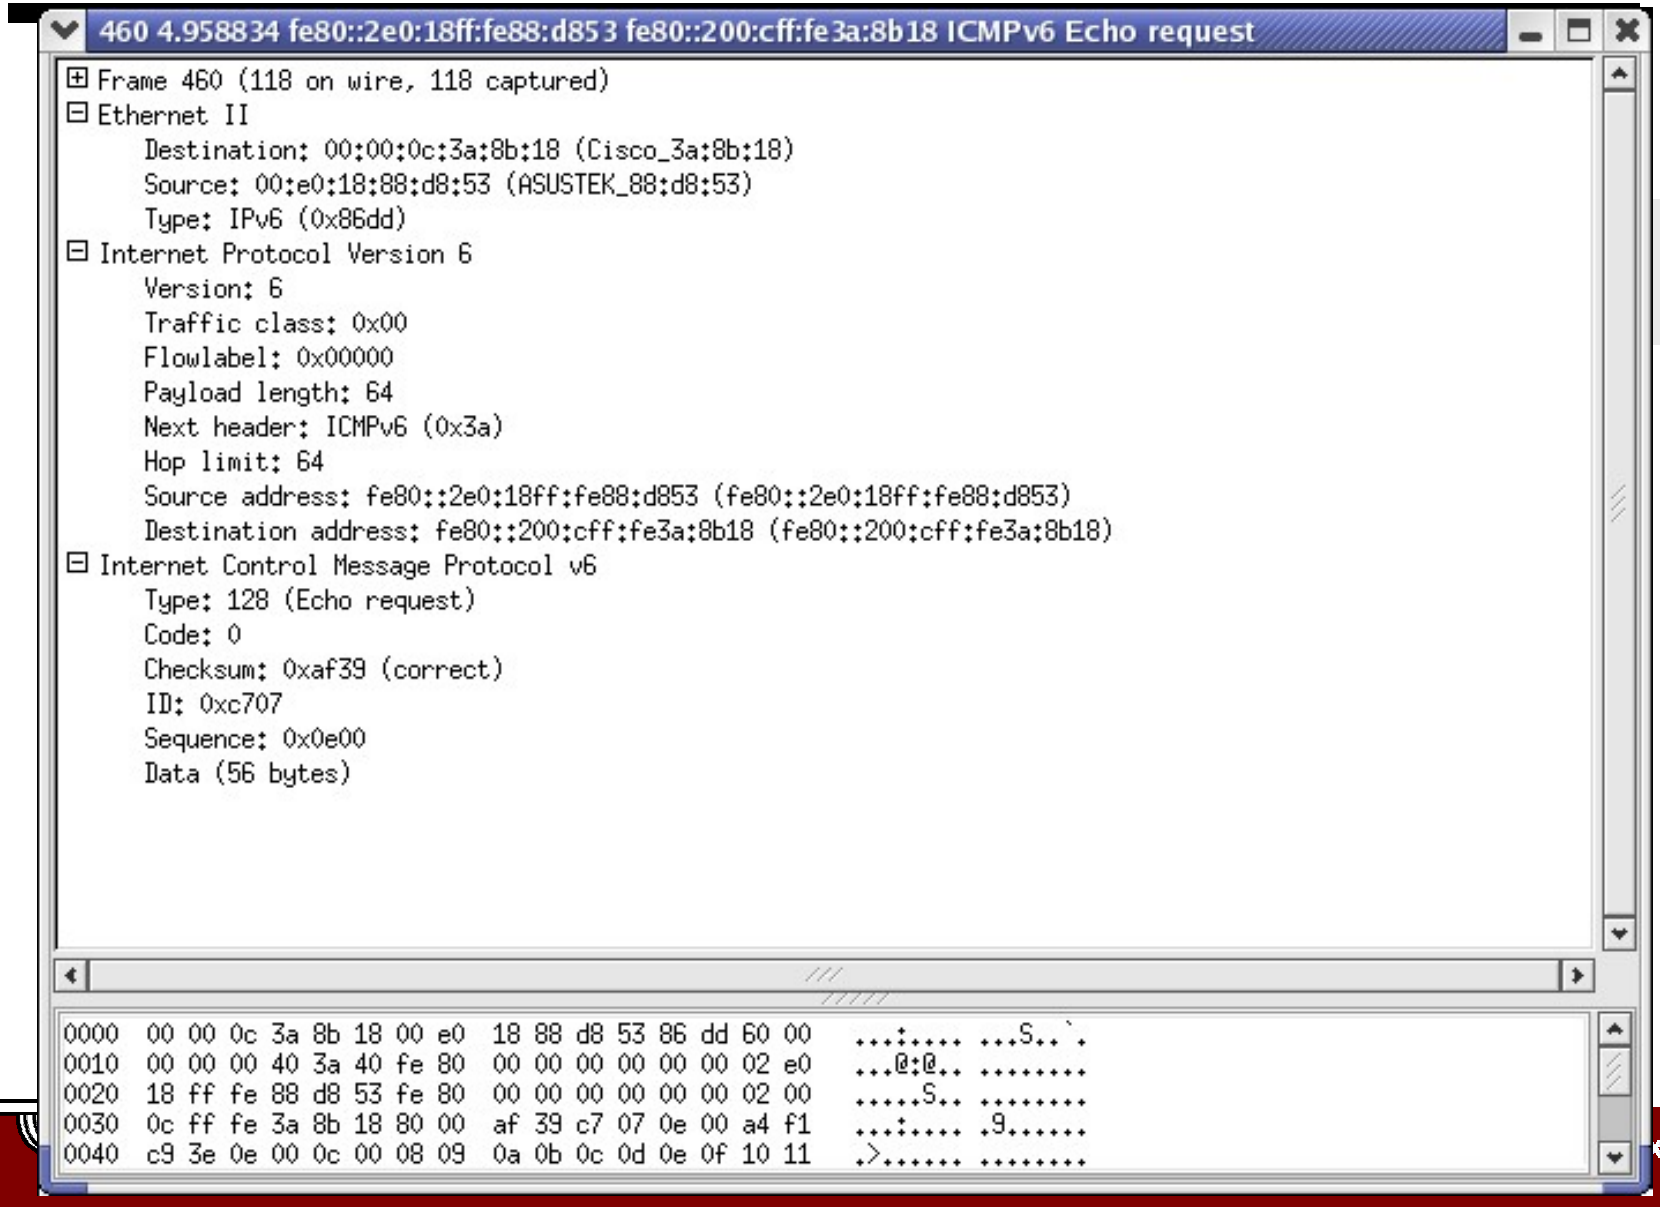
460 4.958834 fe80::2e0:18ff:fe88:d853 fe80::200:cff:fe3a:8b18 ICMPv6 Echo request
Frame 460 (118 on wire, 118 captured)
Ethernet II
Destination: 00:00:0c:3a:8b:18 (Cisco_3a:8b:18)
Source: 00:e0:18:88:d8:53 (ASUSTEK_88:d8:53)
Type: IPv6 (0x86dd)
Internet Protocol Version 6
Version: 6
Traffic class: 0x00
Flowlabel: 0x00000
Payload length: 64
Next header: ICMPv6 (0x3a)
Hop limit: 64
Source address: fe80::2e0:18ff:fe88:d853 (fe80::2e0:18ff:fe88:d853)
Destination address: fe80::200:cff:fe3a:8b18 (fe80::200:cff:fe3a:8b18)
Internet Control Message Protocol v6
Type: 128 (Echo request)
Code: 0
Checksum: 0xaf39 (correct)
ID: 0xc707
Sequence: 0x0e00
Data (56 bytes)
0000 00 00 0c 3a 8b 18 00 e0 18 88 d8 53 86 dd 60 00
0010 00 00 00 40 3a 40 fe 80 00 00 00 00 00 00 02 e0
0020 18 ff fe 88 d8 53 fe 80 00 00 00 00 00 00 02 00
0030 0c ff fe 3a 8b 18 80 00 af 39 c7 07 0e 00 a4 f1
0040 c9 3e 0e 00 0c 00 08 09 0a 0b 0c 0d 0e 0f 10 11

Path MTU Discovery

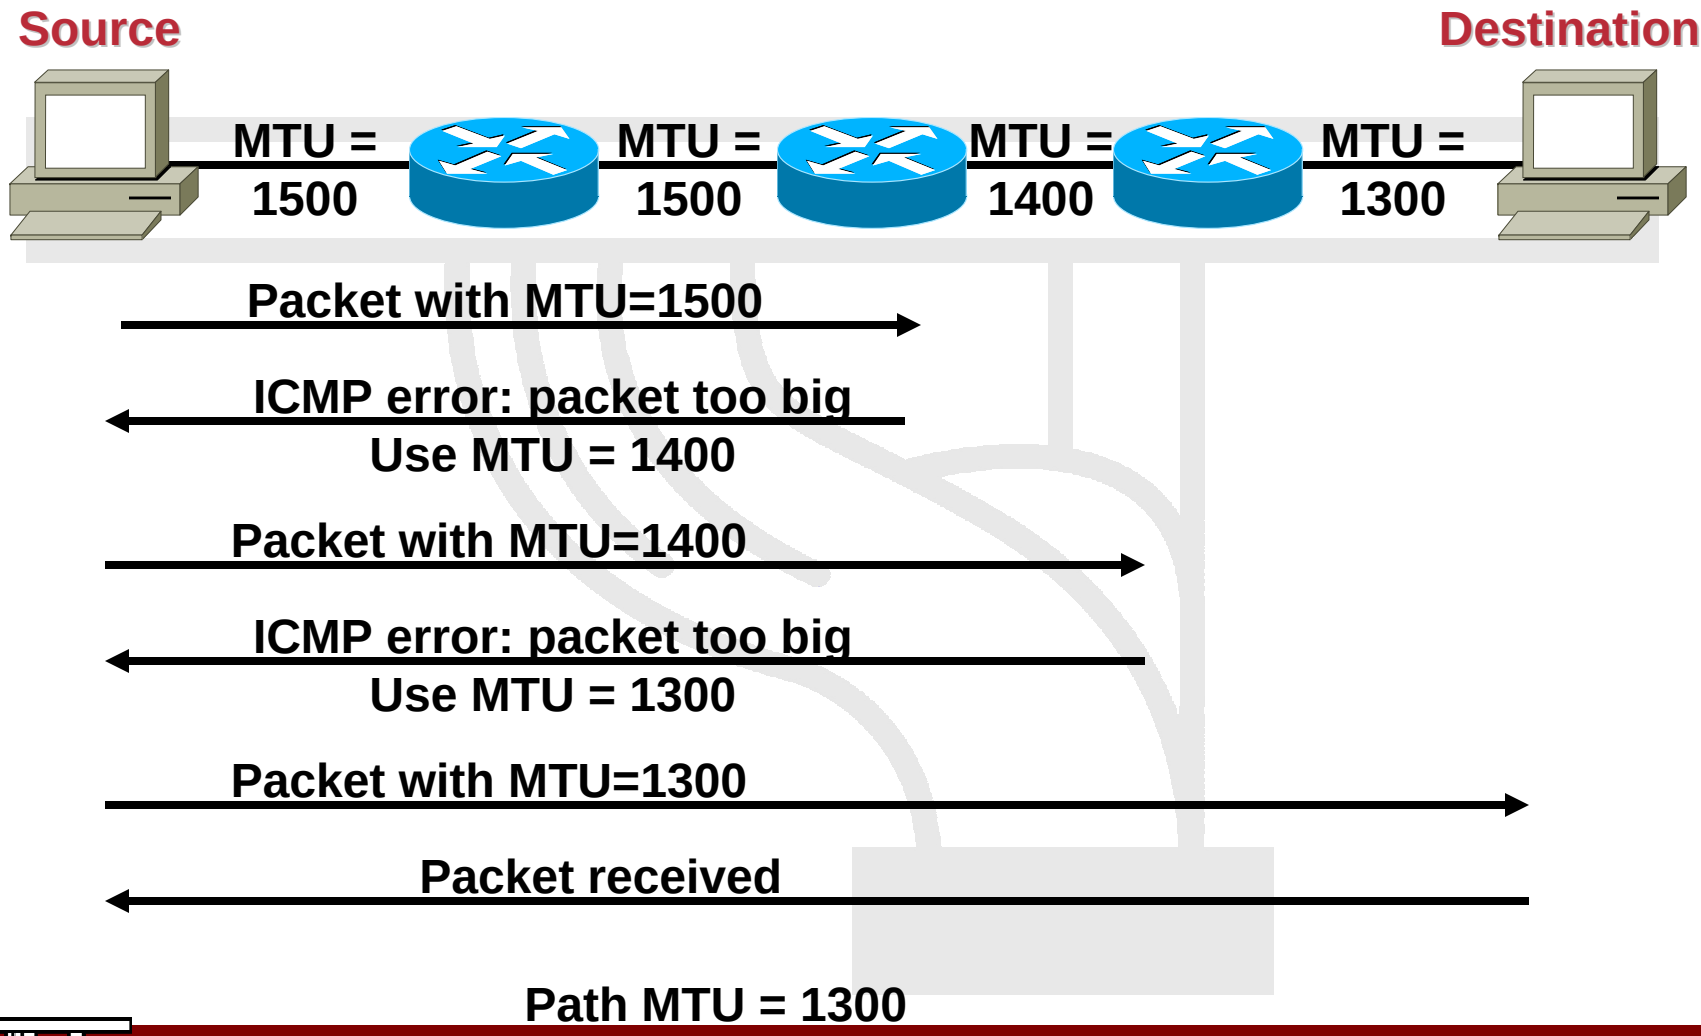

Neighbor Discovery (RFC 2461)

- **Protocol built on top of ICMPv6 (RFC 2463)**
  - **Combination of IPv4 protocols (ARP, ICMP,…)**
- **Neighbor Discovery:**
  - **Determines the link-layer address of a neighbor on the same link**
  - **Duplicate Address Detection**
  - **Finds neighbor routers, Keeps track of neighbors**
- **Defines 5 ICMPv6 packet types**
  - **Router Solicitation / Router Advertisements**
  - **Neighbor Solicitation / Neighbor Advertisements**
  - **Redirect**

Neighbor Discovery (RFC 2461)

- **Defined mechanisms between nodes attached on the same link**
  - **Router discovery**
  - **Prefix discovery**
  - **Parameters discovery, ie: link MTU, hop limit,…**
  - **Address autoconfiguration**
  - **Address Resolution (same function as ARP)**
  - **Next-hop determination**
  - **Neighbor Unreachability Detection (useful for default routers)**
  - **Duplicate Address Detection**
  - **Redirect**

Solicited-Node Multicast Address

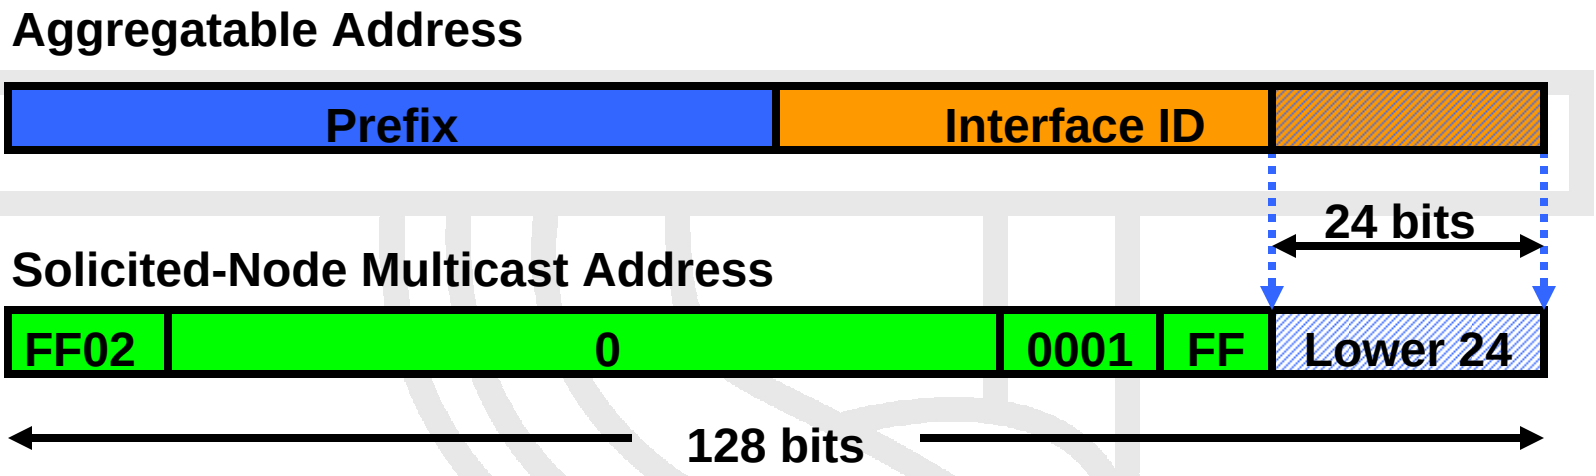

- **A solicited-node address is a:**

  Multicast address with a link-local scope

  Formed by a prefix and the rightmost 24 bits of the aggregatable address

Neighbor Solicitation

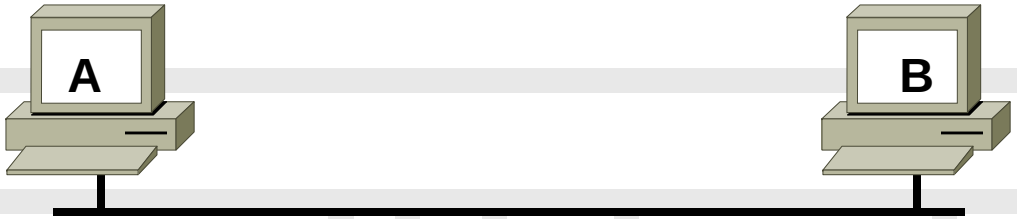

Neighbor Solicitation

ICMP type = 135
Src = A
Dst = Solicited-node multicast of B
Data = link-layer address of A
Query = what is your link address?

Neighbor Advertisement

ICMP type = 136
Src = B
Dst = A
Data = link-layer address of B

A and B can now exchange packets on this link

Neighbor Solicitation

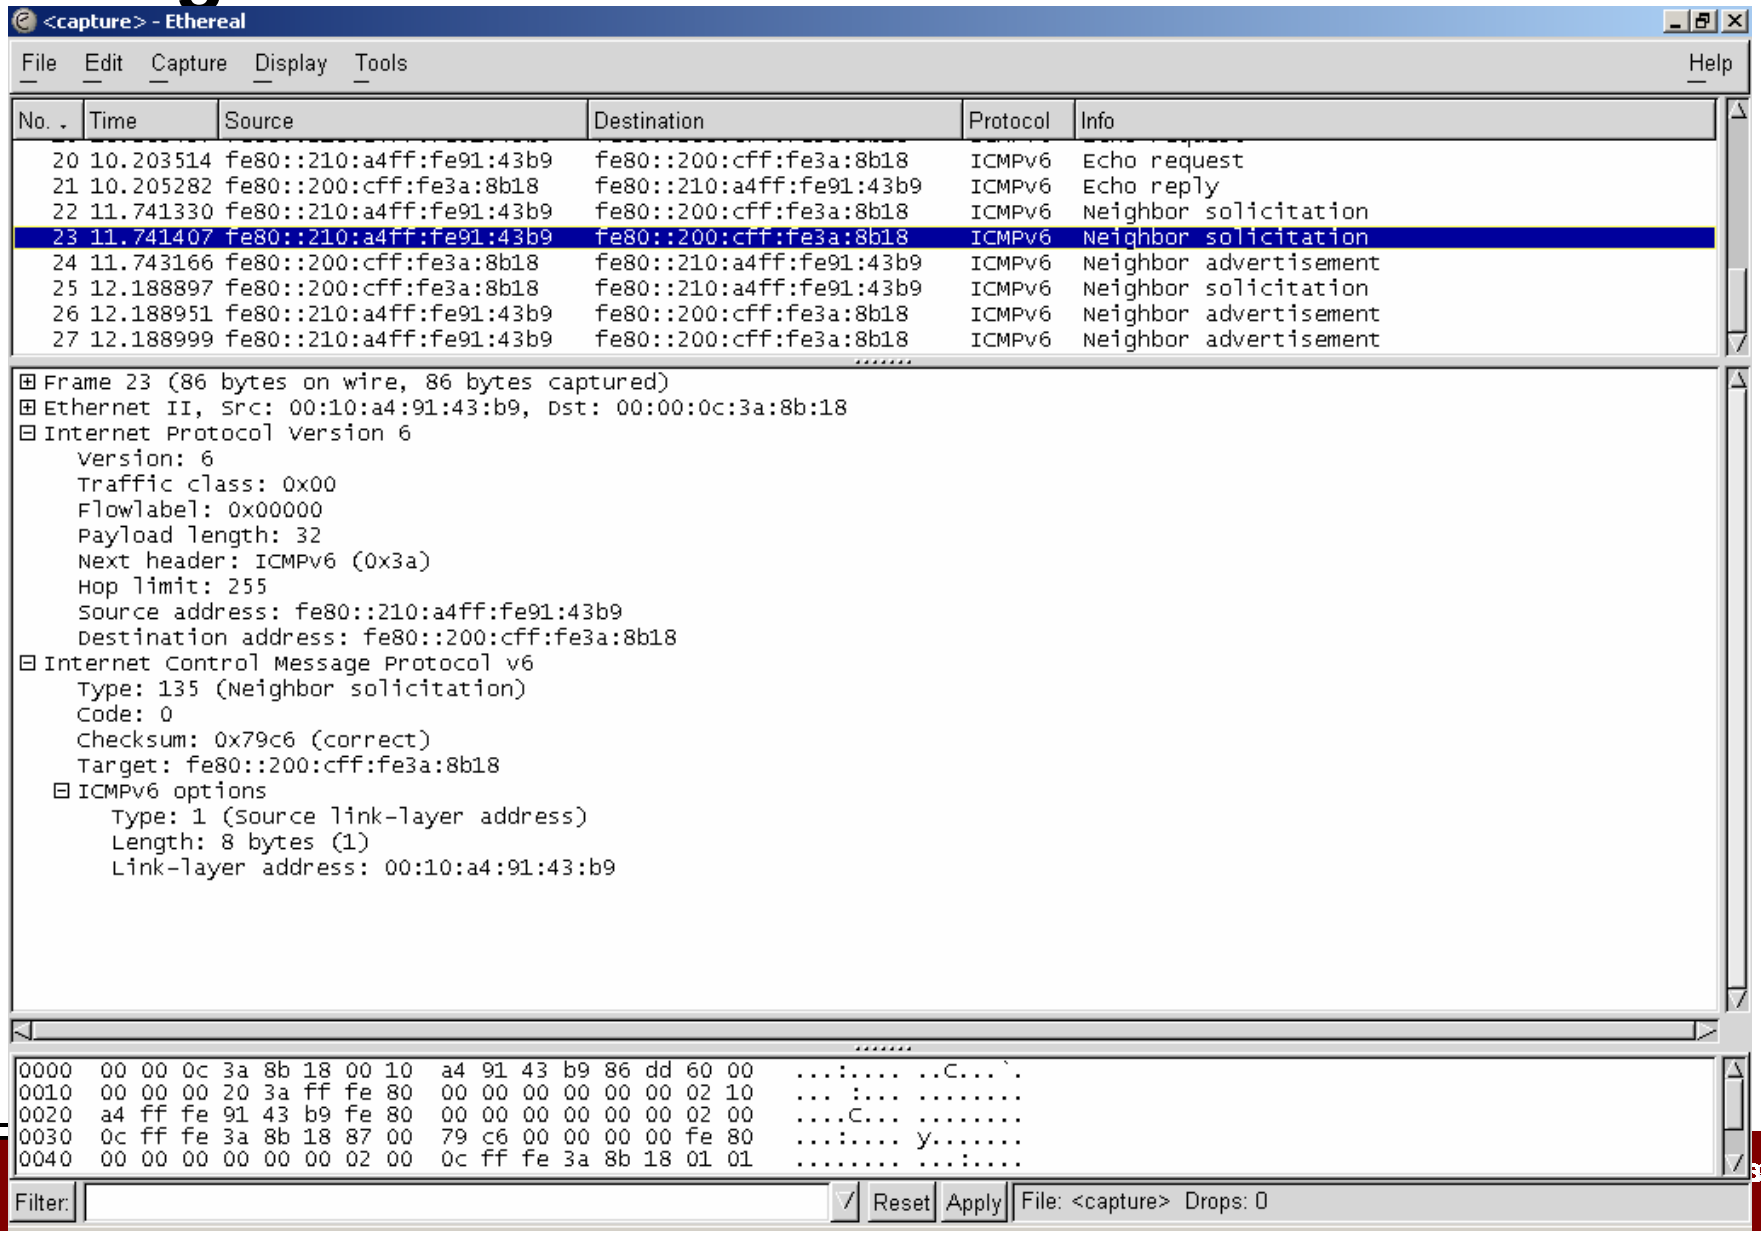

Neighbor Advertisement

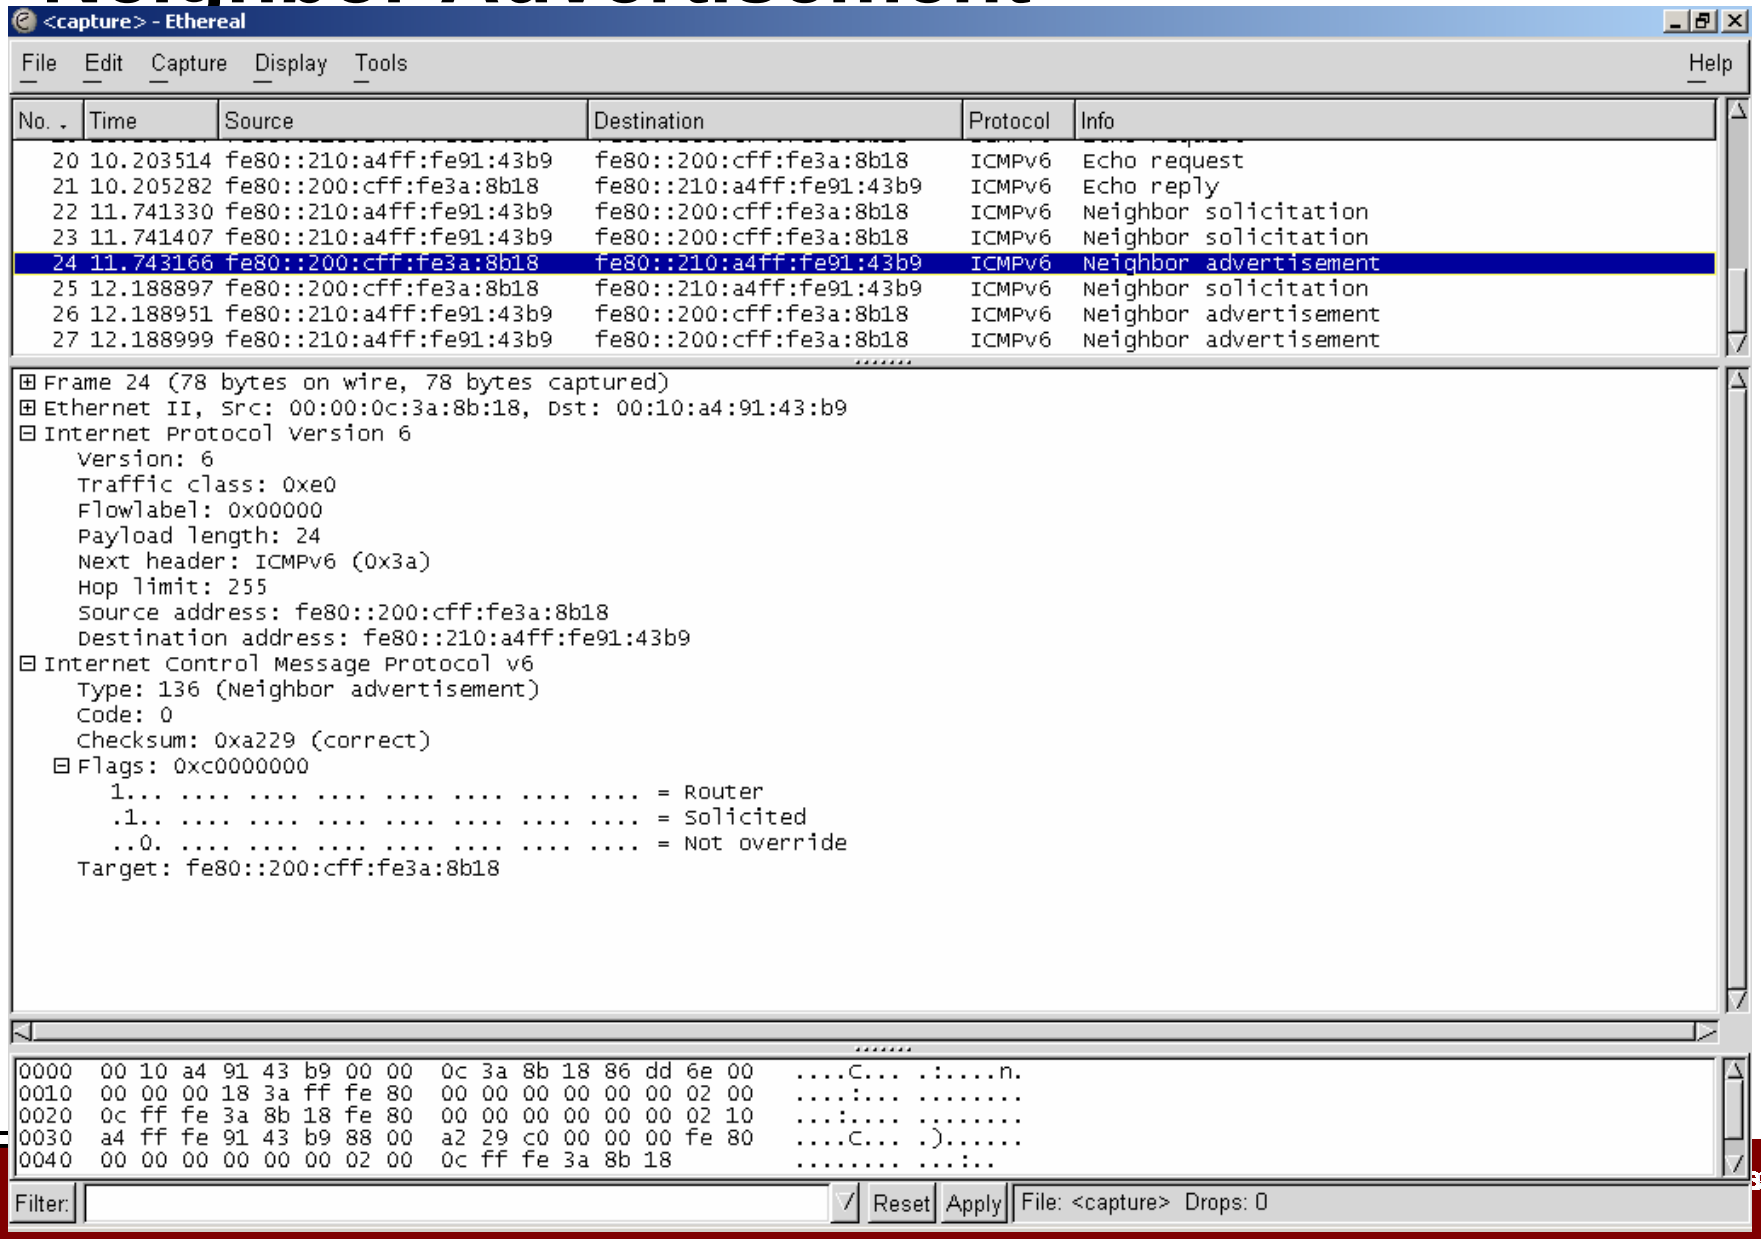

Router Advertisements (RA)

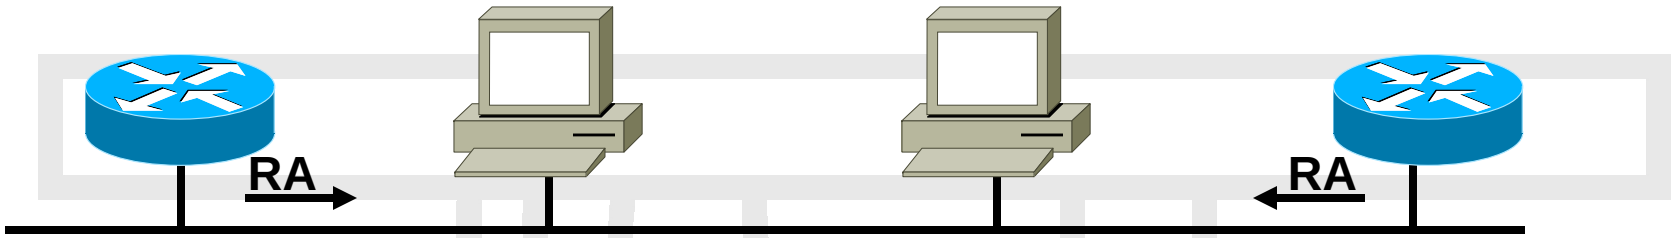

RA packet definitions:

ICMP Type = 134

Src = Router Link-local Address

Dst = All-nodes multicast address

Data= options, prefix, lifetime, autoconfig flag

- **Routers send periodic Router Advertisements (RA) to the all-nodes multicast address.**

Router Solicitations

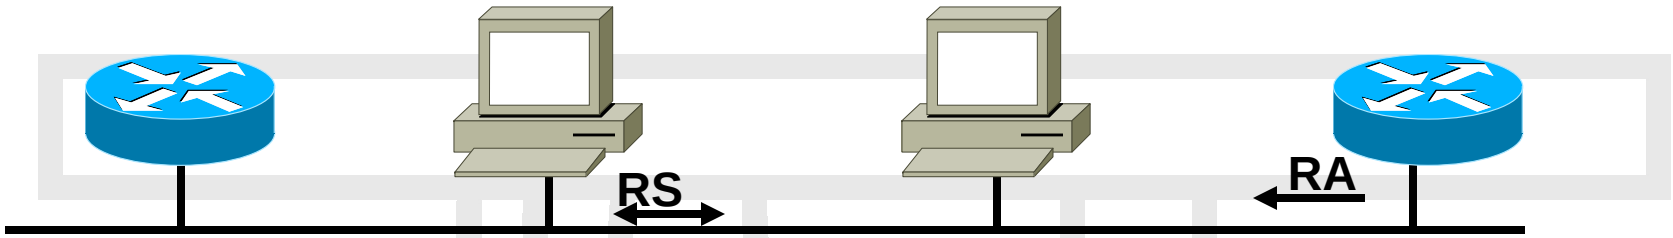

RS packet definitions:

ICMP Type = 133

Src = Unspecified Address

Dst = All-routers multicast address

- **At boot time, nodes sends Router Solicitations to receive promptly Router Advertisements.**

Duplicate Address Detection

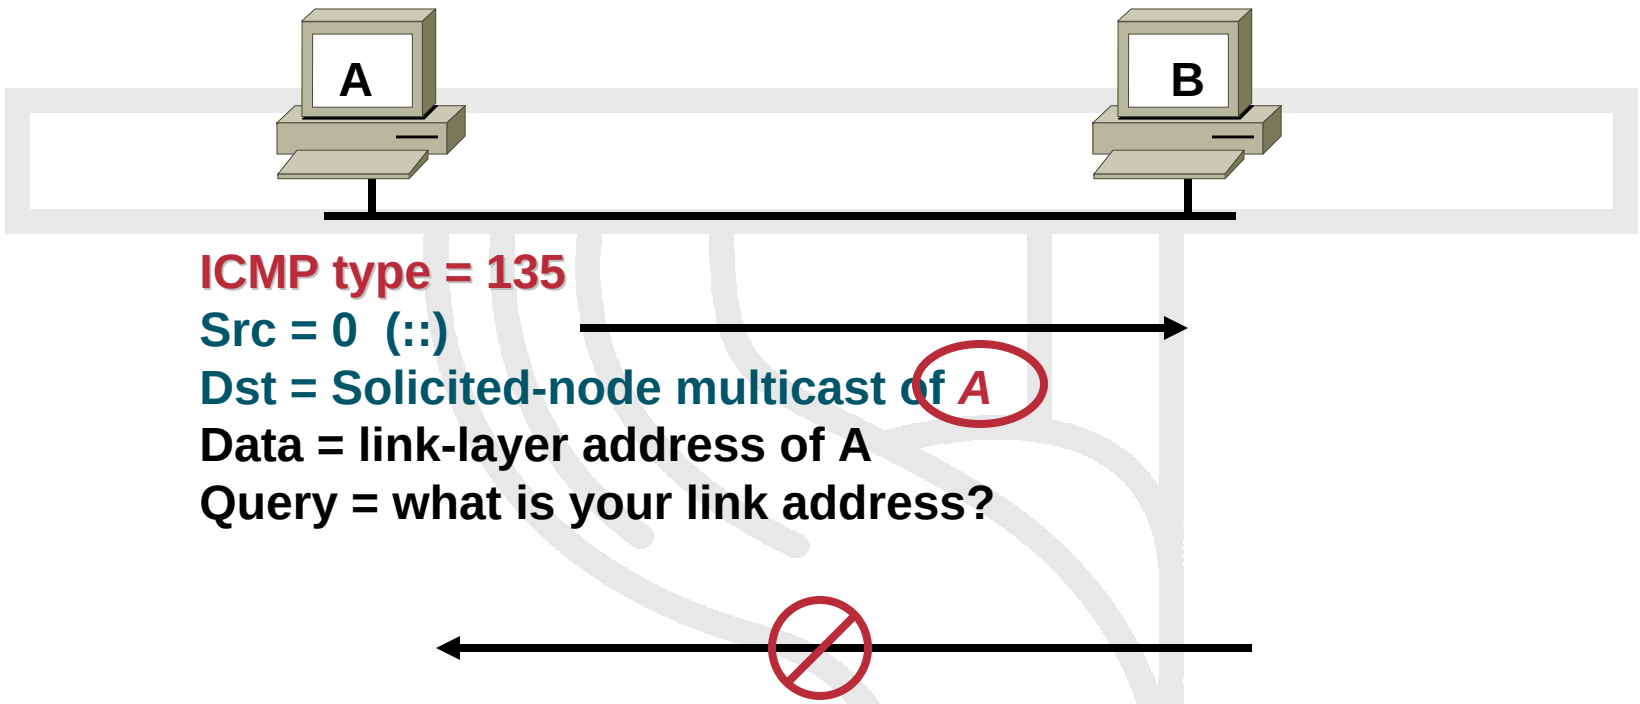

- **Duplicate Address Detection (DAD) uses neighbor solicitation to verify the existence of an address to be configured.**

IPv6 Auto-Configuration

- **Stateless** (RFC2462)

  Host autonomously configures its own Link-Local address

  Router solicitation are sent by booting nodes to request RAs for configuring the interfaces.

- **Stateful**

  DHCPv6

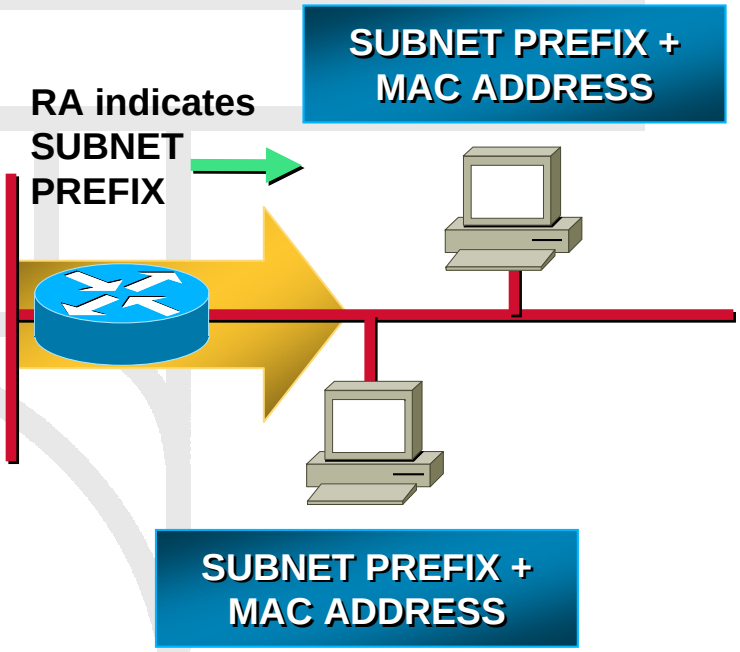

At boot time, an IPv6 host build a Link-Local address, then its global IPv6 address(es) from RA

Stateless Autoconfiguration

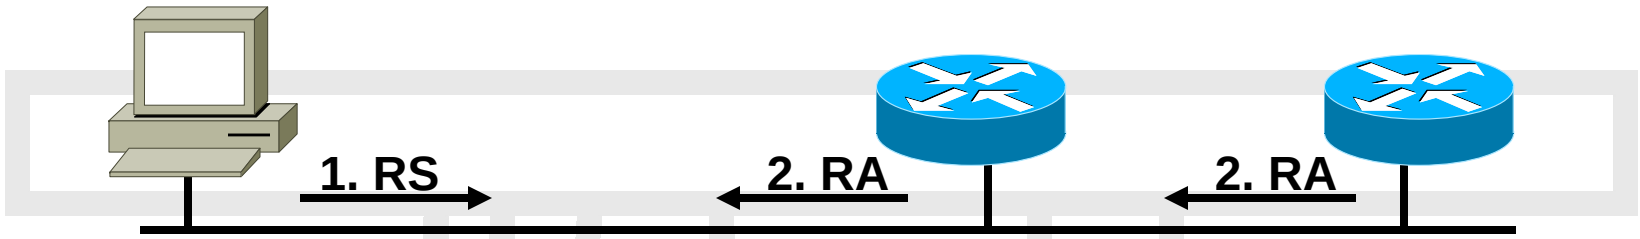

1. **RS**:

   ICMP Type = 133

   Src = ::

   Dst = All-Routers multicast Address

   query= please send RA

2. **RA**:

   ICMP Type = 134

   Src = Router Link-local Address

   Dst = All-nodes multicast address

   Data= options, prefix, lifetime, autoconfig flag

- **Router solicitations are sent by booting nodes to request RAs for configuring the interfaces.**

Address Lifetime

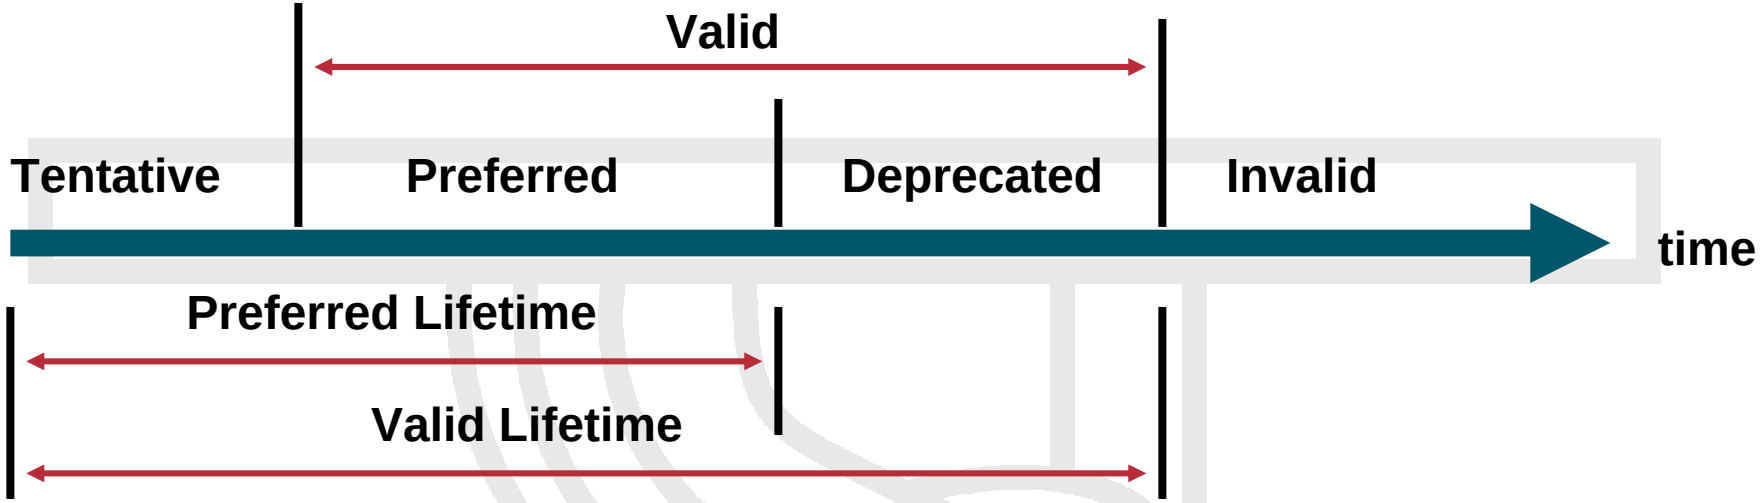

- **Tentative** : the address is in the process of being verified as unique
- **Preferred** : a node can send and receive unicast traffic to and from a preferred address
- **Deprecated** : the address is still valid, but using it for new communication is discouraged
- **Invalid** : the address can no longer send unicast traffic to or receive it from a node. An address enters this state after the valid lifetime expires.

DHCPv6

- Updated version of DHCP for IPv4.
- Supports new addressing.
- Enables more control than stateless auto-configuration.
- Can be used for renumbering.
- Can be used for automatic domain name registration of hosts using dynamic DNS.

DHCPv6 - Process

- **Same as in IPv4, but:**

  Client first detect the presence of routers on the link.

  If found, then examines router advertisements to determine if DHCP can be used.

  If no router found or if DHCP can be used, then

    DHCP Solicit message is sent to the All-DHCP-Agents multicast address

    Using the link-local address as the source address

Renumbering

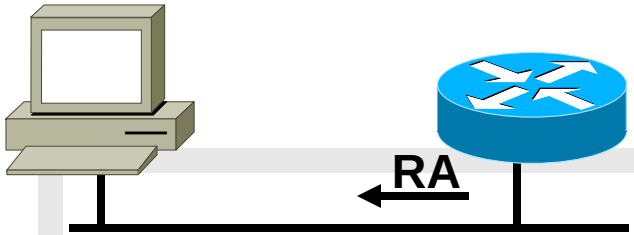

RA packet definitions:

ICMP Type = 134

Src = Router Link-local Address

Dst = All-nodes multicast address

Data= 2 prefixes:

Current prefix (to be deprecated) with short lifetime

New prefix (to be used) with normal lifetime

- **Renumbering is achieved by modifying the RA to announce the old prefix with a short lifetime and the new prefix.**

IOS Prefix Renumbering Scenario

ipv6 nd prefix-advertisement <routing-prefix>/<length> <valid-lifetime> <preferred-lifetime> [onlink | autoconfig]

Router configuration before renumbering:

```
interface Ethernet0
  ipv6 nd prefix-advertisement 3ffe:b00:c18:1::/64 43200 43200 onlink autoconfig
```

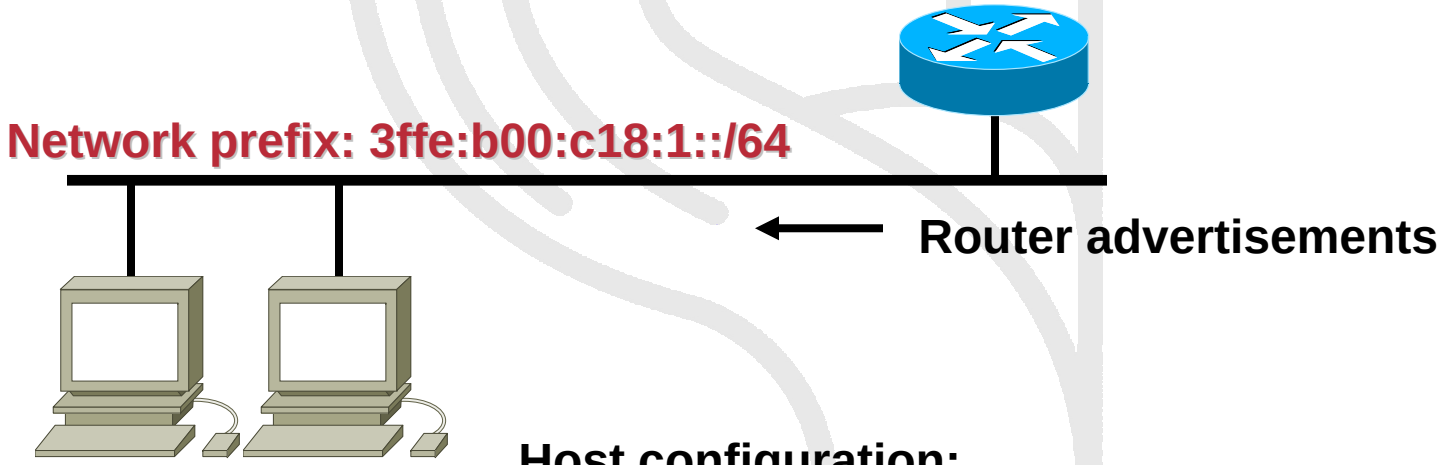

Autoconfiguring IPv6 hosts

Host configuration:

```
preferred address 3ffe:b00:c18:1:260:8ff:fede:8fbe
```

IOS Prefix Renumbering Scenario

ipv6 nd prefix-advertisement <routing-prefix>/<length> <valid-lifetime> <preferred-lifetime> [onlink | autoconfig]

Router configuration after renumbering:

```
interface Ethernet0
  ipv6 nd prefix-advertisement 3ffe:b00:c18:1::/64 43200 0 onlink autoconfig
  ipv6 nd prefix-advertisement 3ffe:b00:c18:2::/64 43200 43200 onlink autoconfig
```

Autoconfiguring IPv6 hosts

Host configuration:

```
deprecated address 3ffe:b00:c18:1:260:8ff:fede:8fbe
preferred address 3ffe:b00:c18:2:260:8ff:fede:8fbe
```

Neighbor Discovery

- Configure a router to send router advertisements
- Configure a network to use global addresses
  - Configure EUI-64 addresses
  - Configure with a fixed manual interface-id
- Renumber a network

Neighbor Discovery

- Configure a global IPv6 address on the LAN interface using the /64 prefix of your subnet and the EUI-64 format for the host part of the address

Neighbor Discovery

- Look at the config of the PC and see if it has received a global address.
- Enable the debugging mode for IPv6 Neighbor Discovery.
- Enable router advertisements by using the neighbor discovery command with the subnet prefix assigned to your LAN.
- Since we don’t want infinite lifetimes, use 5 minutes (300 seconds) for the lifetime (both preferred and valid).
- Verify connectivity using the ping command on the Cisco router to the PC using the new assigned address of the PC as the ping destination address.

Neighbor Discovery
Renumbering the Local Network

- **Configure the new address of the LAN interface by using the new global subnet prefix assigned to your LAN (See Topology slide 2)**

  Use a manual interface-id instead of EUI-64

- **Enable router advertisements by using the neighbor discovery command with the NEW subnet prefix assigned to your LAN**
- **Verify that your PC has a new address with the new prefix.**

Neighbor Discovery
Renumbering the Local Network

- Modify the Neighbor Advertisements for the old prefix by making the preferred lifetime equals to zero
- Verify that your PC now deprecates the use of the “old” address prefix and prefer the new one.
- Verify the connectivity with a ping from the Cisco to the PC using the new assigned address of the PC.

IOS Prefix Renumbering
Original Prefix 2006:1/64

```
!
interface Ethernet0
 ipv6 address 2006:1::1/64
 ipv6 nd ra-interval 60
 ipv6 nd prefix-advertisement 2006:1::/64 10000 6000 onlink autoconfig
!
```

IOS Prefix Renumbering Example
Original Prefix 2006:1/64

```
R1#sh ipv6 int ethernet 0
Ethernet0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::200:CFF:FE3A:8B18
  Global unicast address(es):
    2006:1::1, subnet is 2006:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF3A:8B18
    FF02::1:FF00:1
  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.
R1#
```

```
R1#sh int ethernet 0
Ethernet0 is up, line protocol is up
  Hardware is Lance, address is
0000.0c3a.8b18 (bia 0000.0c3a.8b18)
```

IOS Prefix Renumbering Example
Windows

```
D:\>ipv6 if 4
Interface 4 (site 1): Local Area Connection
  uses Neighbor Discovery
  link-level address: 00-10-a4-91-43-b9
    tentative address 2006:1::210:a4ff:fe91:43b9, 10000s/6000s (addrconf)
    tentative address fe80::210:a4ff:fe91:43b9, infinite/infinite
    multicast address ff02::1, 1 refs, not reportable
    multicast address ff02::1:ff91:43b9, 2 refs, last reporter, 9 seconds until
report
  link MTU 1372 (true link MTU 1372)
  current hop limit 64
  reachable time 35500ms (base 30000ms)
  retransmission interval 1000ms
  DAD transmits 1

D:\>
```

IOS Prefix Renumbering Example
Windows

```
D:\>ipv6 if 4
Interface 4 (site 1): Local Area Connection
  uses Neighbor Discovery
  link-level address: 00-10-a4-91-43-b9
    preferred address 2006:1::210:a4ff:fe91:43b9, 9963s/5963s (addrconf)
    preferred address fe80::210:a4ff:fe91:43b9, infinite/infinite
    multicast address ff02::1, 1 refs, not reportable
    multicast address ff02::1:ff91:43b9, 2 refs, last reporter
  link MTU 1372 (true link MTU 1372)
  current hop limit 64
  reachable time 35500ms (base 30000ms)
  retransmission interval 1000ms
  DAD transmits 1

D:\>
```

IOS Prefix Renumbering Example
New Prefix 2008:1/64

```
!
interface Ethernet0
 ipv6 address 2006:1::1/64
 ipv6 nd ra-interval 60
 ipv6 nd prefix-advertisement 2006:1::/64 10000 0 onlink autoconfig
 ipv6 nd prefix-advertisement 2008:1::/64 10000 8000 onlink autoconfig
!
```

IOS Prefix Renumbering Example
Windows display

```
D:\>ipv6 if 4
Interface 4 (site 1): Local Area Connection
  uses Neighbor Discovery
  link-level address: 00-10-a4-91-43-b9
    preferred address 2008:1::210:a4ff:fe91:43b9, 9990s/7990s (addrconf)
    deprecated address 2006:1::210:a4ff:fe91:43b9, 9990s/0s (addrconf)
    preferred address fe80::210:a4ff:fe91:43b9, infinite/infinite
    multicast address ff02::1, 1 refs, not reportable
    multicast address ff02::1:ff91:43b9, 3 refs, last reporter
  link MTU 1372 (true link MTU 1372)
  current hop limit 64
  reachable time 35500ms (base 30000ms)
  retransmission interval 1000ms
  DAD transmits 1

D:\>
```

Lab setup

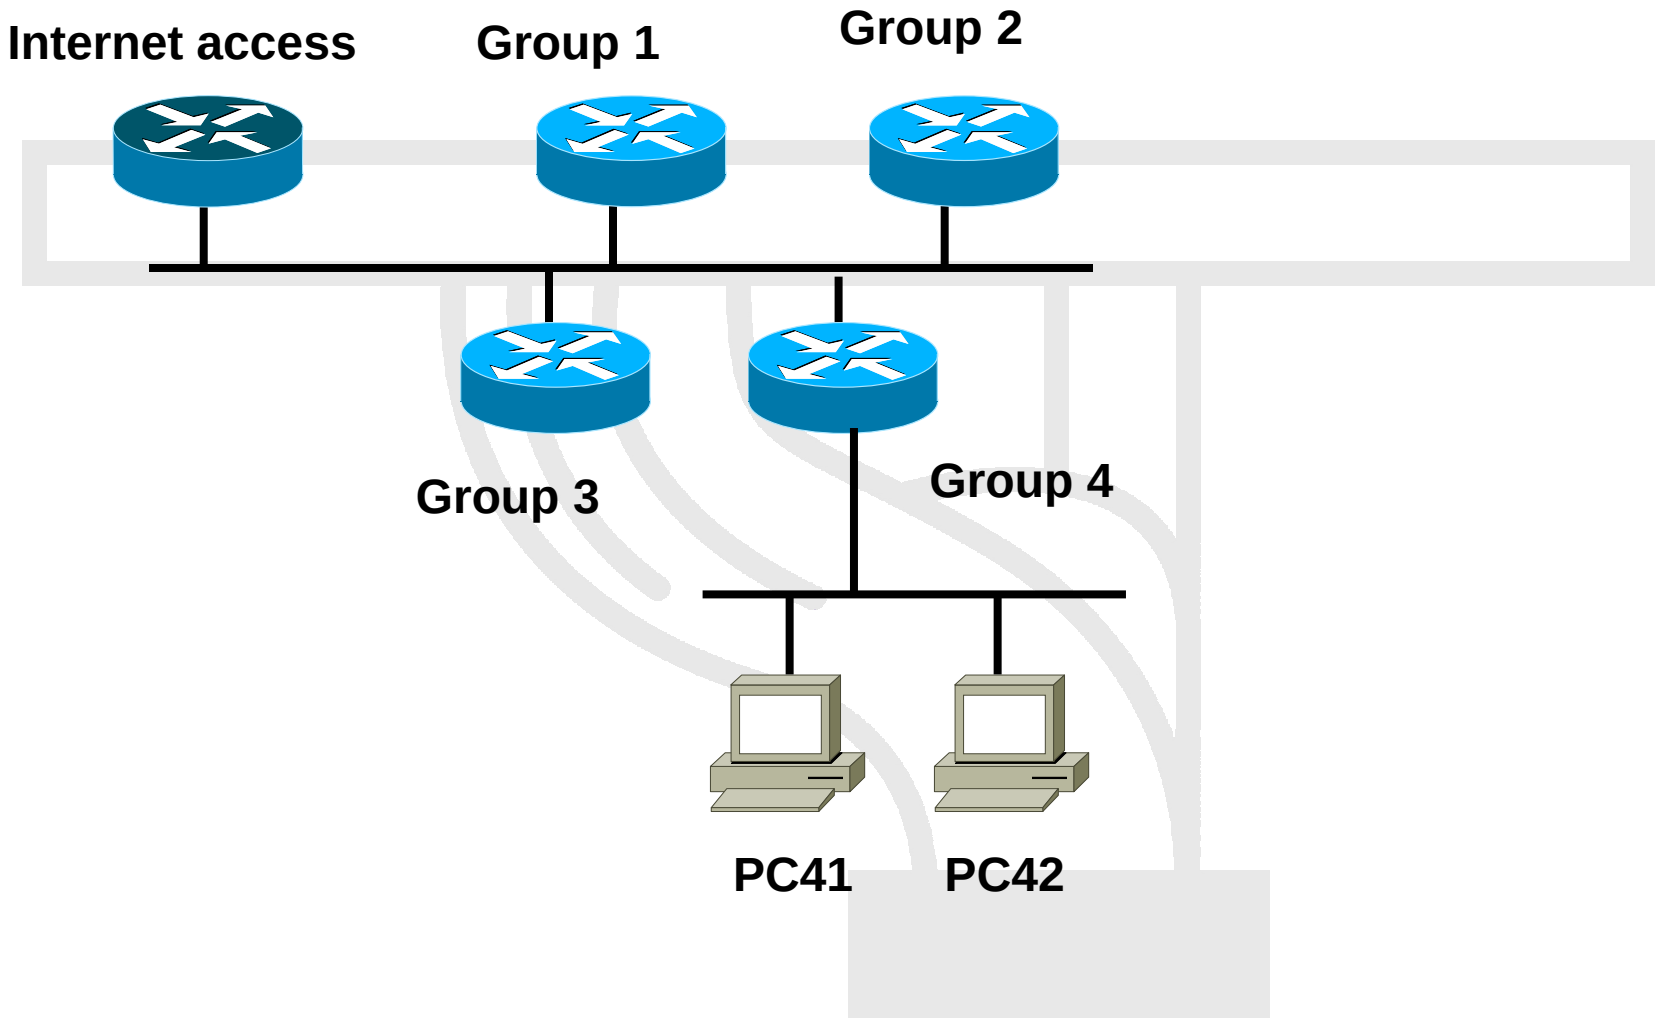

Enable IPv6 on Ethernet 0/0 with EUI-64

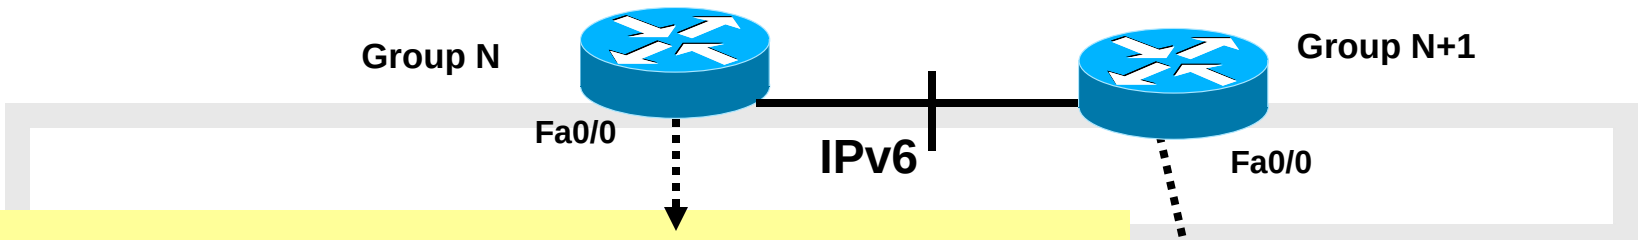


```
!
ipv6 unicast-routing
ipv6 cef
!
interface FastEthernet0/0
ipv6 address 2001:4F8:FFFE:100F::/64 eui-64
ipv6 address 2001:4F8:FFFE:100F::N/64
!
```

```
!
ipv6 unicast-routing
ipv6 cef
!
interface FastEthernet0/0
ipv6 address 2001:4F8:FFFE:100F::/64 eui-64
ipv6 address 2001:4F8:FFFE:100F::<N+1>/64
!
```

Numbering a PC

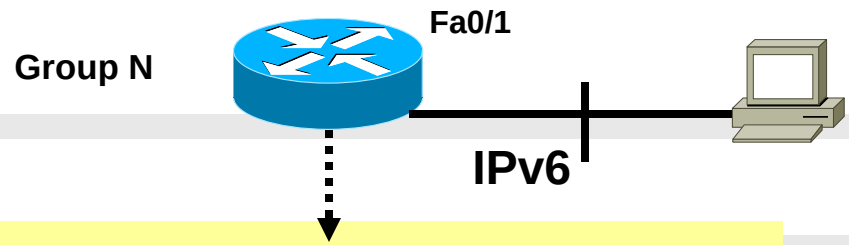


```
!
interface FastEthernet0/1
ipv6 address 2001:4F8:FFFE:100<N>::1/64
!
```

Verifying IPv6

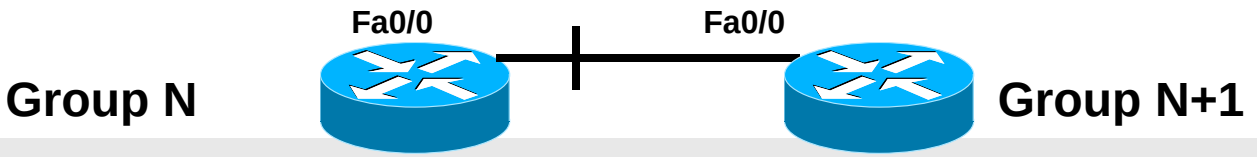

- **Verify**

show ipv6 interface

ping the remote ipv6 interface

show ipv6 route

debug ipv6 nd

Configuring IPv6 static routes

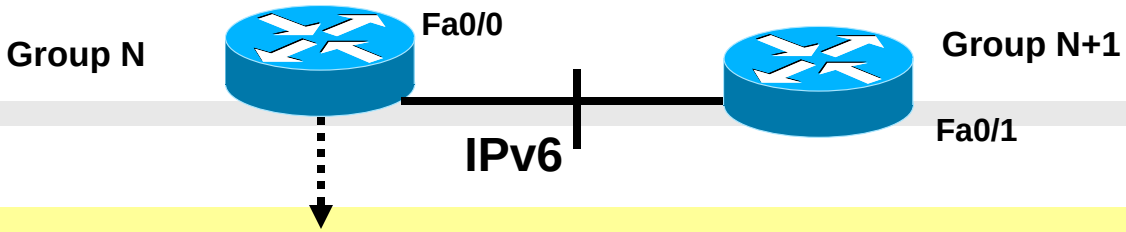


```
!
ipv6 route 2001:4F8:FFFE:100<N+1>::/64  2001:4F8:FFFE:100F::<N+1>
!
ipv6 route 2001:4F8:FFFE:100<N+2>::/64  2001:4F8:FFFE:100F::<N+2>
```

Verifying IPv6 static routes

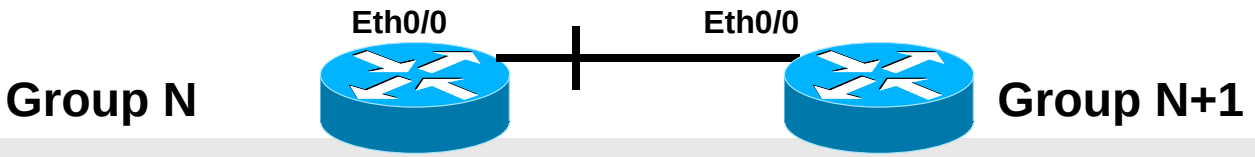

- **Verify**

ping remote fa0/1 interfaces

RIPng

- RFC 2080 describes RIPng
- Same as IPv4
  - Distance-vector, 15 hop Radius, split-horizon, poison reverse, etc..
  - Based on RIPv2
- Updated Features for IPv6
  - Uses IPv6 for transport
  - IPv6 prefix, next-hop IPv6 address
  - Uses the multicast group FF02::9 for RIP updates
  - Updates are sent on UDP port 521

Enhanced Routing Protocol Support
RIPng Configuration and Display

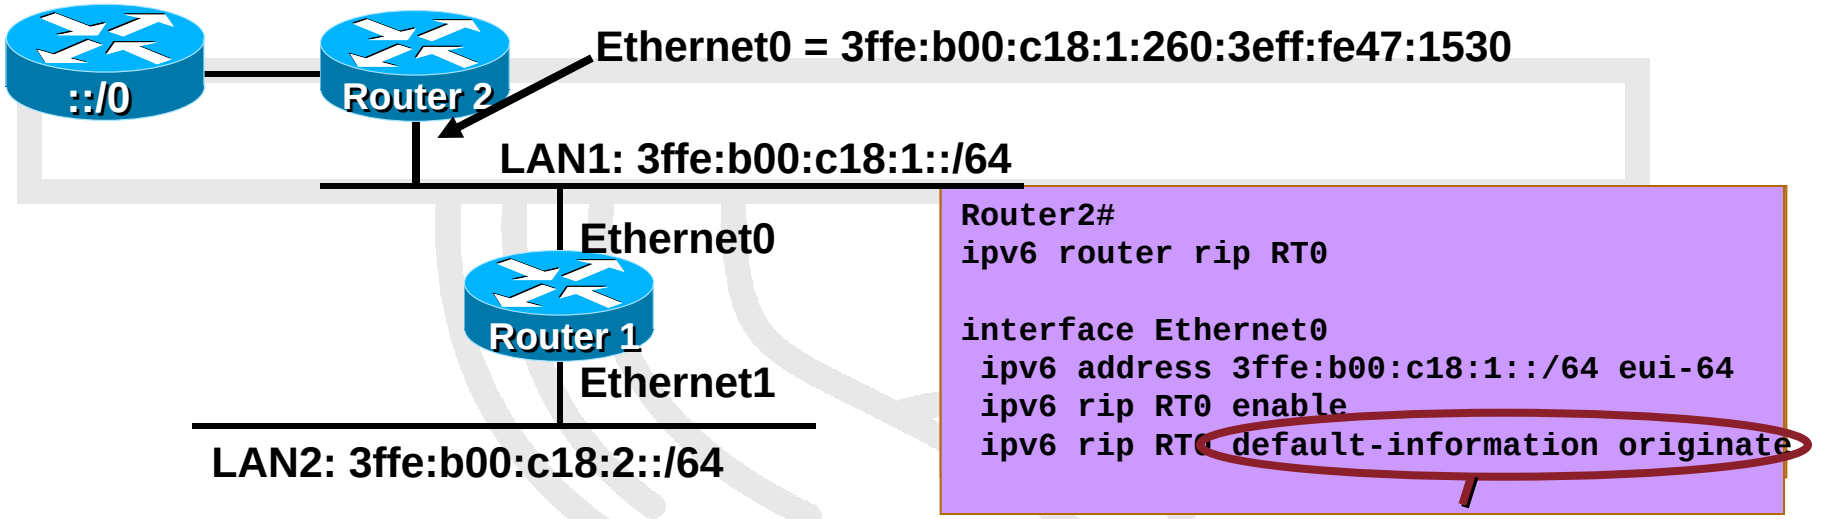


```
Router2#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
 ipv6 rip RT0 default-information originate
```

```
Router1#
ipv6 router rip RT0

interface Ethernet0
 ipv6 address 3ffe:b00:c18:1::/64 eui-64
 ipv6 rip RT0 enable
Interface Ethernet1
 ipv6 address 3ffe:b00:c18:2::/64 eui-64
 ipv6 rip RT0 enable
```

```
Router2# debug ipv6 rip
RIPng: Sending multicast update on Ethernet0 for RT0
       src=FE80::260:3eff:fe47:1530
       dst=FF02::9 (Ethernet0)
       sport=521, dport=521, length=32
       command=2, version=1, mbz=0, #rte=1
       tag=0, metric=1, prefix=::/0
```

Multicast all Rip-Routers

Link-local src address

Enabling RIPng

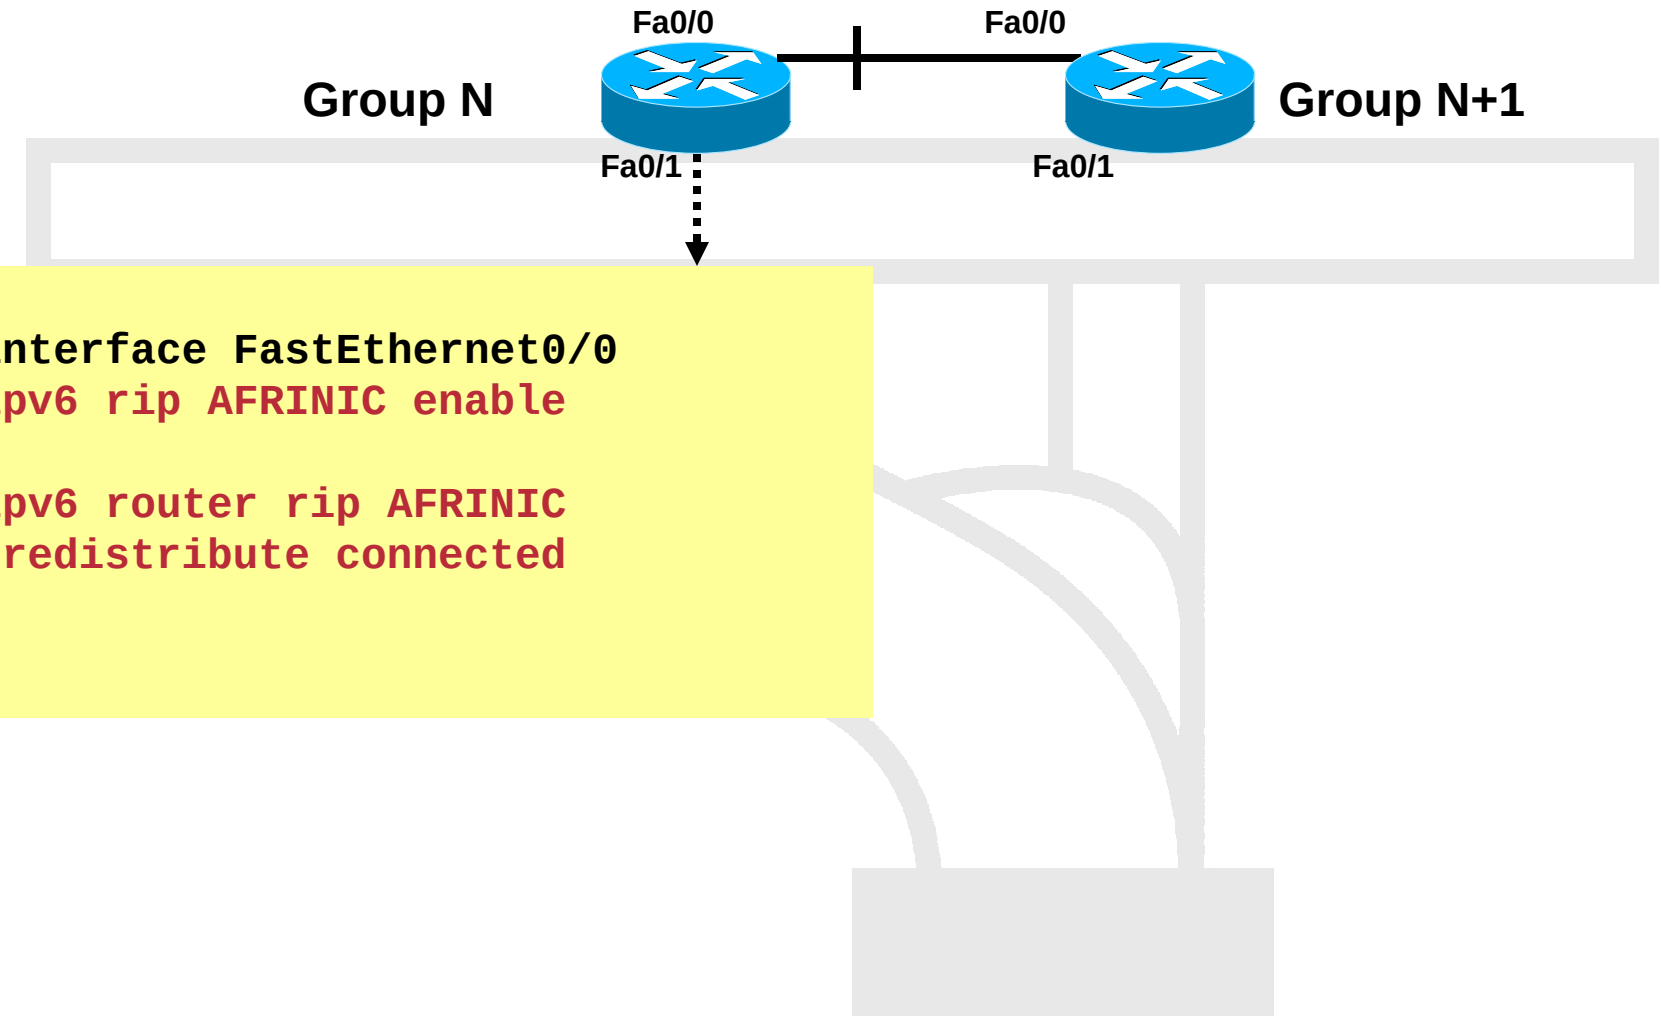


```
!
interface FastEthernet0/0
ipv6 rip AFRINIC enable
!
ipv6 router rip AFRINIC
 redistribute connected
!
```

Enable RIPng - verifying

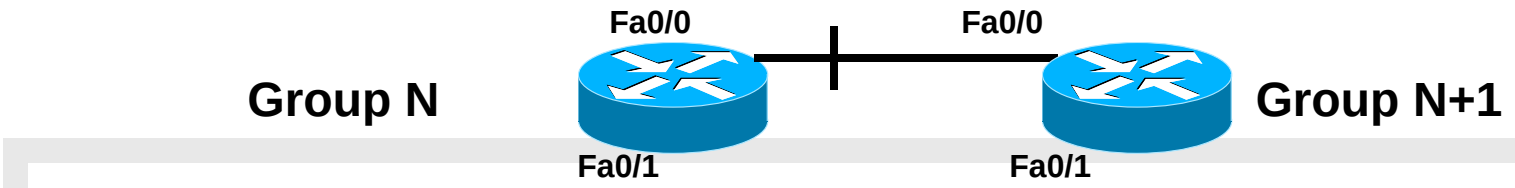

- **Verify**
  - **show ipv6 route rip**
  - **show ipv6 rip**
  - **ping <remote interface>**
  - **debug ipv6 rip**

- **Suggestions**
  - **filter the updates**
  - **announce more prefixes (create lookback interfaces)**

OSPFv2

- April 1998 was the most recent revision (RFC 2328)
- OSPF uses a 2-level hierarchical model
- SPF calculation is performed independently for each area
- Typically faster convergence than DVRPs
- Relatively low, steady state bandwidth requirements

OSPFv3 overview

- OSPF for IPv6
- Based on OSPFv2, with enhancements
- Distributes IPv6 prefixes
- Runs directly over IPv6
- Ships-in-the-night with OSPFv2

OSPFv3 / OSPFv2 Similarities

- **Basic packet types**

  Hello, DBD, LSR, LSU, LSA

- **Mechanisms for neighbor discovery and adjacency formation**
- **Interface types**

  P2P, P2MP, Broadcast, NBMA, Virtual

- **LSA flooding and aging**
- **Nearly identical LSA types**

OSPFv3 / OSPFv2 Differences

- OSPFv3 runs over a link, rather than a subnet
- Multiple instances per link
- OSPFv2 topology not IPv6-specific
  - Router ID
  - Link ID
- Standard authentication mechanisms
- Uses link-local addresses
- Generalized flooding scope
- Two new LSA types

Configuring OSPFv3 in Cisco IOS® Software

- **Similar to OSPFv2**
  - **Prefixing existing Interface and Exec mode commands with "ipv6"**
- **Interfaces configured directly**
  - **Replaces `network` command**
- **"Native" IPv6 router mode**
  - **Not a sub-mode of `router ospf`**

Configuration Modes in OSPFv3

- **Entering router mode**

  [no] ipv6 router ospf <process ID>

- **Entering interface mode**

  [no] ipv6 ospf <process ID> area <area ID>

- **Exec mode**

  [no] show ipv6 ospf [<process ID>]

  clear ipv6 ospf [<process ID>]

Cisco IOS OSPFv3 Specific Attributes

- **Configuring area range**

  [no] area <area ID> range <prefix>/<prefix length>

- **Showing new LSA**

  show ipv6 ospf [<process ID>] database link

  show ipv6 ospf [<process ID>] database prefix

OSPFv3 configuration example

```
Router1#
 interface Ethernet0
   ipv6 address 2001:1:1:1::1/64
   ipv6 ospf 1 area 0

  interface Ethernet1
    ipv6 address 2001:2:2:2::2/64
    ipv6 ospf 1 area 1

  ipv6 router ospf 1
   router-id 1.1.1.1
   area 1 range 2001:2:2::/48
```

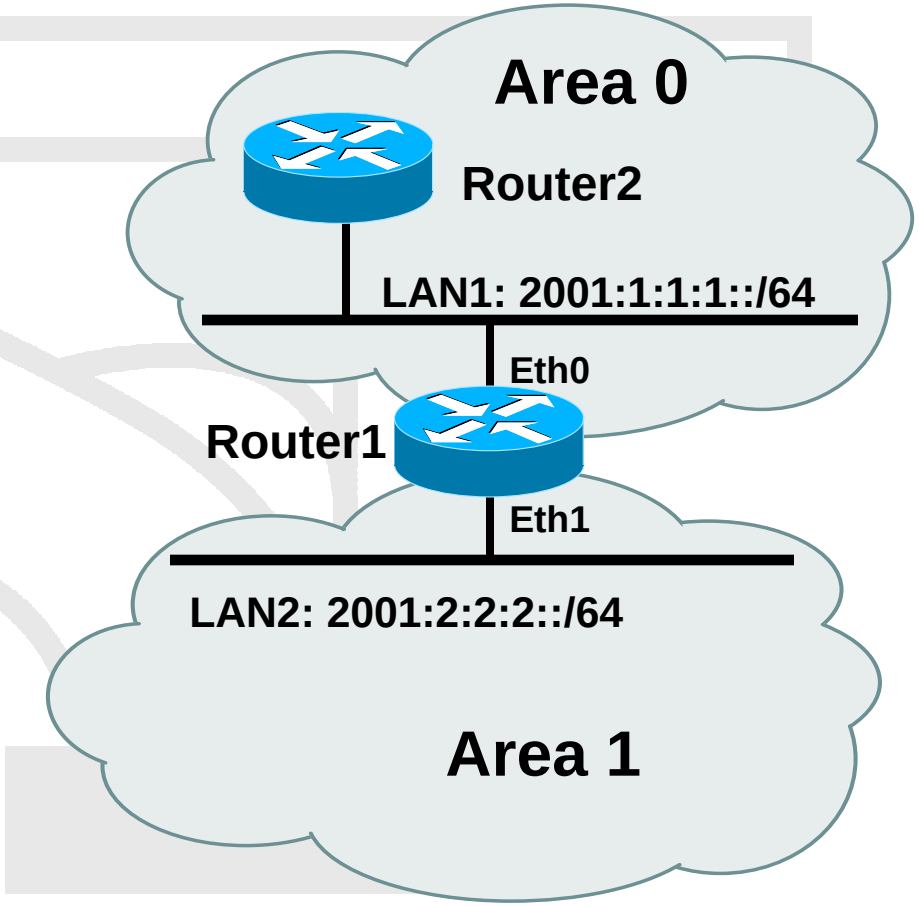

Cisco IOS OSPFv3 Display

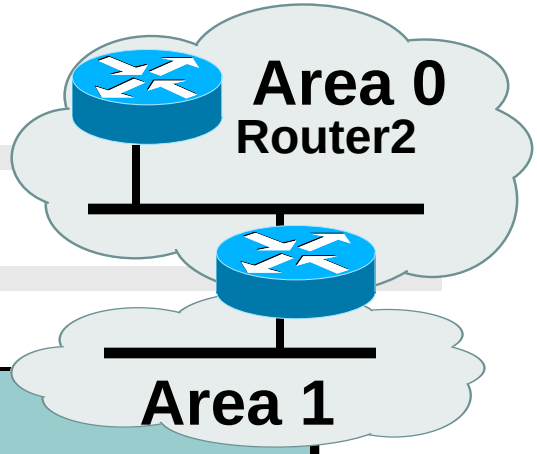


```
Router 2# show ipv6 route ospf
IPv6 Routing Table - 9 entries
Codes: C - Connected, L - Local, S - Static, R - RIP, B - BGP
       U - Per-user Static route
       I1 - ISIS L1, I2 - ISIS L2, IA - ISIS interarea
       O - OSPF intra, OI - OSPF inter, OE1 - OSPF ext 1, OE2 - OSPF ext 2
O   2001:1:1:2::1/128 [110/1]
     via FE80::205:5FFF:FEAF:2C38, Ethernet0
OI  2001:2:2::/48 [110/2]
     via FE80::205:5FFF:FEAF:2C38, Ethernet0
```

Enable OSPFv3

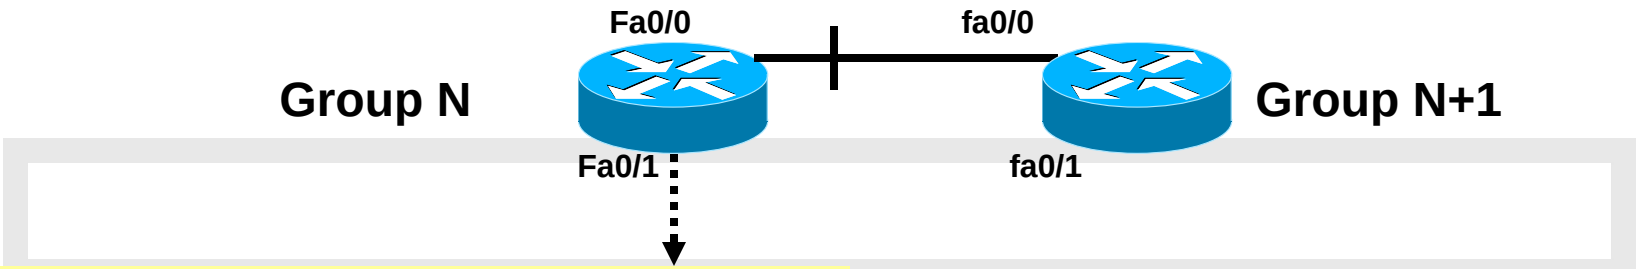


```
interface FastEthernet0/0
ipv6 ospf 1 area 0
!
ipv6 ospf 1
router-id N.N.N.N
redistribute connected
!
```

Verifying OSPFv3

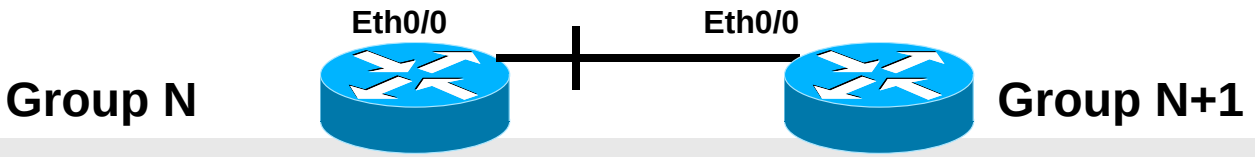

- **Verify**

show ipv6 ospf

show ipv6 ospf interface

show ipv6 ospf neighbor

Show IPv6 ospf database

IPv6 over IPv4 Tunnels

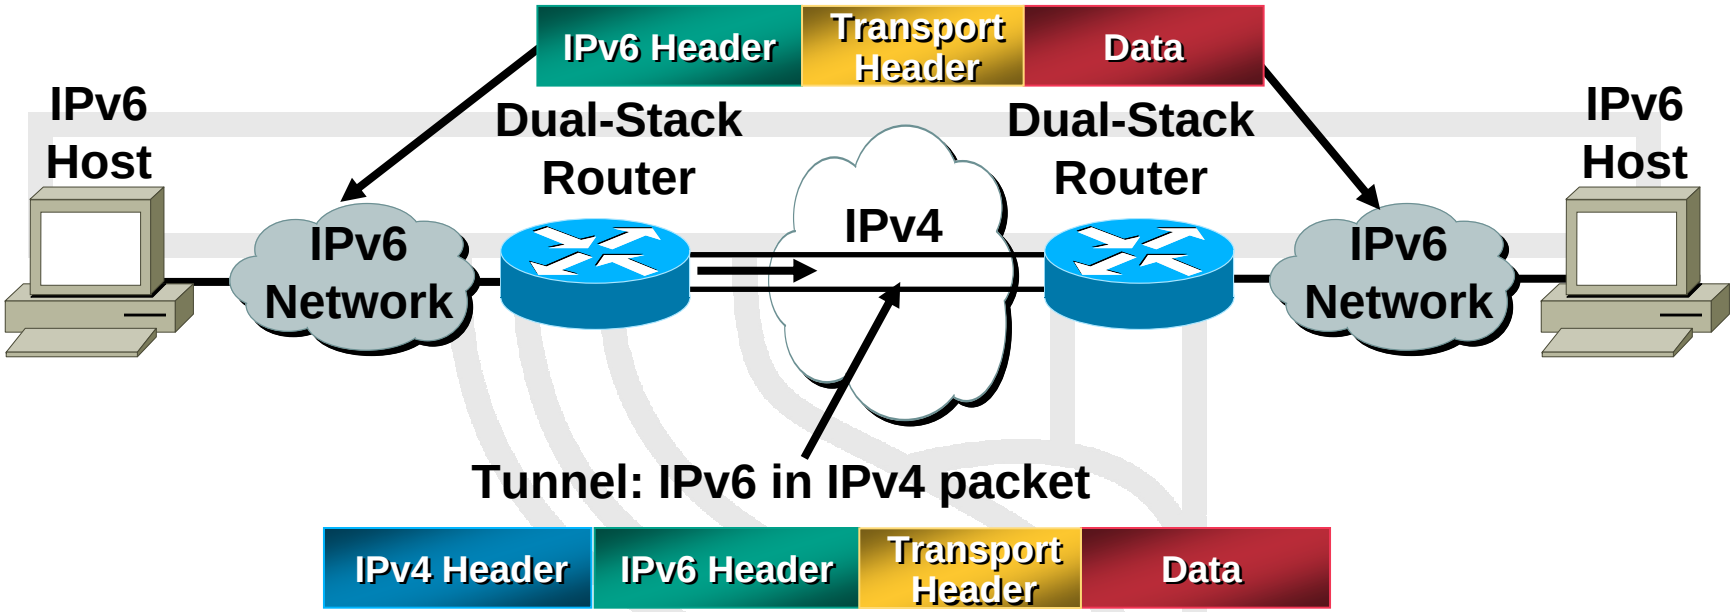

- Tunneling is encapsulating the IPv6 packet in the IPv4 packet
- Tunneling can be used by routers and hosts
- Many techniques are available to establish a tunnel
  - Configured: Manual Tunnel (RFC 2893), GRE (RFC 2473)
  - Automatic: 6to4 (RFC 3056), ISATAP

Manually Configured Tunnel (RFC 2893)

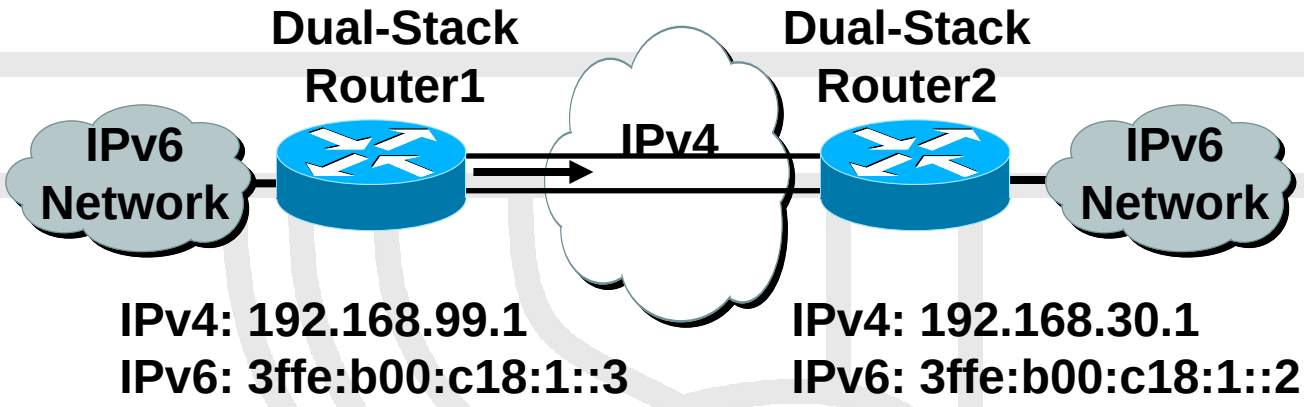


```
router1#

interface Tunnel0
 ipv6 address 3ffe:b00:c18:1::3/64
 tunnel source 192.168.99.1
 tunnel destination 192.168.30.1
 tunnel mode ipv6ip
```

```
router2#

interface Tunnel0
 ipv6 address 3ffe:b00:c18:1::2/64
 tunnel source 192.168.30.1
 tunnel destination 192.168.99.1
 tunnel mode ipv6ip
```

- **Manually Configured tunnels require:**
  - **Dual stack end points**
  - **Both IPv4 and IPv6 addresses configured at each end**

6to4 Tunnel (RFC 3056)

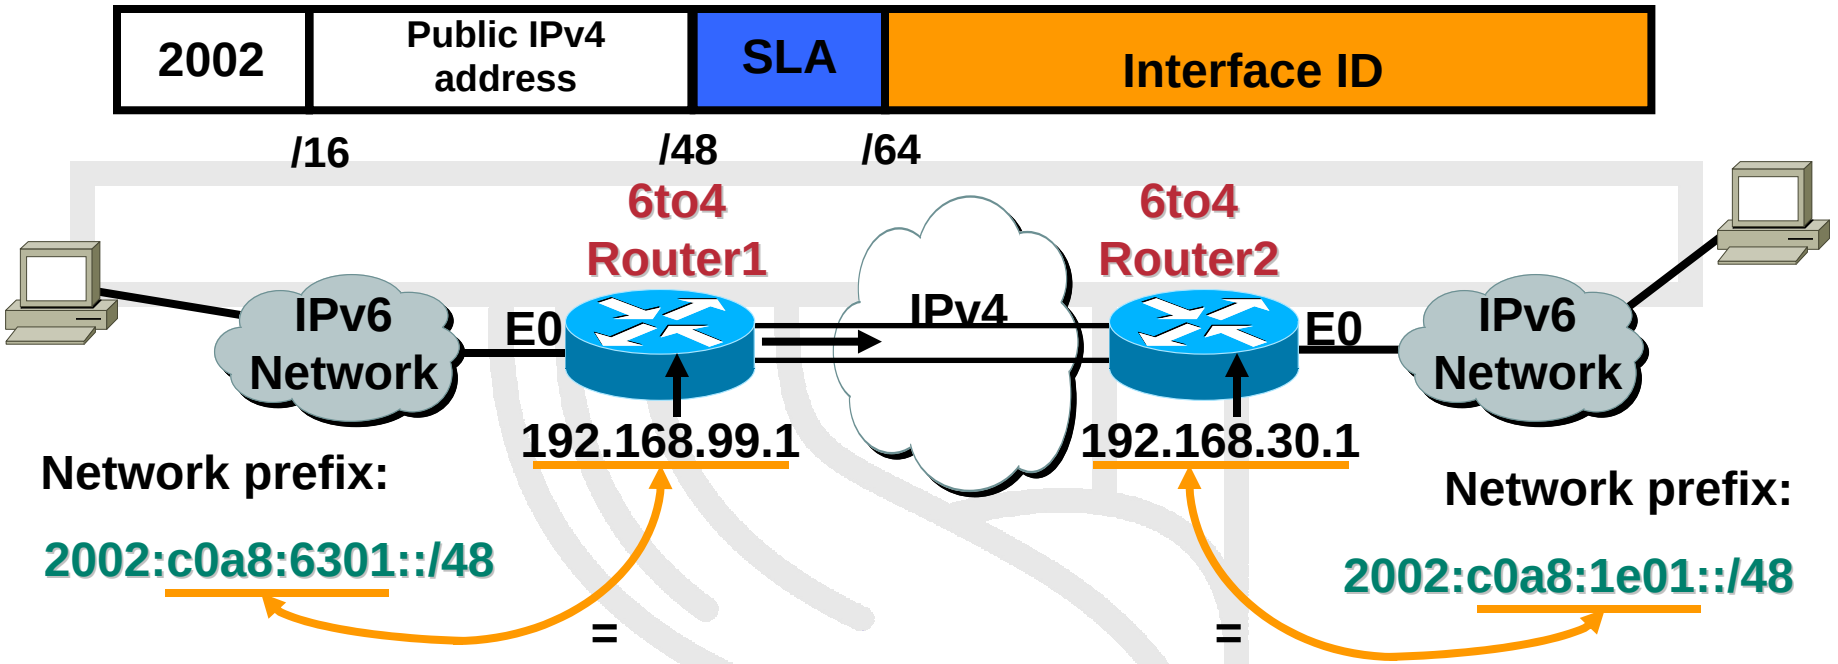

- **6to4 Tunnel:**
  - Is an automatic tunnel method
  - Gives a prefix to the attached IPv6 network
  - 2002::/16 assigned to 6to4
  - Requires one global IPv4 address on each Ingress/Egress site

```
router2#
interface Loopback0
 ip address 192.168.30.1 255.255.255.0
 ipv6 address 2002:c0a8:1e01:1::/64 eui-64
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Loopback0
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
```

6to4 Relay

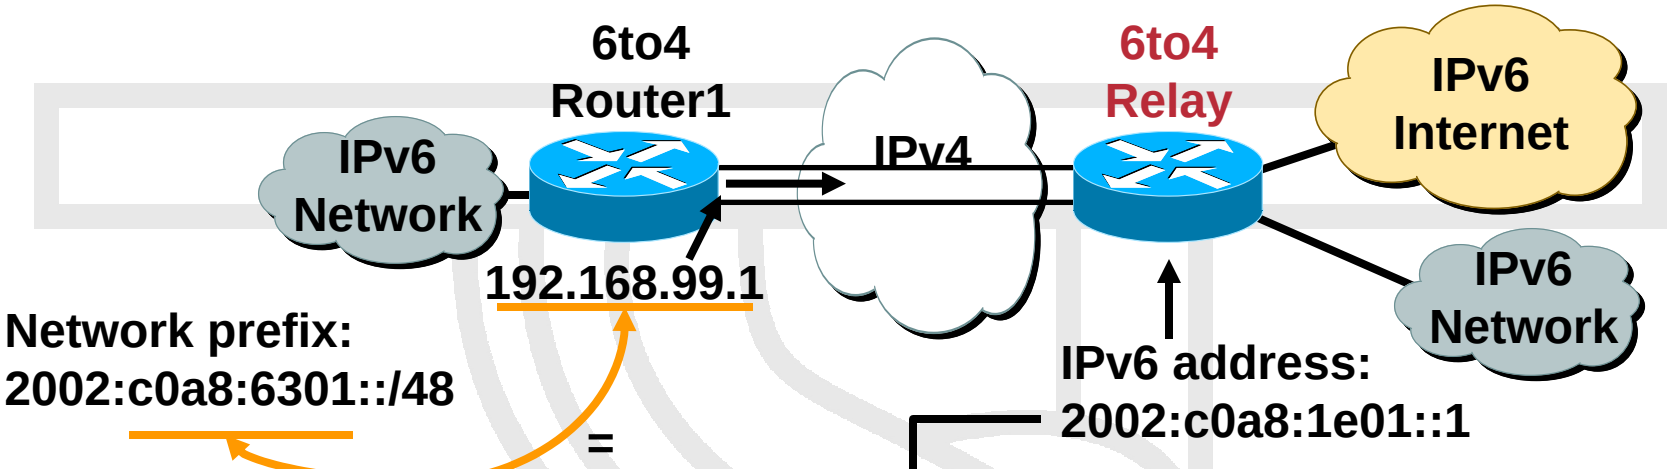


```
router1#
interface Loopback0
 ip address 192.168.99.1 255.255.255.0
 ipv6 address 2002:c0a8:6301:1::/64 eui-64
interface Tunnel0
 no ip address
 ipv6 unnumbered Ethernet0
 tunnel source Loopback0
 tunnel mode ipv6ip 6to4

ipv6 route 2002::/16 Tunnel0
ipv6 route ::/0 2002:c0a8:1e01::1
```

- **6to4 relay:**
  - Is a gateway to the rest of the IPv6 Internet
  - Default router
  - Anycast address (RFC 3068) for multiple 6to4 Relay

Setup an IPv6/IPv4 tunnel

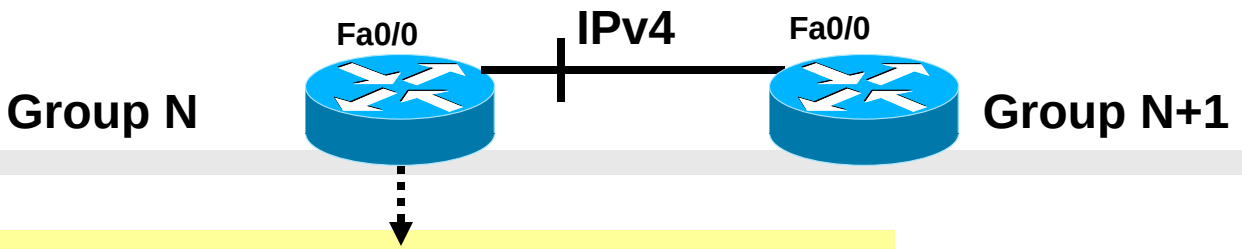


```
!
interface Tunnel0
ipv6 address 2001:4F8:FFFE:10F<N>::<N>/64
tunnel source 196.200.209.N
tunnel destination 196.200.209.<N+1>
ipv6 rip AFRINIC enable
tunnel mode ipv6ip
!
interface FastEthernet0/0
ip address 196.200.209.N 255.255.255.0
<remove IPV6 from the interface>
!
ipv6 router rip AFRINIC
 redistribute connected
```

Setup an IPv6/IPv4 tunnel - verifying

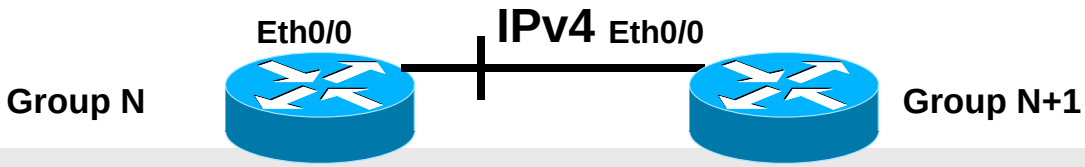

- **Verify**
  - **show ipv6 interface**
  - **show ipv6 rip**
  - **ping <remote loopback interface>**
- **Suggestions**
  - **Replace RIPng by static routing in the tunnel**
  - **Monitor Router Advertisement in the tunnel**

Campus Deployment: ISATAP

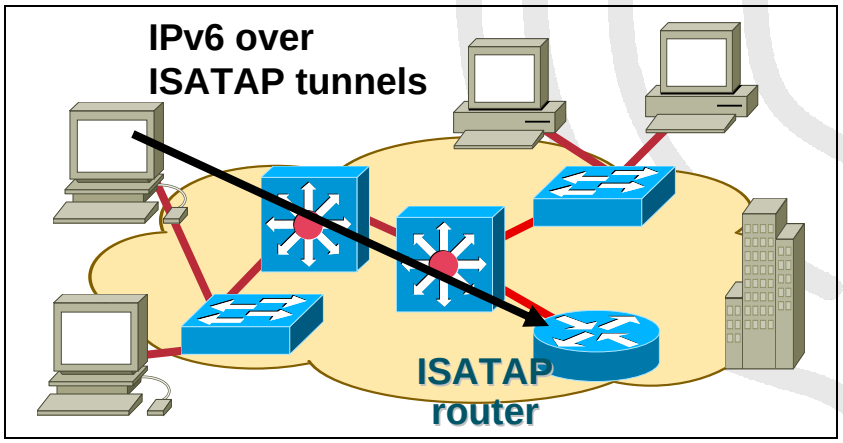

Supported on Microsoft Windows XP Pro SP1, .NET

- **An upgrade of Layer 3 infrastructure to implement Dual Stack can't be performed but IPv6 is needed**
  - **Cost**
  - **Hardware/software dependencies**
  - **Sparse IPv6 population on campus**
- **ISATAP router provides an IPv6 prefix to the IPv6 hosts**
- **Direct connection via tunnels between ISATAP hosts**
- **ISATAP router used to go outside from the LAN**

ISATAP Details

Use IANA‘s OUI 00-00-5E and encode IPv4 address as part of EUI-64

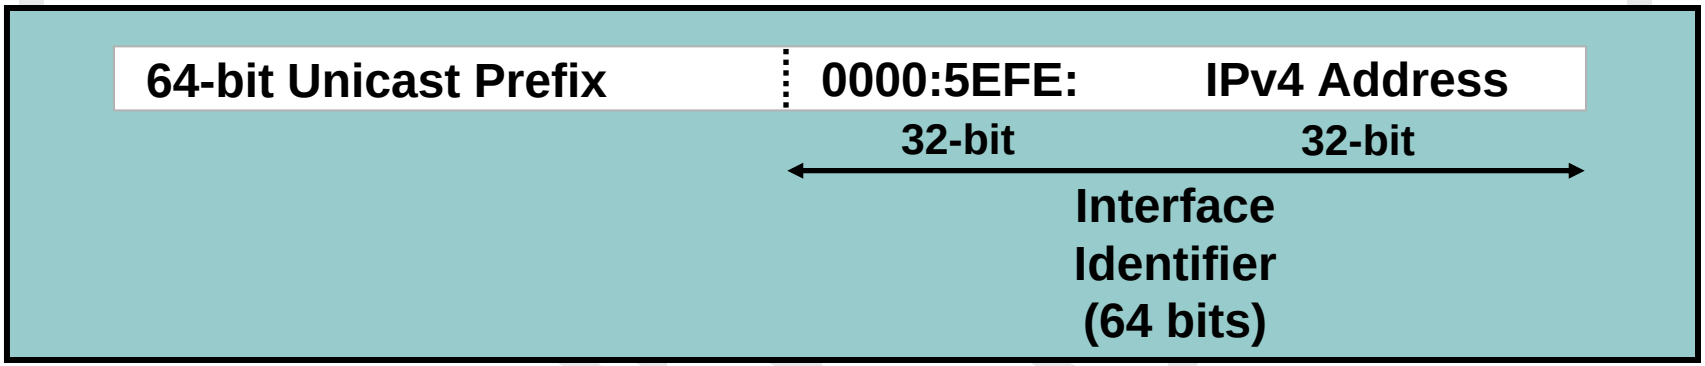

- **Automatic discovery of ISATAP routers**

  DNS "isatap.domainname" A record lookup

  Automatic deprecation when end system receives native IPv6 router advertisements

draft-ietf-ngtrans-isatap-11

draft-ietf-ngtrans-isatap-scenario-01

ISATAP Router

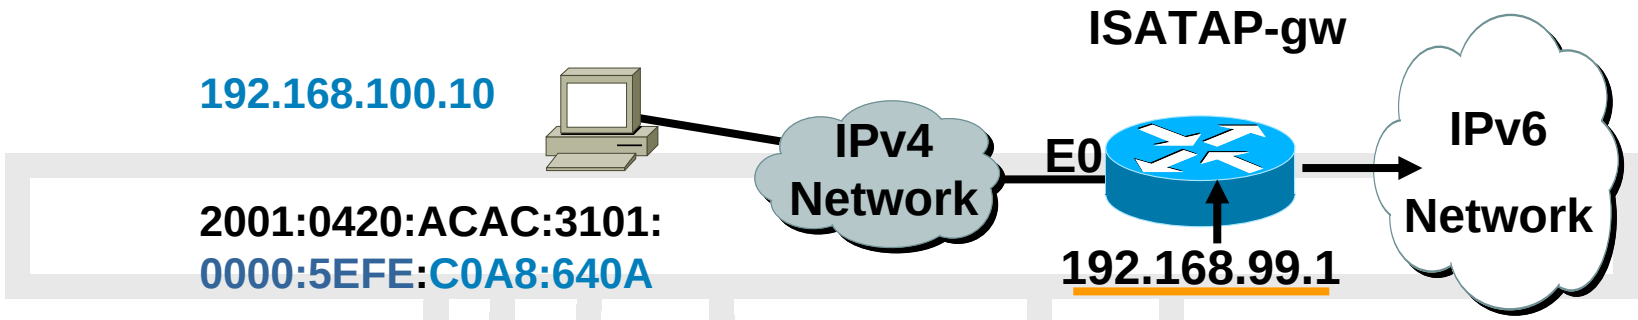


```
ISATAP-gw#
!
interface Ethernet0
 ip address 192.168.99.1 255.255.255.0
!
interface Tunnel0
 ipv6 address 2001:0420:ACAC:3101::/64 eui-64
 no ipv6 nd suppress-ra
 tunnel source Ethernet0
 tunnel mode ipv6ip isatap
!
```

- Cisco IOS 12.2(15)T, 12.2(14)S
- Supported in Windows XP Pro SP1
- The tunnel source command must point to an interface with an IPv4 address configured.
- Configure the ISATAP IPv6 address, and prefixes to be advertised just as you would with a native IPv6 interface.
- The IPv6 address has to be configured as an EUI-64 address since the last 32 bits in the interface identifier is used as the IPv4 destination address.

Windows XP Pro - ISATAP

- **Cisco Router**

  IP (interface Fast0/0) : 10.151.1.1

  ISATAP Default IPv6 : FE80::5EFE:A97:101

- **Windows XP Pro**

  By default ISATAP Adress = FE80::0000:5EFE:w.x.y.z

  Where w.x.y.z is the IPv4 address

  IP : 10.151.4.194

  Default ISATAP IPv6 : FE80::5EFE:10.151.4.194

  Adding global address - ipv6 adu 2/3000:1::10.151.1.194

  Adding route - ipv6 rtu 3000:1::/64 2

Cisco IOS

```
!
interface Tunnel0
 no ip address
 no ip redirects
 ipv6 address 3000:1::1/64
 tunnel source FastEthernet0/0
 tunnel mode ipv6ip isatap
!
interface FastEthernet0/0
 ip address 10.151.1.1 255.255.255.0
 speed 100
 full-duplex
 ipv6 address 2006:1::1/64
 ipv6 nd prefix 2006:1::/64 10000 8000
 ipv6 ospf 1 area 0
!
```

Cisco IOS

```
r1#sh ipv6 int tu0
Tunnel0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::5EFE:A97:101
  Global unicast address(es):
    3000:1::1, subnet is 3000:1::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:0
    FF02::1:FF00:1
    FF02::1:FF97:101
  MTU is 1480 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is not supported
  ND reachable time is 30000 milliseconds
  Subnet-router Anycast is configured
  Hosts use stateless autoconfig for addresses.
r1#
```

Windows XP

```
C:\>ipconfig
Configuration IP de Windows

Carte Ethernet Connexion au réseau local:

        Suffixe DNS propre à la connexion : cisco.com
        Adresse IP. . . . . . . . . . . . : 10.151.4.194
        Masque de sous-réseau . . . . . . : 255.255.255.0
        Passerelle par défaut . . . . . . : 10.151.4.1

C:\>
C:\>ipv6 if 2
Interface 2 :  Pseudo-interface de tunnels automatiques
  n'utilise pas la découverte de voisin
  n'utilise pas la découverte de routeur
  adresse de couche de liaison : 0.0.0.0
  Adresse IPv4 imbriquée EUI-64 : 0.0.0.0
    preferred link-local fe80::5efe:10.151.4.194, vie infinite
  MTU de liaison 1280 (MTU de liaison réelle 65515)
  limite de sauts actuelle 128
  durée d'attente pour la communication 32000ms (base 30000ms)
  intervalle de retransmission 1000ms
  DAD transmet 0

C:\>
```

Windows XP Pro

Ping ISATAP address of R1 10.151.1.1

```
C:\>ping6 fe80::5efe:A97:101%2

Envoi d'une requête 'Ping' fe80::5efe:10.151.1.1%2
à partir de fe80::5efe:10.151.4.194%2 avec 32 octets de données :

Réponse de fe80::5efe:10.151.1.1%2 : octets = 32 temps=32 ms
Réponse de fe80::5efe:10.151.1.1%2 : octets = 32 temps=31 ms
Réponse de fe80::5efe:10.151.1.1%2 : octets = 32 temps=30 ms
Réponse de fe80::5efe:10.151.1.1%2 : octets = 32 temps=31 ms

Statistiques de Ping pour fe80::5efe:10.151.1.1%2 :
    Paquets : envoyés = 4, reçus = 4, perdus = 0 (0% de perte),
Durée approximative des boucles en millisecondes :
    Minimum = 30ms, maximum = 32ms, moyenne = 31ms

C:\>
```

Enable MP-BGP on the tunnel

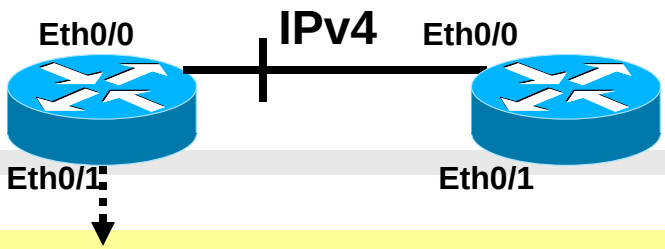


```
router bgp N
 no synchronization
 no bgp default ipv4-unicast
 bgp log-neighbor-changes
neighbor 2001:4F8:FFFE:10F<N+1> remote-as <N+1>
 !
 address-family ipv6
 neighbor 2001:4F8:FFFE:10F<N+1> activate
 redistribute connected
 exit-address-family
!
```

Enable MP-BGP on the tunnel- verifying

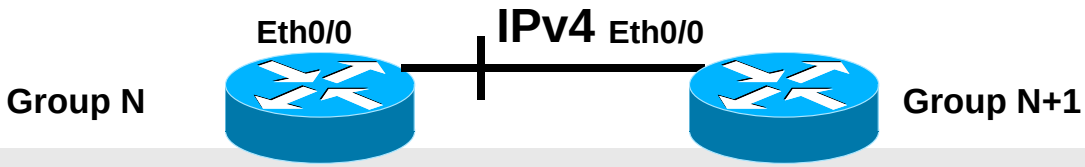

- **Verify**

  show bgp neighbor

  show bgp ipv6

  ping <remote loopback interface>

- **Suggestions**

  add policies and filters

MP-BGP peer over IPv6- verifying

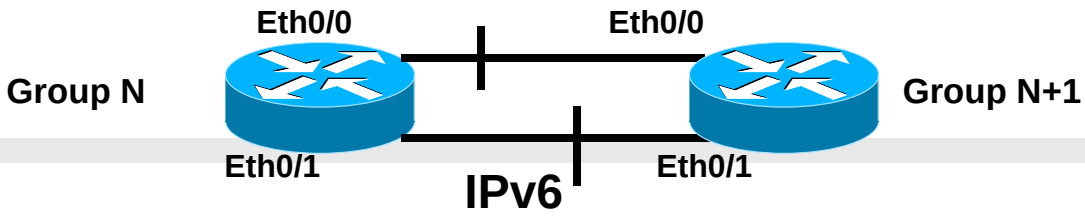

- **Verify**
  - **show bgp neighbor**
  - **Show bgp ipv6**
  - **Ping <remote loopback interface>**
- **Suggestions**
  - **Remove the “update-source” and see what happens**
  - **Add policies and filters, play with attributes**

References

- **IPv6 documentation on CCO**

http://www.cisco.com/univercd/cc/td/doc/product/software/ios124/124cg/hipv6_c/index.htm

- **IPv6 page on CCO: www.cisco.com/go/ipv6**