IPv6

A Comprehensive Tutorial

Author	:	Tony Hill
Date	:	20th December 2013
Version	:	v1-0
Cisco IOS	:	c2600-advipservicesk9-mz124-23.bin
Windows	:	Vista (32-bit) SP2

INDEX

1 Introduction

This paper is a comprehensive IPv6 tutorial. It describes most of the fundamental components of IPv6 and the mechanisms that IPv6 uses for stateless configuration, neighbour discovery and duplicate address detection. It contains detailed explanations of the configuration and operation of the IPv6 protocol. A small lab is used to obtain packet captures to support explanations of protocol functionality. Device configuration examples are provided throughout.

2 Lab Set-Up

2.1 Lab Topology

Figure 2-1 below shows the very simple lab set-up. ISIS is used as the routing protocol primarily because it is relatively simple to configure, it supports IPv4 / IPv6 multi-topology and ISPs use it widely in their networks. Router R1's and R2's Network Entity Titles (NETs) are shown in blue. Each NET contains the embedded IPv4 loopback address of each router.

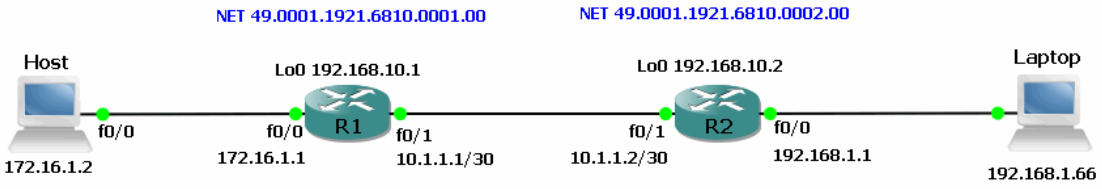

Figure 2-1: Lab IPv4 Topology

2.1.1 Lab IPv4 Configuration

Host (left) has IPv4 address 172.16.1.2 with a default gateway of 172.16.1.1 (R1). Windows Laptop (right) has IPv4 address 192.168.1.66 with a static route towards 172.16.1.0/24 via 192.168.1.1 (R2).

```
C:\windows\system32>route add 172.16.1.0 mask 255.255.255.0 192.168.1.1
OK!
```

Laptop is able to ping Host with IPv4 and vice-versa.

```
C:\windows\system32>ping 172.16.1.2
Pinging 172.16.1.2 with 32 bytes of data:
Reply from 172.16.1.2: bytes=32 time=106ms TTL=253
Reply from 172.16.1.2: bytes=32 time=77ms TTL=253
Reply from 172.16.1.2: bytes=32 time=78ms TTL=253
Reply from 172.16.1.2: bytes=32 time=93ms TTL=253
Ping statistics for 172.16.1.2:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 77ms, Maximum = 106ms, Average = 88ms

Host#ping 192.168.1.66
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 192.168.1.66, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 60/86/156 ms
```

2.1.2 Lab IPv6 Configuration

Figure 2-2 below shows the MAC and IPv6 addresses used in the lab. Router R1 and R2 LAN interface fa0/0 MAC addresses are set manually to AA:AA:11:11:11:11 and AA:A:22:22:22:22 respectively.

R1's IPv6 global unicast LAN address is set manually to 2001:0A0A::1/64 and R2's IPv6 global unicast LAN address to 2001:0B0B::1/64. The link between the two routers uses IPv6 addresses in the unique local range FD00::/8, which is the equivalent of an IPv4 RFC 1918 private address.

Both Host and Laptop use stateless auto-configuration to obtain their IPv6 global unicast prefixes from R1 and R2.

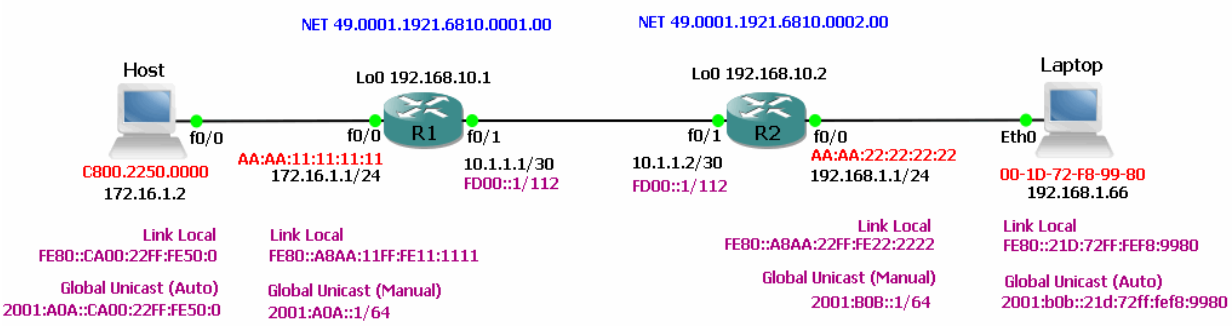

Figure 2-2: Lab IPv6 Topology

Laptop is able to ping Host's IPv6 global unicast address and vice-versa.

```
C:\windows\system32>ping 2001:A0A::CA00:22FF:FE50:0
Pinging 2001:a0a::ca00:22ff:fe50:0 from 2001:b0b::21d:72ff:fef8:9980 with 32 bytes of data:
Reply from 2001:a0a::ca00:22ff:fe50:0: time=109ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=78ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=78ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=62ms

Ping statistics for 2001:a0a::ca00:22ff:fe50:0:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 62ms, Maximum = 109ms, Average = 81ms

Host#ping 2001:b0b::21d:72ff:fef8:9980
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 2001:B0B::21D:72FF:FEF8:9980, timeout is 2 seconds:
!!!!!
Success rate is 100 percent (5/5), round-trip min/avg/max = 52/80/112 ms
```

Laptop's IPv4 and IPv6 addresses are shown below using the "ipconfig/all" command. Note that the Ethernet interface has both a global unicast and link local address.

Neither of these addresses was configured manually. Laptop used the interface MAC address to generate its link local EUI-64 (End User Identifier) address. It used auto-configuration to obtain the /64 network portion of the global unicast address from router R2 and auto-generated the global unicast EUI-64 address using the MAC address. It obtained the IPv6 default gateway (R2's link local address) during router solicitation.

```
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  . : home
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-1D-72-F8-99-80
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:b0b::21d:72ff:fef8:9980(Preferred)
   Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.66(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : 19 December 2013 14:13:56
   Lease Expires . . . . . . . . . . : 20 December 2013 15:37:33
   Default Gateway . . . . . . . . . : fe80::a8aa:22ff:fe22:2222%10
                                       192.168.1.254
```

```
   DHCP Server . . . . . . . . . . . : 192.168.1.254
   DNS Servers . . . . . . . . . . . : 192.168.1.254
                                       192.168.1.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

The link local and default gateway addresses both use the notation fe80**::**<address>**%10**. The %10 is the Windows internal interface index that binds the link local addresses to this specific interface.

```
C:\windows\system32>netsh interface ipv6 show interfaces
Idx  Met   MTU      State        Name
---  ---  -----     -----------  -------------------
  1   50 4294967295 connected    Loopback Pseudo-Interface 1
 11   25   1500     disconnected Wireless Network Connection
 .
 .
 10   20   1500     connected    Local Area Connection
 .
 .
 49   20   1500     connected    VirtualBox Host-Only Network
```

The "netsh interface ipv6 show route" command confirms that Laptop has installed the correct default route (**::**/0) via R2's link local address.

```
C:\windows\system32>netsh interface ipv6 show route
Publish  Type      Met  Prefix                    Idx  Gateway/Interface Name
-------  --------  ---  ------------------------  ---  ------------------------
No       Manual    256  ::/0                       10  fe80::a8aa:22ff:fe22:2222
No       Manual    256  ::1/128                     1  Loopback Pseudo-Interface 1
No       Manual    8    2001:b0b::/64              10  Local Area Connection
No       Manual    8    2001:b0b::/120             10  Local Area Connection
No       Manual    256  2001:b0b::21d:72ff:fef8:9980/128 10  Local Area Connection
.
.
No       Manual    256  fe80::21d:72ff:fef8:9980/128   10  Local Area Connection
.
.
```

Host's IPv6 addresses and default gateway are shown below. Host also uses the interface MAC address to generate its EUI-64 link local address and auto-configuration to obtain from R1 the /64 network portion of the global unicast address in order to auto-generate its EUI-64 global unicast address.

```
Host#show ipv6 int
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::CA00:22FF:FE50:0
  Global unicast address(es):
    2001:A0A::CA00:22FF:FE50:0, subnet is 2001:A0A::/64 [PRE]
      valid lifetime 2591933 preferred lifetime 604733
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF50:0
  MTU is 1500 bytes
  ICMP error messages limited to one every 100 milliseconds
  ICMP redirects are enabled
  ND DAD is enabled, number of DAD attempts: 1
  ND reachable time is 30000 milliseconds
  Default router is FE80::A8AA:11FF:FE11:1111 on FastEthernet0/0
```

Table 2-1 below shows the configuration of routers R1 and R2 in the lab.

Router R1	**Router R2**
ipv6 unicast-routing interface Loopback0 ip address 192.168.10.1 255.255.255.255 ip router isis	ipv6 unicast-routing interface Loopback0 ip address 192.168.10.2 255.255.255.255 ip router isis !

interface FastEthernet0/0 mac-address aaaa.1111.1111 ip address 172.16.1.1 255.255.255.0 ip router isis duplex auto speed auto **ipv6 address 2001:A0A::1/64** **ipv6 router isis** interface FastEthernet0/1 ip address 10.1.1.1 255.255.255.252 ip router isis duplex auto speed auto **ipv6 address FD00::1/112** **ipv6 router isis** **router isis** **net 49.0001.1921.6810.0001.00** **is-type level-2-only** **metric-style wide** **address-family ipv6** **multi-topology** **exit-address-family**	interface FastEthernet0/0 mac-address aaaa.2222.2222 ip address 192.168.1.1 255.255.255.0 ip router isis duplex auto speed auto **ipv6 address 2001:B0B::1/64** **ipv6 router isis** interface FastEthernet1/0 ip address 10.1.1.2 255.255.255.252 ip router isis duplex auto speed auto **ipv6 address FD00::2/112** **ipv6 router isis** **router isis** **net 49.0001.1921.6810.0002.00** **is-type level-2-only** **metric-style wide** **address-family ipv6** **multi-topology** **exit-address-family**

Table 2-1: R1 and R2 Configurations

2.1.3 Lab Windows Configuration

By default, Windows uses a mechanism called temporary IPv6 addresses. In addition to the unicast EUI-64 address, Windows generates an additional, temporary unicast address with a randomly generated interface ID. The purpose of the temporary address is to provide anonymity and security. The temporary address allows a particular Windows device to avoid being identified by its interface ID.

For clarity, and to avoid confusion when examining packet captures in this tutorial, temporary addresses are disabled on Laptop. The command to determine whether temporary addresses are disabled is:

```
C:\windows\system32>netsh interface ipv6 show privacy
Querying active state...

Temporary Address Parameters
----------------------------------------------
Use Temporary Addresses             : enabled
Duplicate Address Detection Attempts: 5
Maximum Valid Lifetime              : 7d
Maximum Preferred Lifetime          : 1d
Regenerate Time                     : 5s
Maximum Random Time                 : 10m
Random Time                         : 0s
```

The following commands disable temporary addresses on Laptop across reboots using the "store=persistent" parameter:

```
C:\windows\system32>netsh interface ipv6 set privacy state=disabled store=persistent

C:\windows\system32> netsh interface ipv6 set global randomizeidentifiers=disabled
store=persistent
```

3 IPv6 Overview

IPv6 prefixes contain 128 bits providing a total of 2^{128} (3.4×10^{38}) possible addresses. However, the fact that there are 2^{128} bit combinations does not mean that there are this many IPv6 addresses available for allocation in a live, production network. The reasons for this are explained in the following sections.

3.1 IPv6 Header Information

Figure 3-1 below shows the 40-byte base header information of an IPv6 packet.

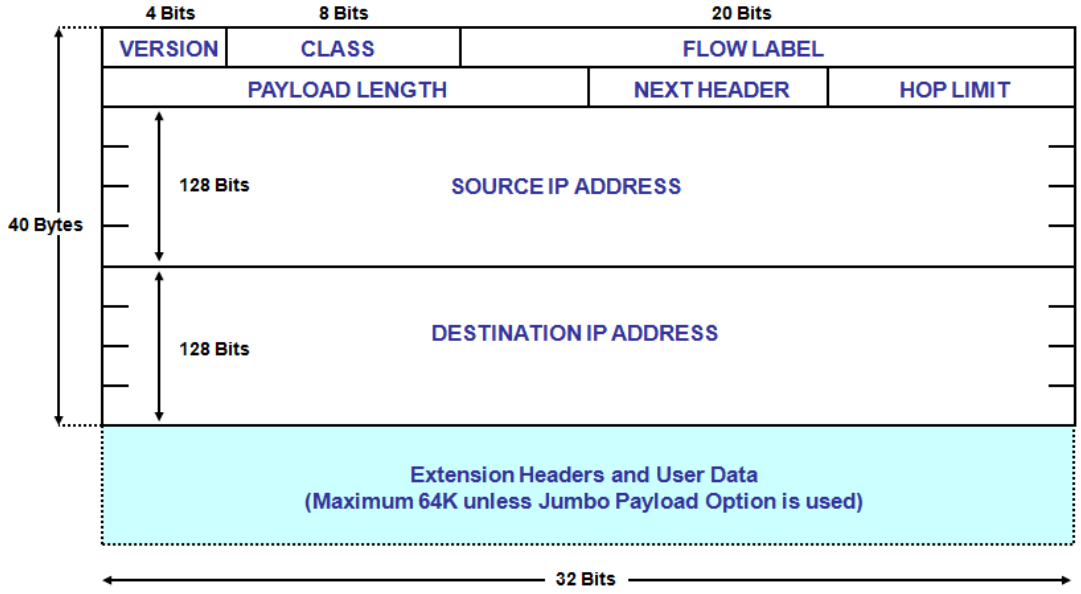

Figure 3-1: IPv6 Packet Header

The significance of each field is described below:

- **VERS** IP Version = 6 [0110 Binary]
- **CLASS** Replaces the IPv4 Precedence and ToS bits
- **FLOW** The Flow field identifies packets belonging to the same flow
- **PAYLOAD LENGTH** Total length of the Extension Headers and Data (16 bits = max 64K bytes)
- **NEXT HEADER** Points to the next (Extension) header after the IPv6 header
- **HOP LIMIT** Same function as the TTL field in IPv4
- **SOURCE IP** 128-bit (16 bytes) Source IPv6 Address
- **DESTINATION IP** 128-bit (16 bytes) Destination IPv6 Address

Figure 3-2 below shows how extension headers are appended to the base IPv6 header.

Figure 3-2: IPv6 Extension Headers

The extension header format has the following properties:

- Replaces the Options field in the IPv4 Header
- Means that the IPv6 “Base” Header is always fixed at 40 bytes
- Allows extensibility of the protocol if enhancements need to be made later
- Reduces processing - no extension headers mean no Options and faster switching

Figure 3-3 below shows how each extension header references the next extension header to provide the functionality required.

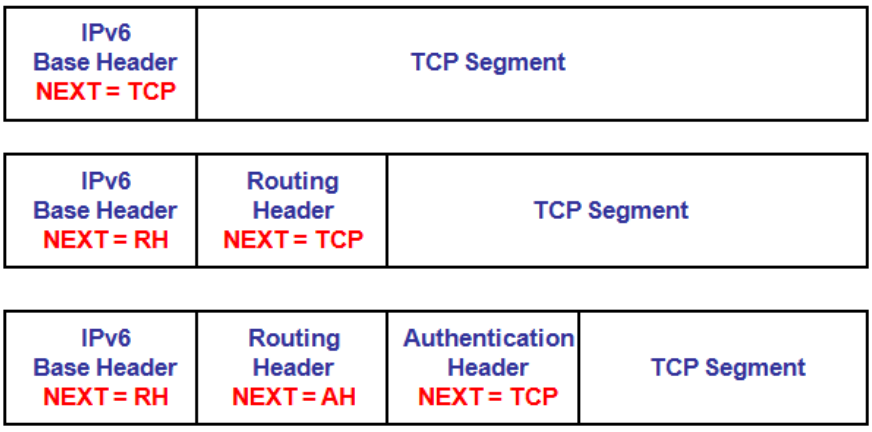

Figure 3-3: Extension Header Referencing

Please refer to the following URL for a complete list of the latest allocated next header values:

http://www.iana.org/assignments/protocol-numbers/protocol-numbers.xhtml

The next header mechanism allows for much easier incorporation of protocol functionality. For example, IPSEC has become an integral part of IPv6 and has its own next header references. Table 3-1 below shows the next header numbers for ESP and AH and their corresponding RFC numbers.

Header	Type	Description	RFC
50	ESP	Encapsulating Security Payload	RFC4303
51	AH	Authentication Header	RFC4302

Table 3-1: ESP & AH Next Header Numbers

One of the many fundamental differences between IPv6 and IPv4 is the way that fragmentation is handled. In IPv6, only hosts fragment packets - routers DO NOT!

If a host is not to fragment a packet, the minimum Maximum Transmission Unit (MTU) for an IPv6 packet (data + header) traversing a network is 1280 octets. In other words, the MTU of the end-to-end path (the PMTU) must not be less than 1280 octets at any given point. Hosts use PMTU discovery to ascertain the minimum, end-to-end MTU of the path but unlike IPv4, IPv6 does not have a “Do not Fragment” (DF) option. In IPv6, a host assumes that the end-to-end MTU is that of the attached link and starts to transmit packets. If a packet hits a router with a link that has a MTU less than 1280 octets, the router returns an ICMPv6 Packet Too Big (ICMPv6 Type 2) response to the host.

Figure 3-4 below illustrates how this is accomplished.

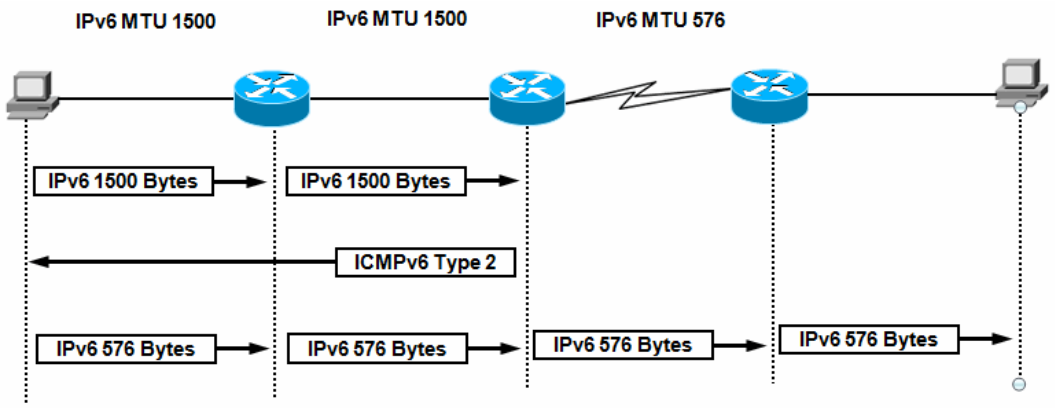

Figure 3-4: IPv6 PMTU Discovery

The IPv6 specification currently includes a fragmentation header option that hosts can use to fragment IPv6 packets at source so that the destination host can reassemble them. However, fragmentation and re-assembly is a resource intensive activity and is best avoided if possible. As at time of writing there is a move towards deprecating the fragmentation header in favour of hosts using only PMTU discovery.

Please refer to the following URL for the latest working group recommendation (dated July 2013) to deprecate the fragmentation header.

http://tools.ietf.org/pdf/draft-bonica-6man-frag-deprecate-02.pdf

3.2 IPv6 Address Types

Please refer to the URL www.ripe.net/ipv6-address-types for a detailed description of the various IPv6 address types. Table 3-2 below summarises the various IPv6 address ranges and their allocations.

Prefix	Description	IPv4 Equivalent
::/128	**Unspecified Address** - Used as a source address by a device that has not yet learnt its own address.	0.0.0.0
::1/128	**Loopback**	127.0.0.1
::FFFF/96 **e.g. ::FFFF:192.10.1.20**	**IPv4 Mapped** - IPv4 address embedded inside an IPv6 address. See RFC 3048.	N/A
FC00::/7 (FC00::/8 & FD::00/8) **e.g. FD00::1/112**	**Unique Local Address** - Similar to IPv4 private addresses. Not used for global routing. Split into two ranges FC00::/8 and FD00::/8. See RFC 4193.	RFC 1918
FE80::/10 **e.g. FE80::21D:72FF:FEF8:9980**	**Link Local** - Network segment specific address. Not routed outside of the segment.	N/A
2001:0000::/32 **e.g. 2001:0000:0001:0001:21D:72FF:FEF8:9980**	**Teredo** - Mechanism for tunnelling IPv6 through IPv4.	N/A
2001:0002::/48 **e.g. 2001:0002:0001::0001**	**Benchmarking** - Used only in documentation. Not used as source or destination addresses in live networks.	198.18.0.0/15

2001:0010::/28 **e.g. 2001:10:0200:00AA::0001**	**Orchid** - Used for fixed-term experimentation.	N/A
2002::/16 **e.g. 2002:ABCD:12EE:0001::0001**	**6to4** - A 6to4 gateway adds its IPv4 address to 2002**::**/16 creating a /48 IPv6 prefix.	N/A
2001:0DB8::/32 **e.g. 2001:0DB8:0001:0002::0001**	**Documentation** - Used only in documentation and examples. Not used as source or destination addresses in live networks.	192.0.2.0/24 198.51.100.0/24 203.0.113.0/24
2000::/3 **e.g. 2001:0B0B::21D:72FF:FEF8:9980**	**Global Unicast** - Globally routable IPv6 addresses.	Any non-private IPv4 address.
FF00::/8 **e.g. FF01:0:0:0:0:0:0:2**	**Multicast** - Used to identify multicast groups, such as all nodes, all routers, node solicited etc. etc.	224.0.0.0/4

Table 3-2: IPv6 Addresses and Types

This paper shows examples and explains the use of the unspecified address [**::**/128], unique local [FD00**::**/8], link local [FE80**::**/10], global unicast [2000**::**/3], 6to4 [2002**::**16] and multicast [FF00**::**/8] addresses in an end-to-end routed IPv6 deployment.

3.3 IPv6 Address Notation

Figure 3-5 below shows an example of long-hand IPv6 address notation using an IANA delegated global unicast address in the 2000**::**/3 range. IPv6 addresses are expressed in hexadecimal (upper or lower case may be used). The network portion is always expressed as a prefix length; the IPv4 network mask notation is never used.

2001:0A0A:0000:0000:CA00:0CFF:FEF4:0000

2	0	0	1	0	A	0	A	0	0	0	0	0	0	0	0	C	A	0	0	0	C	F	F	F	E	F	4	0	0	0	0
0010	0000	0000	0001	0000	1010	0000	1010	0000	0000	0000	0000	0000	0000	0000	0000	1100	1010	0000	0000	0000	1100	1111	1111	1111	1110	1111	0100	0000	0000	0000	0000
Bits 001 + Global Routing Prefix 45 x bits												**Sub-net ID 16 x bits**				**Interface ID 64 x bits**															

Figure 3-5: Example Notation for an IPv6 Global Unicast Address

IPv6 addresses may be expressed in short-hand by omitting leading zeros, so the above address could be expressed as:

2001**:**A0A**:**0**:**0**:**CA00**:**CFF**:**FEF4**:**0

A further short cut is to replace consecutive zeros with a double colon:

2001**:**A0A**::**CA00**:**CFF**:**FEF4**:**0

Substituting consecutive zeros must only take place once within an address otherwise there will be ambiguity about precisely where and how many zeros have been substituted.

3.4 IPv6 Link Local Address

By default, every IPv6 interface possesses an auto-generated link local address. The link local address is generated before any other IPv6 addresses are allocated to an interface - global unicast addresses, for example.

The 48-bit MAC addresses of the fa0/0 LAN interfaces on routers R1 and R2 are configured manually to **AA**:AA:11:11:11:11 and **AA**:AA:22:22:22:22 respectively. This makes packets easier to distinguish in the packet captures shown later and helps to clarify bit manipulation when the interfaces construct their link local and global unicast End User Identifier (EUI-64) addresses.

```
R1(config-if)#mac-address AAAA.1111.1111

R2(config-if)#mac-address AAAA.2222.2222
```

Ethernet is canonical, which means that bytes are transmitted on the wire left to right but the bits within each byte are transmitted right to left i.e. bit-1, the low-order bit, is transmitted first. The first two low-order bits of the first byte of a MAC address have special significance.

Figure 3-6 below shows the bits and bytes of R2's MAC address. Bit-1 of byte-1 is the **I**ndividual/**G**roup bit -value 0 means that this is a unicast rather than multicast MAC address. Bit-2 of byte-1 is the **U**niversal/**L**ocal bit - value 1 means that this is a manually assigned MAC address.

byte-1								byte-2	byte-3	byte-4	byte-5	byte-6
A				A								
b8	b7	b6	b5	b4	b3	b2	b1	AA	22	22	22	22
1	0	1	0	1	0	1	0					
						U/L	I/G					

Figure 3-6: MAC Address U/L and I/G Bits

Enabling IPv6 on a device's interface causes the interface to automatically assign itself an IPv6 link local address. A link local address is unique and is never routed outside of the network segment to which the interface is attached. Devices use link local addresses to communicate with each other on a specific network segment without an administrator having to manually assign an IPv6 address.

Each interface automatically derives its link local address using the interface's MAC address. Figure 3-7 below shows the IPv6 link local address FE80::A8AA:22FF:FE22:2222 that R2 automatically assigns to interface fa0/0. It derives this by inserting FFFE after byte-3 and before byte-4 of the 48-bit MAC address and by flipping the universal/local bit (bit-2) of byte-1. It then appends the resulting 8 x bytes to the reserved IPv6 network address FE80::/10. The network address + interface ID is known as a EUI-64 address; the first 64 x bits denote the network portion and the remaining 64 x bits the interface ID of the 128-bit IPv6 link local address.

```
FE80:0000:0000:0000:A8AA:22FF:FE22:2222
```

EUI-64 Interface Identifier														
byte-1								byte-2	byte-3	byte-4	byte-5	byte-6	byte-7	byte-8
A				8										
b8	b7	b6	b5	b4	b3	b2	b1	AA	22	FF	FE	22	22	22
1	0	1	0	1	0	0	0							

Figure 3-7: R2's Link Local EUI-64 Address

Interfaces use a similar mechanism to derive their auto-configured EUI-64 global unicast addresses. Stateful and stateless auto-configuration is explained in the next section.

3.5 IPv6 Stateful and Stateless Auto-Configuration

IPv6 unicast prefixes are allocated using "stateless" or "stateful" auto-configuration. Sample DHCPv6 configurations and packet captures are shown in Section 4.4 DHCPv6.

3.5.1 Stateful Auto-Configuration

DHCPv6 stateful auto-configuration is somewhat similar to DHCP in IPv4 in that it allocates addresses, manages leases and disseminates services information, such as DNS server addresses and domain names etc. But, apart from the difference in the format and length of IPv4 and IPv6 addresses, there are some fundamental operational differences that DHCPv6 must address - pun fully intended.

Central DHCPv6 servers may allocate a unicast sub-net and address to a router's LAN interface leaving the clients to auto-configure their IPv6 unicast addresses from the router. Or, if a unicast address has been pre-configured on the router's LAN interface, the DHCPv6 server may allocate unicast addresses direct to the clients. Neighbour Discovery (ND) parameters are configurable on the router's LAN interface for it to instruct the clients how to behave. For example, if a router transmits a router advertisement with the "nd managed address" flag set to 1, the clients know they must acquire their addresses using DHCPv6. If the router advertisement has the "nd other configuration" flag set to 1, the clients know they must acquire DNS and other services information using DHCP. If the "nd managed address" flag is set to 0 and the "nd other configuration" flag is set to 1, the clients obtain their IPv6 unicast address from the router using auto-configuration and services information using DHCP. This is exactly how the DHCPv6 demo in this tutorial has been set up.

Another difference is that IPv6 devices on a network segment generate their own link local addresses; a device's default router address on the segment is a link local address. A centralised DHCPv6 server cannot possibly keep track of and manage both unicast and link local addresses so DHCPv6 servers do not allocate a "default gateway" address to clients, they only allocate unicast addresses and services addresses.

3.5.2 Stateless Auto-Configuration

Stateless auto-configuration requires little or no manual intervention other than to configure a private or global unicast IPv6 address on the LAN interface of a router. The hosts attached to the LAN generate their own link local addresses and the router provides the hosts with the LAN sub-net information using router advertisements. It is stateless because there is no central device allocating the unicast addresses or managing leases and renewals. This is fine for small networks.

3.5.3 IPv6 Address Delegation

IPv6 address delegation is an elegant solution that uses DCHPv6 to allocate unicast addresses to routers leaving the clients to auto-configure their unicast addresses and to obtain services information, such as DNS addresses and a domain name, using DHCPv6.

Figure 3-8 below shows how address delegation would work in a service provider network. The DHCPv6 server assigns a unicast prefix block to the SP router, which, in turn, sub-divides the block and allocates a sub-net range to the customer CE router. The customer hosts acquire their unicast prefixes from the CE router using stateless auto-configuration in the usual way i.e. router solicitations and advertisements.

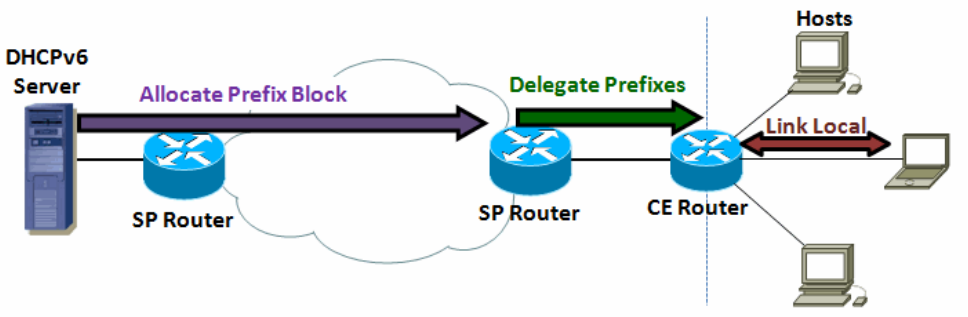

Figure 3-8: IPv6 Address Delegation

There isn't a DHCPv6 server in the lab. However, it is easy to simulate the set-up above by configuring R1 as a DHCPv6 "Delegating" Server and R2 as its DHCPv6 "Client". The CE Client is itself configured as a DHCPv6 server but only to disseminate DNS information to hosts on the local LAN segment.

The configurations and packet captures for the above scenario are in Section 4.4 DHCPv6.

4 IPv6 Functionality

Unlike IPv4, IPv6 never uses layer-2 or layer-3 broadcasts and relies exclusively on the following packet types to function:

- Unicast
- Multicast
- Anycast

The application of these is identical to IPv4. Anycast is simply the use of identical IPv6 destination addresses at different points in the network. A source device sends packets to the topologically closest Anycast destination determined by the routing protocol. This is similar, for example, to the way that Anycast Rendezvous Points are reached in IPv4 multicast networks.

4.1 Multicast Addresses

IPv6 operates using a number of multicast addresses and multicast address types. Table 4-1 below shows some of the well known Link Local-Scope multicast addresses discussed in this tutorial i.e. the All Nodes, All Routers, Multicast Listener Discovery and Solicited Node multicast addresses. IPv6 employs many other multicast address types and scopes, such as the Node Local and Site Local scopes. Please visit the following URL for a full list of the currently allocated IPv6 multicast addresses and scopes.

http://www.iana.org/assignments/ipv6-multicast-addresses/ipv6-multicast-addresses.xhtml

Multicast Address	Description
FF02:0:0:0:0:0:0:1	All Nodes Address
FF02:0:0:0:0:0:0:2	All Routers Address
FF02:0:0:0:0:0:0:16	All MLDv2-capable routers
FF02::1:FF00:0000/104	Solicited-Node Address
FF02:0:0:0:0:2:FF00::/104	Node Information Queries

Table 4-1: IPv6 Link Local-Scope Multicast Addresses Discussed in this Paper

The IPv6 multicast address has the format shown in Table 4-2 below.

Format Prefix								Flags				Scope				Group ID - 112 bits
F				F				Variable				Variable				Permanent or transient multicast group
1	1	1	1	1	1	1	1	0	R	P	T	0/1	0/1	0/1	0/1	

Table 4-2: IPv6 Multicast Address Format

The Format Prefix is always FF.

The Flag bits have the following significance. RFCs define the extended formats of multicast packets that use the R and / or P flags. A description of these extended formats is beyond the scope of this paper.

0 - Reserved
R - The packet contains the embedded IPv6 address of a Rendezvous Point.
P - The packet contains an embedded IPv6 unicast prefix for dynamic address allocation.
T - Transient (0 = permanent well-known multicast address, 1 = dynamic multicast address).

The Scope bits have the following significance:

0 - Reserved
1 - Interface-Local scope
2 - Link-Local scope
3 - Reserved
4 - Admin-Local scope
5 - Site-Local scope
6 - Unassigned
7 - Unassigned
8 - Organization-Local scope
9 - Unassigned
A - Unassigned
B - Unassigned
C - Unassigned
D - Unassigned
E - Global scope
F - Reserved

For example, the All Routers multicast address **FF02:**0:0:0:0:0:0:2 has the following Flag and Scope bits set:

Flags (0000):
0 = Reserved
R = 0 (no embedded RP address present)
P = 0 (no embedded unicast prefix present)
T = 0 (permanent, well known address)

Scope (0010):
2 = Link Local

4.2 Multicast Listener Discovery Protocol

The lab is not running multicast streams so it is not necessary to understand MLDv2's full functionality and capabilities for the purpose of this tutorial but as this paper includes a number of packet captures that contain MLD messages a very brief description of MLD is provided below.

MLD consists of two types of device:

- MLD capable Routers
- MLD capable Hosts (listeners)

The MLD protocol uses three types of message:

- Query (routers to hosts) ICMPv6 type 130
- Report (hosts to routers) ICMPv6 type 131
- Done (host leave messages) ICMPv6 type 132

MLDv2 is the equivalent of IGMPv3 in IPv4. IPv6 hosts use the MLDv2 report message to inform MLD capable routers about the multicast groups they have joined, or wish to join, on the attached network segment. MLD capable routers send query messages to hosts as and when they need to ascertain information about a host's specific multicast group membership. The MLD router needs only to discover multicast groups for which there is at least one interested host listener.

The host generated MLDv2 report message contains the multicast address that is of interest to the host together with an optional specification of the sources from which the host wishes to accept traffic. A host can also specify an INCLUDE or EXCLUDE filter informing the MLD router whether it wishes to receive multicast traffic from all of the specified sources or from all except the specified sources.

MLD messages are embedded in ICMPv6. The IPv6 hop-count is set to 1 so MLD messages never leave the local network segment. The IPv6 source address of a MLD report message is either the unspecified address "**::**" or the link local address of an interface, and the destination address is FF02**::**16 - the "MLD capable routers" address.

The IPv6 router alert option allows MLD routers to receive and process MLD messages from hosts even if the routers are not themselves listening on a specific destination multicast address.

4.2.1 MLD Protocol - Things to Note in the Lab

R1 and R2 in the lab are MLDv2 capable routers whereas Host and Laptop are simply MLD host listeners. Host and Laptop send MLDv2 reports to their local router (MLD capable routers multicast address FF02**::**16) but there are no multicast streams running in the lab so these report messages have no bearing on the IPv6 operations described in this paper.

Host listeners on the same network segment do not communicate with each other via the local MLD router. They communicate direct with each other provided, of course, that they have joined the relevant groups and are listening for packets with the corresponding multicast addresses. For example, a node that is listening for the solicited node address FF02**::**1:FFXX:XXXX will receive and process multicast packets targeted at this address irrespective of whether it sends a report message to the MLD router for this group or otherwise.

A MLD report message confusingly show hosts setting specific multicast addresses to "exclude" for traffic they actually wish to receive, such as solicited node messages and all nodes multicasts. Although confusing, this is purely down to the specification of the INCLUDE / EXLUDE filter. For neighbour discovery, the source address list in a host's MLD report message is always empty (set to zero) because a host cannot know in advance and specify a pre-defined list of all the source addresses for all of the devices on the network segment. Therefore, the host cannot use the INCLUDE filter. A zero length source list and EXCLUDE filter effectively means that the host wants to receive the specified multicast traffic from any source.

Figure 4-1 below shows Laptop interface Eth0 sending a MLDv2 report message to multicast destination FF02**::**16. It is informing the MLD capable router that it has joined two multicast groups; FF02**::**1:FFF8:9980 (solicited node) and FF02**::**C (SSDP - Simple Service Discovery Protocol). The interface has changed its filter to exclude, is specifying a zero length source list and is, therefore, listening out for traffic to the specified multicast groups from any source address.

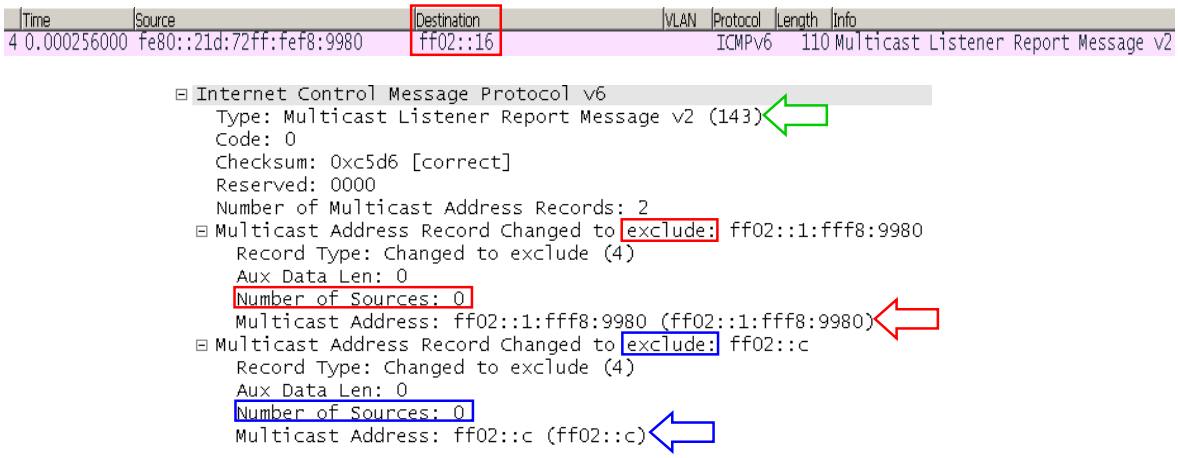

Figure 4-1: Example of MLDv2 Listener Report Message

4.3 Neighbour Discovery

The IPv6 neighbour discovery (ND) mechanism uses ICMPv6 in neighbour solicitation multicast packets to discover local segment neighbours and to perform duplicate address detection (DAD). Neighbour

discovery replaces the IPv4 ARP protocol and is the means by which a sending device determines the MAC address of a neighbour on the network segment.

The ND messages are:

- Neighbour Solicitation ICMPv6 type 135
- Neighbour Advertisement ICMPv6 type 136
- Router Solicitation ICMPv6 type 133
- Router Advertisement ICMPv6 type 134
- Neighbour Redirect ICMPv6 type 137

Please refer to the following URL for a full list of the ICMPv6 types and codes:

http://www.iana.org/assignments/icmpv6-parameters/icmpv6-parameters.xhtml

The packet captures in this section are taken between Laptop and R2.

Router R2's fa0/0 interface has two IPv6 addresses - the link local address FE80**::**A8AA**:**22FF**:**FE22**:**2222 derived from the interface MAC address and the manually configured global unicast address 2001**:**B0B**::**1. Neighbour discovery and duplicate address detection takes place for each of these addresses using exactly the same process.

```
R2#show ipv6 interface fa0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8AA:22FF:FE22:2222
  Global unicast address(es):
    2001:B0B::1, subnet is 2001:B0B::/64
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:FF00:1
    FF02::1:FF22:2222
  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.
```

Interface fa0/0 joins the following multicast groups (shown in red above) in order to receive and process packets that are sent to these multicast addresses:

- FF02**::**1 (short-hand for FF02**::**0001) All Nodes
- FF02**::**2 (short-hand for FF02**::**0002) All Routers
- **FF02::1:FF**00:1 (short-hand for FF02**::**1:FF00:0001) Solicited Node
- **FF02::1:FF**22:2222 Solicited Node

The All Nodes and All Routers groups are self explanatory but the last two groups require some explanation. When an interface becomes active it appends the last 24 bits of each of its link local and unicast IPv6 addresses to the Solicited Node multicast address **FF02::1:FF** and joins both multicast groups. The last 24 bits of interface fa0/0's link local address are 22**:**2222 and the last 24 bits of the global unicast address are 00**:**0001 so interface fa0/0 joins both Solicited Node multicast groups **FF02::1:FF**22**:**2222 and **FF02::1:FF**00**:**1.

Laptop's Eth0 interface does exactly the same with its link local address and auto-configured global unicast prefix (shown below in blue).

```
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  . : home
```

```
Description . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
Physical Address. . . . . : 00-1D-72-F8-99-80
DHCP Enabled. . . . . : Yes
Autoconfiguration Enabled . . . . : Yes
IPv6 Address. . . . . : 2001:b0b::21d:72ff:fef8:9980(Preferred)
Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
IPv4 Address. . . . . : 192.168.1.66(Preferred)
Subnet Mask . . . . . : 255.255.255.0
Lease Obtained. . . . . : 19 December 2013 14:13:56
Lease Expires . . . . . : 20 December 2013 15:37:33
Default Gateway . . . . . : fe80::a8aa:22ff:fe22:2222%10
                            192.168.1.254
DHCP Server . . . . . : 192.168.1.254
DNS Servers . . . . . : 192.168.1.254
                            192.168.1.254
NetBIOS over Tcpip. . . . . : Enabled
```

However, Laptop joins only Solicited Node multicast group **FF02::1:FF**F8**:**9980 because its IPv6 global unicast address is not manually configured so the last 24 bits of both the link local and unicast addresses are identical.

4.3.1 Duplicate Address Detection

An interface determines whether its address is unique on a network segment by transmitting an ICMPv6 neighbour solicitation message to the solicited node multicast group using the unspecified address "**::**" as the source. If a reply is not received, the interface knows that no other interface on the segment possesses the same address.

Figure 4-2 below shows Laptop interface Eth0 transmitting two solicited node multicasts from unspecified source "**::**" - one packet is duplicate address detection for Eth0's link local address the other is duplicate address detection for Eth0's global unicast address.

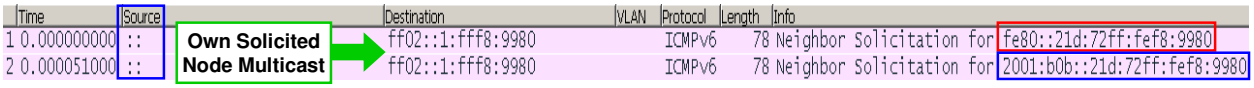

Figure 4-2: Link Local & Unicast Address DAD

Figure 4-3 below shows the ICMPv6 type 135 content of the first packet. It contains the complete link local address of interface Eth0 so a receiving device can do a full comparison of its own and Eth0's link local address.

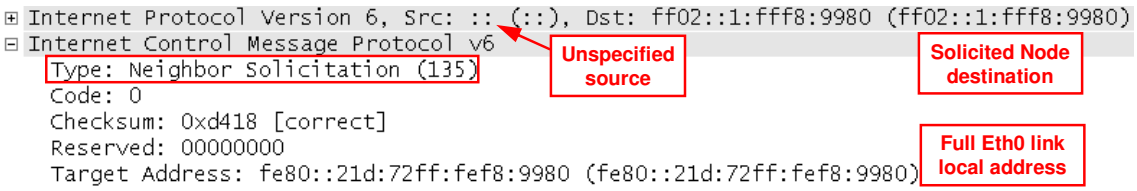

Figure 4-3: Link Local Address DAD - ICMPv6 Content

The DAD multicast mechanism is much more efficient than if broadcasts were used. Only those interfaces on the network segment that have addresses that terminate with the same 24 bits of the solicited node address receive and process the above packets. Every other interface ignores the packets and is unaware that DAD is taking place.

Figure 4-4 below shows layer-2 content of the neighbour solicitation message. Note the destination MAC address. The last 32 bits of the solicited node multicast address are appended to the layer-2 multicast address **33:33:**XX:XX:XX:XX. The source MAC address is the physical address of interface Eth0.

```
⊟ Ethernet II, Src: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80), Dst: 33:33:ff:f8:99:80 (33:33:ff:f8:99:80)
  ⊟ Destination: 33:33:ff:f8:99:80 (33:33:ff:f8:99:80)
      Address: 33:33:ff:f8:99:80 (33:33:ff:f8:99:80)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
  ⊟ Source: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      Address: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
```

Figure 4-4: Link Local Address DAD - Layer-2 Content

The same process is repeated for the global unicast address. Figure 4-5 below shows the ICMPv6 type 135 contents of the second packet. It contains the complete global unicast address of interface Eth0 so a receiving device can do a full comparison of its own and Eth0's global unicast address.

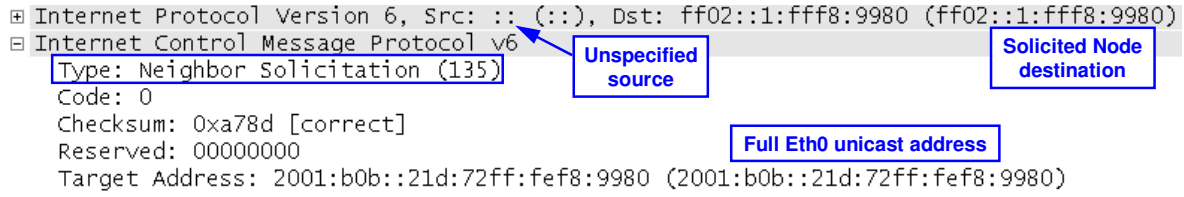

Figure 4-5: Unicast Address DAD ICMPv6 Contents

The layer-2 multicast information is identical to that in Figure 4-4 because Laptop uses its MAC address as the source. The last 32 bits of the solicited node multicast address are appended to the layer-2 multicast address **33:33:**XX:XX:XX:XX.

4.3.2 Neighbour Solicitation

IPv6 neighbour solicitation replaces the IPv4 ARP mechanism. Figure 4-6 below shows the sequence of events when Laptop pings Host. The exchange of packets populates the IPv6 neighbour cache on every device with the MAC address of the adjacent device. The source of the ping is Laptop's unicast address 2001:B0B::21D:72FF:FEF8:9980 and the destination is Host unicast address 2001:A0A::CA00:22FF:FE50:0.

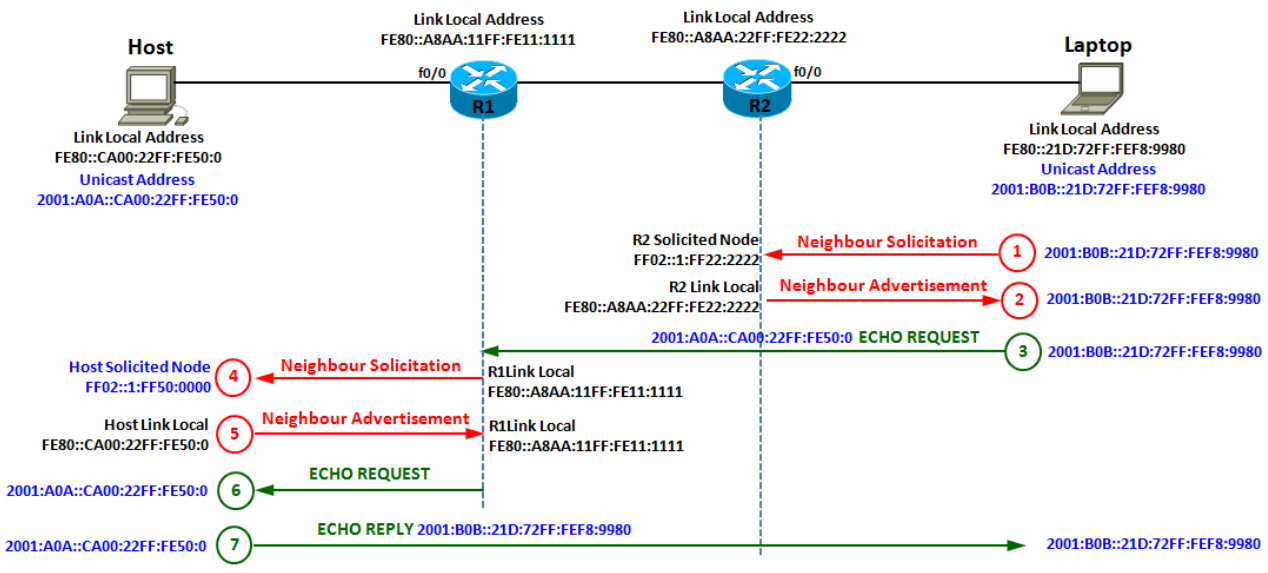

Figure 4-6: Neighbour Solicitation

1) Laptop sends a Neighbour Solicitation from its unicast address to the Solicited Node address of R2's link local address
2) R2 responds with its MAC address using its link local address to Laptop's unicast address
3) Laptop sends an ICMPv6 Echo Request to R2, which forwards it to R1
4) R1 sends a Neighbour Solicitation using its link local address to the Solicited Node address of Host's unicast address
5) Host responds with its MAC address using its link local address to R1's link local address
6) R1 forwards the IMCPv6 Echo Request to Host
7) Host responds with an ICMPv6 Echo Reply from its unicast address to Laptop's unicast address

The output below shows Laptop's IPv6 neighbour cache immediately after the Neighbour Solicitation process completes. The MAC address of R2's link local address is changed from "Stale" to "Reachable". Laptop knows that R2 is a "(Router)" because it had previously received a Router Advertisement message from R2.

```
C:\windows\system32>netsh interface ipv6 show neighbors "Local Area Connection"
Interface 10: Local Area Connection
Internet Address                              Physical Address   Type
--------------------------------------------  -----------------  -----------
fe80::48c4:610b:b522:63ef                     00-1d-72-d6-ca-d2  Stale
fe80::a8aa:22ff:fe22:2222                     aa-aa-22-22-22-22  Reachable (Router)
ff02::2                                       33-33-00-00-00-02  Permanent
ff02::c                                       33-33-00-00-00-0c  Permanent
ff02::16                                      33-33-00-00-00-16  Permanent
ff02::fb                                      33-33-00-00-00-fb  Permanent
ff02::1:2                                     33-33-00-01-00-02  Permanent
ff02::1:3                                     33-33-00-01-00-03  Permanent
ff02::1:ff03:894b                             33-33-ff-03-89-4b  Permanent
ff02::1:ff05:1e0b                             33-33-ff-05-1e-0b  Permanent
ff02::1:ff0b:2d3d                             33-33-ff-0b-2d-3d  Permanent
ff02::1:ff22:2222                             33-33-ff-22-22-22  Permanent
ff02::1:ff22:63ef                             33-33-ff-22-63-ef  Permanent
ff02::1:ff35:ff1e                             33-33-ff-35-ff-1e  Permanent
ff02::1:ff71:457e                             33-33-ff-71-45-7e  Permanent
ff02::1:ff81:1e06                             33-33-ff-81-1e-06  Permanent
ff02::1:ff83:3c59                             33-33-ff-83-3c-59  Permanent
ff02::1:ff90:fae                              33-33-ff-90-0f-ae  Permanent
ff02::1:ff9c:1e85                             33-33-ff-9c-1e-85  Permanent
ff02::1:ffb5:17a4                             33-33-ff-b5-17-a4  Permanent
ff02::1:ffc5:bd95                             33-33-ff-c5-bd-95  Permanent
ff02::1:fff8:9980                             33-33-ff-f8-99-80  Permanent
```

Note the entries with "Permanent" against them. Permanent in this context means static entries that are created during protocol operation. These are all the MAC address of the well-known multicast destinations, including the All Routers address FF02::2, R2's Solicited Node address FF02::1:FF22:2222 etc.

A neighbour's reachability state is one of five possible values:

INCOMLETE: Address resolution is in progress.

DELAY: The neighbour is no longer reachable. Traffic was recently sent to the neighbour but upper layer protocols may still be processing requests. Wait briefly until sending further unicast neighbour solicitation probes.

PROBE: The neighbour is no longer reachable: Neighbour solicitation probes are currently being sent to verify reachability.

REACHABLE: The neighbour is reachable. The neighbour has recently responded to neighbour solicitation probes and traffic has been sent / received.

STALE: The neighbour is no longer reachable. No attempts will be made to verify neighbour reachability until traffic is sent to the neighbour.

The amount of time a neighbour remains in the REACHABLE state before it reverts to STALE is dependent on the device operating system. For router R2 (Cisco 2621) the default value is 30 seconds. This is configurable using interface neighbour discovery parameters.

```
show ipv6 interface fa0/0
  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.
```

The "Age" of a neighbour entry is the number of minutes since it was last reachable. In the output below from R2 it is 8 minutes since Laptop pinged Host. The neighbour entries are Laptop's unicast address (the ping source) and MAC address via interface fa0/0, and R1's link local address and MAC address via fa0/1.

```
R2#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
.
.
2001:B0B::21D:72FF:FEF8:9980                8 001d.72f8.9980  STALE Fa0/0
FE80::CA01:6FF:FEB8:1                       8 c801.06b8.0001  STALE Fa0/1
```

The following output from R2 was taken as soon as Laptop initiated the ping to Host. The neighbour entries for Laptop and R1 transition from STALE to DELAY to REACHABLE (Age = 0) before returning to STALE after 30 seconds.

```
R2#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
.
.
2001:B0B::21D:72FF:FEF8:9980               59 001d.72f8.9980  STALE Fa0/0
FE80::CA01:6FF:FEB8:1                      59 c801.06b8.0001  STALE Fa0/1

R2#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
.
.
2001:B0B::21D:72FF:FEF8:9980               59 001d.72f8.9980  DELAY Fa0/0
FE80::CA01:6FF:FEB8:1                      59 c801.06b8.0001  DELAY Fa0/1

R2#show ipv6 neighbors
IPv6 Address                              Age Link-layer Addr State Interface
.
.
2001:B0B::21D:72FF:FEF8:9980                0 001d.72f8.9980  REACH Fa0/0
FE80::CA01:6FF:FEB8:1                       0 c801.06b8.0001  REACH Fa0/1
```

Figure 4-7 below shows the IPv6 content of the ICMPv6 type 135 neighbour solicitation packet from Laptop to the default router R2. The source address is Laptop's unicast address and the destination is R2's link local solicited node address.

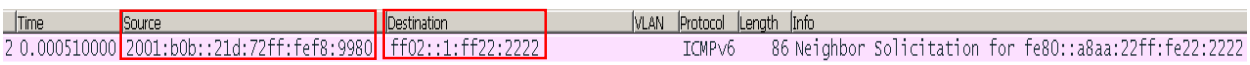

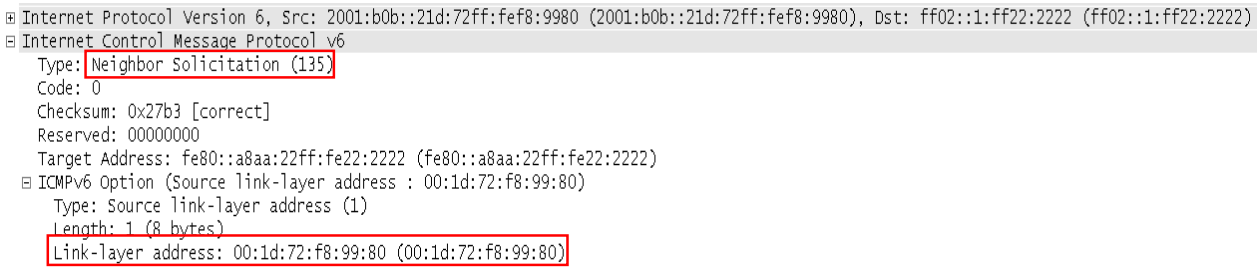

Figure 4-7: Neighbour Solicitation ICMPv6 Content

Figure 4-8 below shows the layer-2 content of the neighbour solicitation packet. The source address is Laptop's MAC address and the destination address is the last 32 bits of R2's solicited node address mapped into MAC address **33:33:**XX:XX:XX:XX.

Figure 4-8: Neighbour Solicitation Layer-2 Content

4.3.3 Neighbour Advertisement

Figure 4-9 below shows the IPv6 content of the ICMPv6 type 136 neighbour advertisement response from R2 to Laptop. The source address is R2's link local address and the destination is Laptop's unicast address.

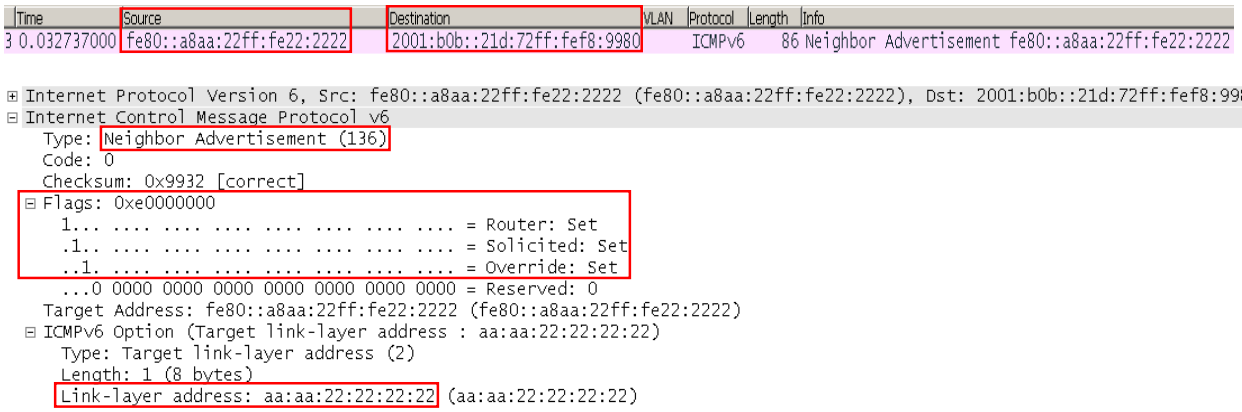

Figure 4-9: Neighbour Advertisement ICMPv6 Content

The Flags in the neighbour advertisement have the following significance:

R Router flag: The R-bit indicates that the sender is a router.

S Solicited flag: The S-bit indicates that the advertisement was sent in response to a Neighbour Solicitation.

O Override flag: The O-bit indicates that the receiver should override an existing IPv6 neighbour cache entry with the link-layer address in this packet.

Figure 4-10 below shows the layer-2 content of the neighbour advertisement packet. The source address is R2's MAC address and the destination address is Laptop's MAC address.

```
⊟ Ethernet II, Src: aa:aa:22:22:22:22 (aa:aa:22:22:22:22), Dst: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
  ⊟ Destination: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      Address: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
  ⊟ Source: aa:aa:22:22:22:22 (aa:aa:22:22:22:22)
      Address: aa:aa:22:22:22:22 (aa:aa:22:22:22:22)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
```

Figure 4-10: Neighbour Advertisement Layer-2 Content

4.3.4 Router Solicitation

Figure 4-11 below shows an ICMP type 133 router solicitation packet from Laptop interface Eth0 to the All Routers multicast address FF02::02. Having completed DAD, interface Eth0 now specifies it full link local address as the source address. Router solicitations are transmitted only from the link local address because they relate solely to addresses on the local network segment.

```
| Time          | Source                | Destination | VLAN | Protocol | Length | Info                                          |
3 0.000141000   fe80::21d:72ff:fef8:9980  ff02::2              ICMPv6     70 Router Solicitation from 00:1d:72:f8:99:80

⊞ Internet Protocol Version 6, Src: fe80::21d:72ff:fef8:9980 (fe80::21d:72ff:fef8:9980), Dst: ff02::2 (ff02::2)
⊟ Internet Control Message Protocol v6
    Type: Router Solicitation (133)
    Code: 0
    Checksum: 0x6202 [correct]
    Reserved: 00000000
  ⊟ ICMPv6 Option (Source link-layer address : 00:1d:72:f8:99:80)
      Type: Source link-layer address (1)
      Length: 1 (8 bytes)
      Link-layer address: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
```

Full Eth0 link local address

All Routers destination

Figure 4-11: Router Solicitation ICMPv6 Contents

Figure 4-12 below shows how the IPv6 multicast information is mapped onto the layer-2 addresses. Note the destination MAC address. The last 32 bits of the all routers multicast address are appended to the layer-2 multicast address **33:33:**XX:XX:XX:X**2**. The source MAC address is the physical address of interface Eth0.

```
⊟ Ethernet II, Src: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80), Dst: 33:33:00:00:00:02 (33:33:00:00:00:02)
  ⊟ Destination: 33:33:00:00:00:02 (33:33:00:00:00:02)
      Address: 33:33:00:00:00:02 (33:33:00:00:00:02)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
  ⊟ Source: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      Address: 00:1d:72:f8:99:80 (00:1d:72:f8:99:80)
      .... ..0. .... .... .... .... = LG bit: Globally unique address (factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
```

Figure 4-12: Router Solicitation Layer-2 Contents

4.3.5 Router Advertisement

Router R2 replies to the router solicitation message with an ICMPv6 type 134 router advertisement message to the All Nodes multicast destination FF02::01. The router advertisement message contains the network segment's unicast prefix information. Figure 4-13 below shows the IPv6 contents of the router advertisement message with R2's link local address as the source.

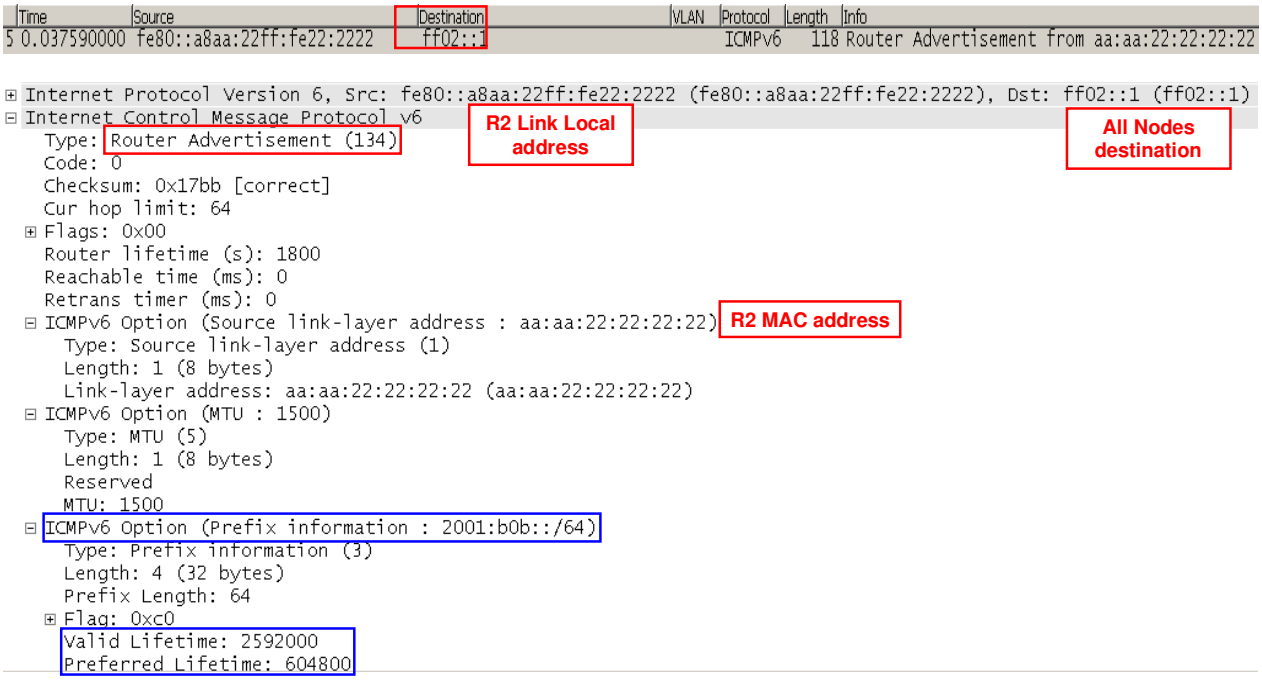

Figure 4-13: Router Advertisement ICMPv6 Contents

Routers also transmit router advertisements periodically. In the lab, R1 and R2 send them every 200 seconds. This interval is configurable using an interface neighbour discovery parameter.

```
show ipv6 interface fa0/0
  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.
```

The lifetime values of the unicast prefix that the router specifies are configurable. The output below from R2 shows the lifetime values it receives from R1 (FE80::CA01:BFF:FEF0:1) for the private prefix FD00::/12.

```
R2#show ipv6 routers
Router FE80::CA01:BFF:FEF0:1 on FastEthernet0/1, last update 0 min
  Hops 64, Lifetime 1800 sec, AddrFlag=0, OtherFlag=0, MTU=1500
  HomeAgentFlag=0, Preference=Medium
  Reachable time 0 msec, Retransmit time 0 msec
  Prefix FD00::/112 onlink autoconfig
    Valid lifetime 2592000, preferred lifetime 604800
```

Figure 4-14 below shows the relationship between the valid and preferred prefix lifetime values, which are expressed in seconds. The valid lifetime must be greater than or equal to the preferred lifetime.

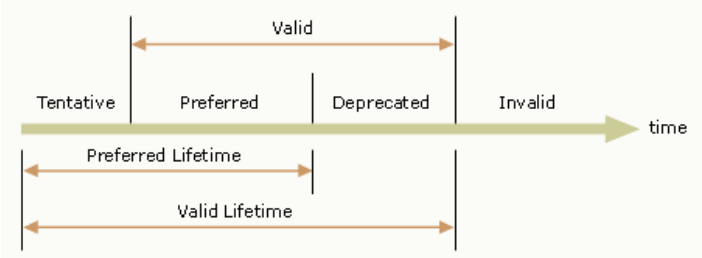

Figure 4-14: Router Advertisement Valid and Preferred Lifetimes

Stateless auto-configured addresses are in one or more of the following states:

- **Tentative** - duplicate address detection is in progress and the address has not yet been confirmed as being unique.
- **Preferred** - duplicate address detection has completed. The device can send / receive unicast traffic to / from a preferred address.
- **Deprecated** - the address is still valid but its use is not recommended for new sessions. Existing sessions can continue to send / receive unicast traffic to / from a deprecated address.
- **Valid** - the device can send / receive traffic to / from the valid address.
- **Invalid** - the device can no longer send / receive unicast traffic to / from the invalid address. The address becomes invalid once the valid lifetime timer expires.

Figure 4-15 below shows how the IPv6 multicast information is mapped onto the layer-2 addresses. Note the destination MAC address. The last 32 bits of the all nodes multicast address are appended to the layer-2 multicast address **33:33:**XX:XX:XX:X**1**. The source MAC address is the physical address of R2's interface fa0/0.

```
⊞ Frame 5: 118 bytes on wire (944 bits), 118 bytes captured (944 bits) on interface 0
⊟ Ethernet II, Src: aa:aa:22:22:22:22 (aa:aa:22:22:22:22), Dst: 33:33:00:00:00:01 (33:33:00:00:00:01)
  ⊟ Destination: 33:33:00:00:00:01 (33:33:00:00:00:01)
      Address: 33:33:00:00:00:01 (33:33:00:00:00:01)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...1 .... .... .... .... = IG bit: Group address (multicast/broadcast)
  ⊟ Source: aa:aa:22:22:22:22 (aa:aa:22:22:22:22)
      Address: aa:aa:22:22:22:22 (aa:aa:22:22:22:22)
      .... ..1. .... .... .... .... = LG bit: Locally administered address (this is NOT the factory default)
      .... ...0 .... .... .... .... = IG bit: Individual address (unicast)
    Type: IPv6 (0x86dd)
```

Figure 4-15: Router Advertisement Layer-2 Contents

4.4 DHCPv6

4.4.1 Lab DHCPv6 Address Delegation

Figure 4-16 below shows the topology and connections for DHCPv6 address delegation in the lab. R1 is configured as a DHCPv6 server and R2 as a DHCPv6 Client. R1 delegates block 2001:B1B::/64 to R2. R2 applies this sub-net to interface fa0/0 and assigns fa0/0 IPv6 unicast address 2001:B1B::1. All hosts attached to fa0/0 obtain their IPv6 unicast addresses from R2 using auto-configuration (router solicitations / advertisements). R2 sets the "nd other-config-flag" in its router advertisements to tell hosts that they should use DHCP to obtain services information. R2 acts as a DHCP server to allocate only DNS IPv6 addresses to hosts.

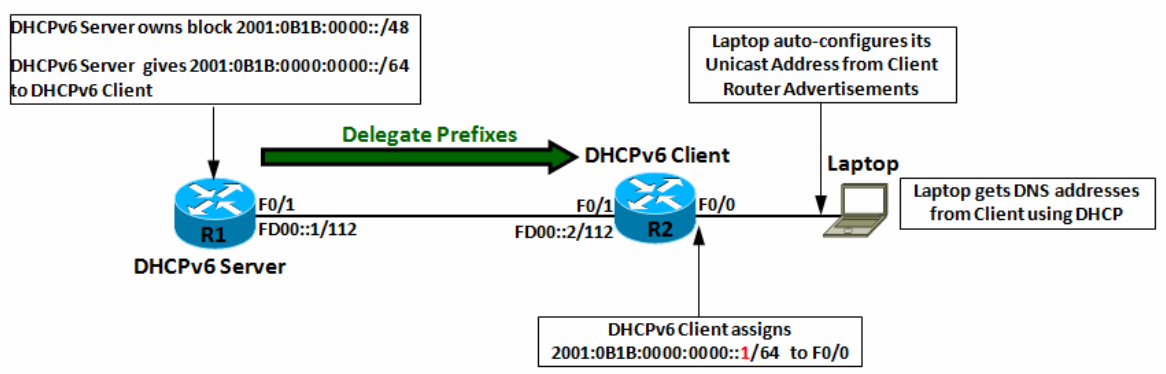

Figure 4-16: DHCPv6 Address Delegation

The relevant parts of R1's and R2's configuration are shown below.

Router R1 (DHCPv6 Server)

```
!
! Pool ISP defines the valid & preferred lifetimes for prefixes
! and references local prefix pool ISP-ADDRESS that holds
! the prefix block to delegate to DHCv6 Client R2
!
ipv6 dhcp pool ISP
 prefix-delegation pool ISP-ADDRESSES lifetime 3600 1800

!
! Interface fa0/1 connects to DHCPv6 Client R2 and
! references pool ISP with the statement "ipv6 dhcp
! server ISP"
!
interface FastEthernet0/1
 ip address 10.1.1.1 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 ipv6 address FD00::1/112
 ipv6 dhcp server ISP
 ipv6 router isis

!
! Local pool ISP holds the prefix to delegate to R2. The
! pool has one prefix of length /48 out of which it delegates
! a /64 to Client R2
!
ipv6 local pool ISP-ADDRESSES 2001:B1B::/48 64
```

Router R2 (DHCPv6 Client)

```
!
! R2 has its own pool from which it assigns DNS addresses
! to hosts on LAN interface fa0/0
!
ipv6 dhcp pool DNS
 dns-server 2001:4860:4860::8888
 dns-server 2001:4860:4860::8844

!
! Interface fa0/0 connects to the local LAN. It references
! its local DHCP pool DNS to allocate DNS addresses to hosts.
!
! The FROM-ISP statement specifies that address ::1 is
! allocated to fa0/0 from the prefix block that DHCPv6
! server R1 delegates
!
! The "nd other-config-flag" tells R2 to set the other config
! flag in router advertisements so that hosts on the LAN
! use DHCP to obtain services information, such as DNS
! addresses.
!
interface FastEthernet0/0
 mac-address aaaa.2222.2222
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 ipv6 address FROM-ISP ::1/64
 ipv6 nd other-config-flag
 ipv6 dhcp server DNS
 ipv6 router isis

!
! Interface fa0/1 connects to DHCPv6 Server R1. The "ipv6
! dhcp client pd FROM-ISP" statement instructs R2 to get
! its IPv6 addresses from "prefix delegator" R1.
!
interface FastEthernet0/1
 ip address 10.1.1.2 255.255.255.252
 ip router isis
 duplex auto
 speed auto
 ipv6 address FD00::2/112
 ipv6 dhcp client pd FROM-ISP
 ipv6 router isis
```

The output below on R1 shows DHCPv6 pool ISP referencing prefix pool ISP-ADDRESS and setting the preferred and valid lifetime for the prefix block it allocates to R2.

```
R1#show ipv6 dhcp pool
DHCPv6 pool: ISP
  Prefix pool: ISP-ADDRESSES
               preferred lifetime 1800, valid lifetime 3600
  Active clients: 1
```

The output below on R1 shows that DHCPv6 Client R2 is connected to R1 over interface fa0/1 using client's IPv6 address FD00::2. R1 has delegated to R2 the prefix block 2001:B1B::/64 with the specified preferred and valid lifetimes. The DHCP Unique Identifier (DUID) is R2 client's fa0/0 MAC address appended to the value 00030001. The DUID allows individual server / client sessions to be identified should multiple server / client sessions exist.

```
R1#show ipv6 dhcp binding
Client: FD00::2
  DUID: 00030001AAAA22222222
  Username : unassigned
  Interface : FastEthernet0/1
  IA PD: IA ID 0x00050001, T1 900, T2 1440
    Prefix: 2001:B1B::/64
            preferred lifetime 1800, valid lifetime 3600
            expires at Dec 20 2013 02:00 AM (2943 seconds)
```

The output below on R2 shows the local pool that it uses to disseminate DNS only information to the local clients.

```
R2#show ipv6 dhcp pool
DHCPv6 pool: DNS
  DNS server: 2001:4860:4860::8888
  DNS server: 2001:4860:4860::8844
  Active clients: 0
```

The output below on R2 shows that interface fa0/0 is in server mode using pool DNS to distribute DNS IPv6 information. Interface fa0/1 is in client mode connecting to Server R1 at address FD00::1. The DUID is server R1's fa0/0 MAC address and R1 has delegated the prefix block 2001:B1B::/64.

```
R2#show ipv6 dhcp interface
FastEthernet0/0 is in server mode
  Using pool: DNS
  Preference value: 0
  Hint from client: ignored
  Rapid-Commit: disabled
FastEthernet0/1 is in client mode
  State is OPEN
  Renew will be sent in 00:08:04
  List of known servers:
    Reachable via address: FD00::1
    DUID: 00030001AAAA11111111
    Preference: 0
    Configuration parameters:
      IA PD: IA ID 0x00050001, T1 900, T2 1440
        Prefix: 2001:B1B::/64
                preferred lifetime 1800, valid lifetime 3600
                expires at Dec 20 2013 02:00 AM (2343 seconds)
  Prefix name: FROM-ISP
  Rapid-Commit: disabled
```

The output below on R2 shows that the server delegated IPv6 block has been allocated to interface fa0/0 with the specified address "**::1**". Configuring "ipv6 nd other-config-flag" on fa0/0 causes R2 to tell hosts on the LAN that they should "use DCHP to obtain other information" e.g. DNS information. Note that configuring DHCPv6 on R2 has caused it to join two new multicast groups - FF02::1:2 All DHCP Agents and FF05::1:3 All DHCP Servers. FF02::1:2 has link local scope and FF05::1:3 has site local scope.

```
R2#show ipv6 interface fa0/0
FastEthernet0/0 is up, line protocol is up
  IPv6 is enabled, link-local address is FE80::A8AA:22FF:FE22:2222
  Global unicast address(es):
```

```
    2001:B1B::1, subnet is 2001:B1B::/64 [PRE]
      valid lifetime 3593 preferred lifetime 1793
  Joined group address(es):
    FF02::1
    FF02::2
    FF02::1:2
    FF02::1:FF00:1
    FF02::1:FF22:2222
    FF05::1:3
  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.
  Hosts use DHCP to obtain other configuration.
```

The output below shows Laptop's Ethernet interface before R2 is powered on. It has only an IPv4 address, IPv4 default gateway, IPv4 DNS addresses and an IPv6 link local address.

```
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  . : home
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-1D-72-F8-99-80
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.66(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : 25 December 2013 07:52:12
   Lease Expires . . . . . . . . . . : 26 December 2013 07:52:12
   Default Gateway . . . . . . . . . : 192.168.1.254
   DHCP Server . . . . . . . . . . . : 192.168.1.254
   DHCPv6 IAID . . . . . . . . . . . : 218111346
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-11-51-0F-62-00-1D-72-F8-99-80
   DNS Servers . . . . . . . . . . . : 192.168.1.254
                                       192.168.1.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

The output below shows Laptop's Ethernet interface after R2 is powered on. The interface has the IPv4 information shown previously but has acquired an IPv6 unicast address and IPv6 default gateway using auto-configuration, and the two IPv6 DNS addresses using DHCP.

```
Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . : home
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-1D-72-F8-99-80
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:b1b::21d:72ff:fef8:9980(Preferred)
   Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.66(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : 25 December 2013 07:52:12
   Lease Expires . . . . . . . . . . : 26 December 2013 07:52:12
   Default Gateway . . . . . . . . . : fe80::a8aa:22ff:fe22:2222%10
                                       192.168.1.254
   DHCP Server . . . . . . . . . . . : 192.168.1.254
   DHCPv6 IAID . . . . . . . . . . . : 218111346
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-11-51-0F-62-00-1D-72-F8-99-80
   DNS Servers . . . . . . . . . . . : 2001:4860:4860::8888
                                       2001:4860:4860::8844
```

```
                                    192.168.1.254
                                    192.168.1.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

The output below shows Laptop's IPv6 configuration parameters. Router discovery is enabled so that it can transmit neighbour solicitations to perform auto-configuration. However, the "Managed Address" and "Other Stateful" configuration parameters are now set according to the flags that R2 set in its router advertisements - managed set to 0 and other to 1 i.e. do not use DHCP to obtain an IPv6 unicast address but use DCHP to obtain DNS information.

```
C:\windows\system32>netsh interface ipv6 show interface 10

Interface Local Area Connection Parameters
----------------------------------------------
IfLuid                             : ethernet_5
IfIndex                            : 10
Compartment Id                     : 1
State                              : connected
Metric                             : 20
Link MTU                           : 1500 bytes
Reachable Time                     : 22000 ms
Base Reachable Time                : 30000 ms
Retransmission Interval            : 1000 ms
DAD Transmits                      : 1
Site Prefix Length                 : 64
Site Id                            : 1
Forwarding                         : enabled
Advertising                        : disabled
Neighbor Discovery                 : enabled
Neighbor Unreachability Detecion   : enabled
Router Discovery                   : enabled
Managed Address Configuration      : disabled
Other Stateful Configuration       : enabled
Weak Host Sends                    : disabled
Weak Host Receives                 : disabled
Use Automatic Metric               : enabled
Ignore Default routes              : disabled
```

After R1 delegates block 2001:B1B::/64 to R2 and R2 assigns address 2001:B1B::1/64 to fa0/0, Laptop transmits a router solicitation to the All Routers multicast address FF02::2. It receives a router advertisement reply to the All Nodes multicast address FF02::1. Figure 4-17 below shows the IPv6 router solicitation and advertisement and the ICMPv6 contents of the advertisement with the Other Configuration flag set instructing Laptop to obtain its DNS address using DHCP.

	Time	Source	Destination	VLAN	Protocol	Length	Info
1	0.000000000	fe80::21d:72ff:fef8:9980	ff02::2		ICMPv6	70	Router Solicitation from 00:1d:72:f8:99:80
2	0.052257000	fe80::a8aa:22ff:fe22:2222	ff02::1		ICMPv6	118	Router Advertisement from aa:aa:22:22:22:22

```
⊞ Internet Protocol Version 6, Src: fe80::a8aa:22ff:fe22:2222 (fe80::a8aa:22ff:fe22:2222), Dst: ff02::1 (ff02::1)
⊟ Internet Control Message Protocol v6
    Type: Router Advertisement (134)
    Code: 0
    Checksum: 0xccd1 [correct]
    Cur hop limit: 64
  ⊟ Flags: 0x40
      0... .... = Managed address configuration: Not set
      .1.. .... = Other configuration: Set
      ..0. .... = Home Agent: Not set
      ...0 0... = Prf (Default Router Preference): Medium (0)
      .... .0.. = Proxy: Not set
      .... ..0. = Reserved: 0
    Router lifetime (s): 1800
    Reachable time (ms): 0
    Retrans timer (ms): 0
  ⊞ ICMPv6 Option (Source link-layer address : aa:aa:22:22:22:22)
  ⊞ ICMPv6 Option (MTU : 1500)
  ⊞ ICMPv6 Option (Prefix information : 2001:b1b::/64)
```

Figure 4-17: Router Advertisement with Other Configuration Set

As soon as Laptop receives the router advertisement it transmits a DHCPv6 request to the All DHCP Agents multicast address FF02::1:2 requesting the DNS addresses. R2 replies to Laptop's link local address. Figure 4-18 below shows the contents of the DHCPv6 reply containing the two DNS addresses.

```
No.  Time          Source                  Destination             VLAN  Protocol  Length  Info
1    0.000000000   fe80::21d:72ff:fef8:9980 ff02::1:2                     DHCPv6    118     Information-request XID:
2    0.044416000   fe80::a8aa:22ff:fe22:2222 fe80::21d:72ff:fef8:9980     DHCPv6    134     Reply XID: 0xfc14a CID:

⊞ Internet Protocol Version 6, Src: fe80::a8aa:22ff:fe22:2222 (fe80::a8aa:22ff:fe22:2222), Dst: fe80::21d:72ff:fef8:9980
⊞ User Datagram Protocol, Src Port: dhcpv6-server (547), Dst Port: dhcpv6-client (546)
⊟ DHCPv6
    Message type: Reply (7)
    Transaction ID: 0x0fc14a
  ⊞ Server Identifier
  ⊞ Client Identifier
  ⊟ DNS recursive name server
      Option: DNS recursive name server (23)
      Length: 32
      Value: 20014860486000000000000000088882001486048600000...
      DNS server address: 2001:4860:4860::8888 (2001:4860:4860::8888)
      DNS server address: 2001:4860:4860::8844 (2001:4860:4860::8844)
```

Figure 4-18: DHCPv6 Reply with DNS Addresses

5 IPv6 Tunnelling

There are many ways in which IPv6 capable devices can communicate across an IPv4 infrastructure each of which could easily warrant a detailed tutorial of its own. This paper discusses only manual GRE and automatic 6to4 tunnels to introduce the concepts.

5.1 GRE Tunnel

The simplest way to transport IPv6 across an IPv4 network is to manually configure a GRE tunnel between IPv6 islands. Figure 5-1 below shows this extremely easy set-up in the lab.

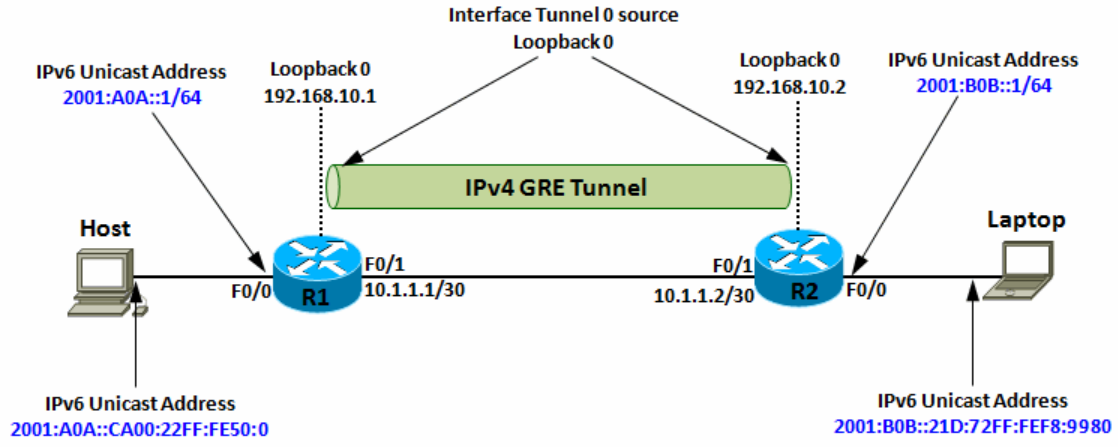

Figure 5-1: IPv4 GRE Tunnel to Transport IPv6

Table 5-1 below shows the configurations of R1 and R2.

Router R1	Router R2
!	!

```
! Interface Loopback 0 is the source address for interface
! Tunnel 0
!
interface Loopback0
 ip address 192.168.10.1 255.255.255.255
 ip router isis

!
! Interface Tunnel 0's IPv6 address on R1 is in the same
! sub-net as Tunnel 0's IPv6 address on R2. Tunnel 0's
! destination IPv4 is on R2 and is reachable because ISIS
! runs between both routers
!
interface Tunnel0
 no ip address
 ipv6 address 2001:C0C::1/64
 tunnel source Loopback0
 tunnel destination 192.168.10.2

!
! R1's LAN interface stays the same as before
!
interface FastEthernet0/0
 mac-address aaaa.1111.1111
 ip address 172.16.1.1 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 ipv6 address 2001:A0A::1/64
 ipv6 router isis

!
! R1's link to R2 has no IPv6 address only an IPv4 address
!
interface FastEthernet0/1
 ip address 10.1.1.1 255.255.255.252
 ip router isis
 duplex auto
 speed auto

!
! ISIS runs between the routers so that each can see
! each other's Loopback IPv4 address
!
router isis
 net 49.0001.1921.6810.0001.00
 is-type level-2-only
 metric-style wide

!
! All that is needed on R1 is an IPv6 static route to R2's LAN
! address via Tunnel 0

ipv6 route 2001:B0B::/32 Tunnel0
```

```
! Interface Loopback 0 is the source address for interface
! Tunnel 0
!
interface Loopback0
 ip address 192.168.10.2 255.255.255.255
 ip router isis

!
! Interface Tunnel 0's IPv6 address on R2 is in the same
! sub-net as Tunnel 0's IPv6 address on R1. Tunnel 0's
! destination IPv4 is on R1 and is reachable because ISIS
! runs between both routers
!
interface Tunnel0
 no ip address
 ipv6 address 2001:C0C::2/64
 tunnel source Loopback0
 tunnel destination 192.168.10.1

!
! R2's LAN interface stays the same as before. R2 continues
! to act as a DNS server for device Laptop to acquire its DNS
! IP addresses
!

interface FastEthernet0/0
 mac-address aaaa.2222.2222
 ip address 192.168.1.1 255.255.255.0
 ip router isis
 duplex auto
 speed auto
 ipv6 address 2001:B0B::1/64
 ipv6 nd other-config-flag
 ipv6 dhcp server DNS

!
! R2's link to R1 has no IPv6 address only an IPv4 address
!
interface FastEthernet0/1
 ip address 10.1.1.2 255.255.255.252
 ip router isis
 duplex auto
 speed auto

!
! ISIS runs between the routers so that each can see
! each other's Loopback IPv4 address
!
router isis
 net 49.0001.1921.6810.0002.00
 is-type level-2-only
 metric-style wide

!
! All that is needed on R2 is an IPv6 static route to R1's
! LAN address via Tunnel 0

ipv6 route 2001:A0A::/32 Tunnel0
```

Table 5-1: R1 and R2 Configuration for IPv6 over IPv4 GRE Tunnel

Laptop connected to R2 successfully pings Host across the IPv4 network.

```
C:\windows\system32>ping 2001:A0A::CA00:22FF:FE50:0
Pinging 2001:a0a::ca00:22ff:fe50:0 from 2001:b0b::21d:72ff:fef8:9980 with 32 byt
es of data:
Reply from 2001:a0a::ca00:22ff:fe50:0: time=248ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=78ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=46ms
Reply from 2001:a0a::ca00:22ff:fe50:0: time=93ms

Ping statistics for 2001:a0a::ca00:22ff:fe50:0:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
    Minimum = 46ms, Maximum = 248ms, Average = 116ms
```

Figure 5-2 below shows the contents of the echo request and reply packets as they traverse the link between R2 and R1. Ethernet transports IPv4, which transports the IPv6 GRE (protocol 0x86dd). The IPv4 source and destination addresses are the IPv4 loopback 0 addresses of R1 and R2, which are the tunnel 0 endpoints.

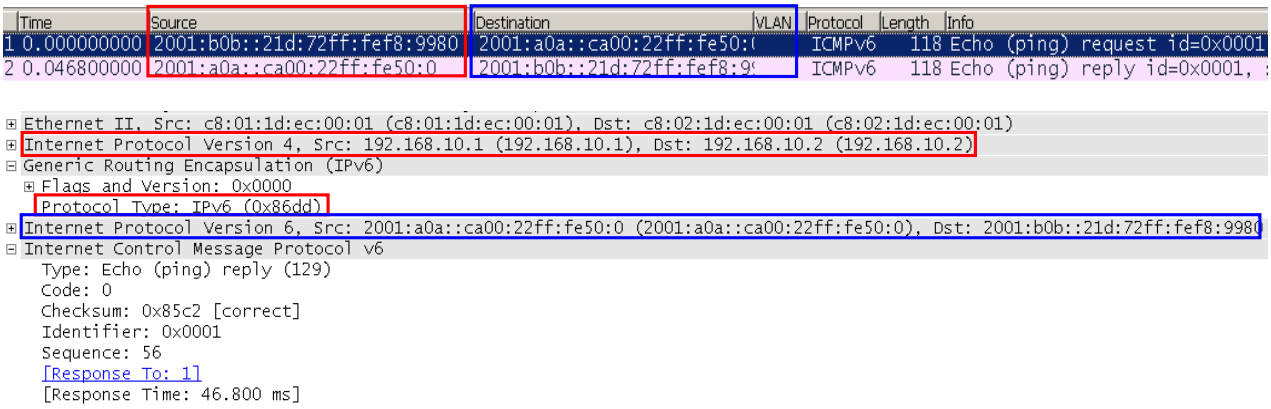

Figure 5-2: IPv6 Encapsulation in IPv4 GRE Packet Contents

The "ipconfig/all" command on Laptop confirms that it continues to acquire its IPv6 unicast address and DNS addresses successfully from R2.

```
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  . : home
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-1D-72-F8-99-80
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2001:b0b::21d:72ff:fef8:9980(Preferred)
   Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.66(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : 26 December 2013 03:50:44
   Lease Expires . . . . . . . . . . : 27 December 2013 15:50:49
   Default Gateway . . . . . . . . . : fe80::a8aa:22ff:fe22:2222%10
                                       192.168.1.254
   DHCP Server . . . . . . . . . . . : 192.168.1.254
   DHCPv6 IAID . . . . . . . . . . . : 218111346
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-11-51-0F-62-00-1D-72-F8-99-80
   DNS Servers . . . . . . . . . . . : 2001:4860:4860::8888
                                       2001:4860:4860::8844
                                       192.168.1.254
                                       192.168.1.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

5.2 Automatic 6to4

Whilst GRE tunnelling is a perfectly acceptable way to inter-connect IPv6 islands in a small network it does not scale for larger networks because multiple point-to-point tunnels must be configured manually between the IPv6 end-points. Automatic 6to4 tunnelling overcomes this disadvantage. With automatic 6to4 tunnelling, the special IPv6 unicast prefix "2002**::/16**" is used to automatically map IPv6 sub-nets to IPv4 prefixes.

Figure 5-3 below illustrates how 6to4 tunnelling could be used to inter-connect three sites each with their own IPv6 LAN sub-nets.

A /32 IPv4 prefix is assigned to the loopback 0 interface on each router - an IPv6 prefix is NOT assigned. A 6to4 tunnel interface on each router uses the loopback 0 interface as its source and no tunnel destination address is specified. The IPv4 network's routing protocol advertises the IPv4 loopback 0 prefixes to each other so every router knows how to reach every other router's loopback 0 interface. Only one 6to4 tunnel interface is configured per router.

The 32 bits of the IPv4 loopback prefix are mapped internally to the special IPv6 2002**::** /**16** prefix resulting in a 2002**::XXXX:XXXX::** /**48** prefix. Each site is allocated one or more IPv6 /64 LAN sub-nets that is a sub-net of (i.e. sits behind) the /48 sub-net of the loopback 0 interface.

When a packet with IPv6 source 2002:C0A8:0101:1::1/64 and destination 2002:AC10:0101:1::1/64 leaves the LAN attached to R2, the router knows it is a 6to4 packet because the destination uses the special 2002::/16 prefix. R2 encapsulates the packet in IPv4 with a destination IPv4 address of 172.16.1.1 and forwards the packet over its 6to4 tunnel interface towards R1. When R1 receives the packet, it extracts the IPv6 packet and forwards it to IPv6 sub-net 2002:AC10:0101:1::/64.

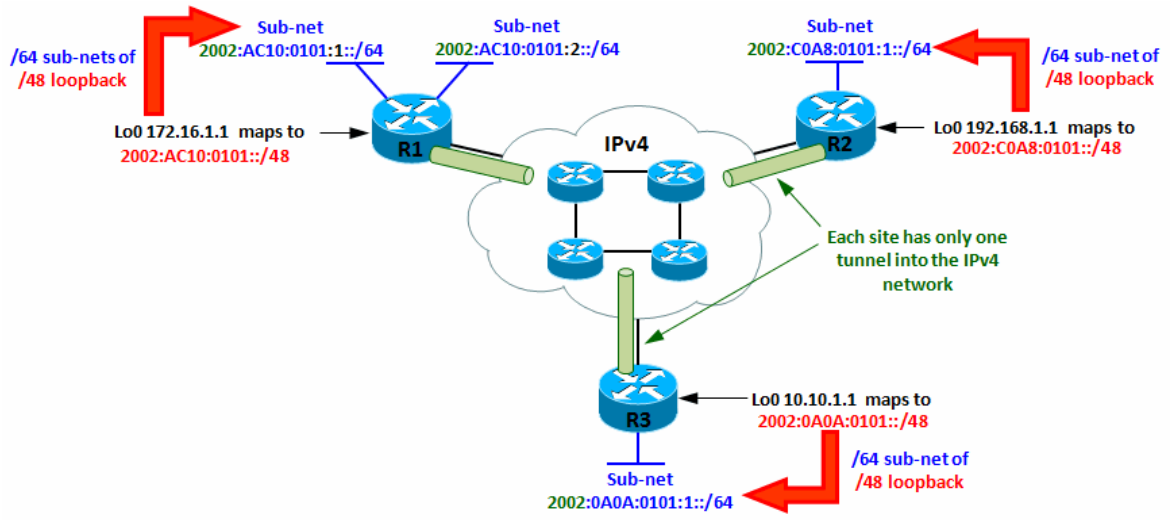

Figure 5-3: Inter-Connecting IPv6 Sites Using 6to4 Tunnelling

Figure 5-4 below shows how R1 loopback 172.16.1.1/32, R2 loopback 192.168.1.1/32 and R3 loopback 10.10.1.1/32 are mapped to the special 6to4 2002**::XXXX:XXXX::** /**48** prefixes and the 16 bits that are used for the LAN sub-nets that sit behind the /48 prefixes.

(16 bits) /16		(32 bits) /48								(16 bits) /64	
R1 6to4		172		16		1		1		IPv6 Sub-Net (16 bits)	
20	02	AC		10		01		01		XX	XX
0010 0000	0000 0010	1010	1100	0001	0000	0000	0001	0000	0001	0000	0001
0010 0000	0000 0010	1010	1100	0001	0000	0000	0001	0000	0001	0000	0010
R2 6to4		192		168		1		1		IPv6 Sub-Net (16 bits)	
20	02	C0		A8		01		01		XX	XX
0010 0000	0000 0010	1100	0000	1010	1000	0000	0001	0000	0001	0000	0001
R3 6to4		10		10		1		1		IPv6 Sub-Net (16 bits)	
20	02	0A		0A		01		01		XX	XX
0010 0000	0000 0010	0000	1010	0000	1010	0000	0001	0000	0001	0000	0001

Figure 5-4: IPv4 Mapping to IPv6 6to4 address

Figure 5-5 below shows the lab set up for 6to4 tunnelling between R1 and R2. Only IPv4 is configured between the two routers with ISIS advertising each loopback 0 IPv4 interface address.

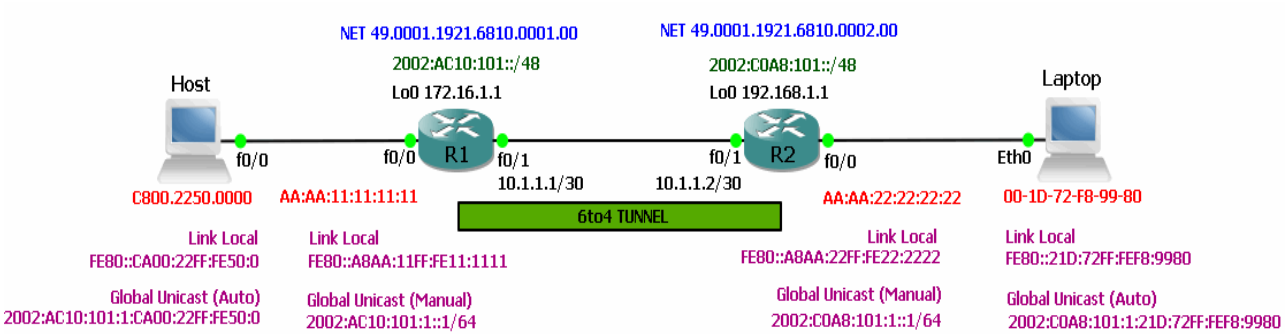

Figure 5-5: 6to4 Tunnel between Lab Routers R1 and R2

Laptop pings Host's 2002:AC10:101:**1**:CA00:22FF:FE50:0 IPv6 address on sub-net "**:1**" successfully.

```
C:\windows\system32>ping  2002:AC10:101:1:CA00:22FF:FE50:0
Pinging 2002:ac10:101:1:ca00:22ff:fe50:0 from 2002:c0a8:101:1:21d:72ff:fef8:9980
 with 32 bytes of data:
Reply from 2002:ac10:101:1:ca00:22ff:fe50:0: time=180ms
Reply from 2002:ac10:101:1:ca00:22ff:fe50:0: time=64ms
Reply from 2002:ac10:101:1:ca00:22ff:fe50:0: time=47ms
Reply from 2002:ac10:101:1:ca00:22ff:fe50:0: time=45ms

Ping statistics for 2002:ac10:101:1:ca00:22ff:fe50:0:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
Approximate round trip times in milli-seconds:
    Minimum = 45ms, Maximum = 180ms, Average = 84ms
```

The output below on Laptop confirms that it has acquired its new 6to4 "2002**::**" prefix from R2 using router solicitations. Laptop continues to obtain its IPv6 DNS addresses from R2 using DHCPv6.

```
Ethernet adapter Local Area Connection:
   Connection-specific DNS Suffix  . : home
   Description . . . . . . . . . . . : Broadcom NetLink (TM) Gigabit Ethernet
   Physical Address. . . . . . . . . : 00-1D-72-F8-99-80
   DHCP Enabled. . . . . . . . . . . : Yes
   Autoconfiguration Enabled . . . . : Yes
   IPv6 Address. . . . . . . . . . . : 2002:c0a8:101:1:21d:72ff:fef8:9980(Preferred)
   Link-local IPv6 Address . . . . . : fe80::21d:72ff:fef8:9980%10(Preferred)
   IPv4 Address. . . . . . . . . . . : 192.168.1.66(Preferred)
   Subnet Mask . . . . . . . . . . . : 255.255.255.0
   Lease Obtained. . . . . . . . . . : 26 December 2013 03:50:44
   Lease Expires . . . . . . . . . . : 27 December 2013 15:50:48
   Default Gateway . . . . . . . . . : fe80::a8aa:22ff:fe22:2222%10
                                       192.168.1.254
   DHCP Server . . . . . . . . . . . : 192.168.1.254
   DHCPv6 IAID . . . . . . . . . . . : 218111346
   DHCPv6 Client DUID. . . . . . . . : 00-01-00-01-11-51-0F-62-00-1D-72-F8-99-80
   DNS Servers . . . . . . . . . . . : 2001:4860:4860::8888
                                       2001:4860:4860::8844
                                       192.168.1.254
                                       192.168.1.254
   NetBIOS over Tcpip. . . . . . . . : Enabled
```

The output below on R2 confirms that packets are traversing 6to4 tunnel 0.

```
R2#show ipv6 tunnel
Tun Route  LastInp              Packets Description
  0   -    00:01:42                 52
```

Table 5-2 below shows the configurations for R1 and R2.

Router R1	Router R2
! **! Interface Loopback 0 is the source address for interface** **! Tunnel 0** **!** interface Loopback0 ip address 172.16.1.1 255.255.255.255 ip router isis **!** **! Interface Tunnel 0 requires only a tunnel source IPv4** **! address and for IPv6 to be enabled. Tunnel mode is** **! set to ipv6ip 6to4. No destination IPv4 is required** **!** interface Tunnel0 no ip address no ip redirects **ipv6 enable** tunnel source Loopback0 **tunnel mode ipv6ip 6to4** **!** **! LAN interface f0/0 is a /64 sub-net of the Loopback 0** **! interface** **!** interface FastEthernet0/0 mac-address aaaa.1111.1111 no ip address duplex auto speed auto **ipv6 address 2002:AC10:101:1::1/64** **!** **! LAN interface f0/1 only transports IPv4** **!** interface FastEthernet0/1 ip address 10.1.1.1 255.255.255.252 ip router isis duplex auto speed auto ! router isis net 49.0001.1921.6810.0001.00 is-type level-2-only metric-style wide **!** **! All that is needed on R2 is an IPv6 static route to direct** **! 6to4 2002:: prefixes via Tunnel 0** **ipv6 route 2002::/16 Tunnel0**	**!** **! Interface Loopback 0 has the source address for interface** **! Tunnel 0** **!** interface Loopback0 ip address 192.168.1.1 255.255.255.255 ip router isis **!** **! Interface Tunnel 0 requires only a tunnel source IPv4** **! address and for IPv6 to be enabled. Tunnel mode is** **! set to ipv6ip 6to4. No destination IPv4 is required** **!** interface Tunnel0 no ip address no ip redirects **ipv6 enable** tunnel source Loopback0 **tunnel mode ipv6ip 6to4** **!** **! LAN interface f0/0 is a /64 sub-net of the Loopback 0** **! interface** **!** interface FastEthernet0/0 mac-address aaaa.2222.2222 no ip address duplex auto speed auto **ipv6 address 2002:C0A8:101:1::1/64** ipv6 nd other-config-flag ipv6 dhcp server DNS **!** **! LAN interface f0/1 only transports IPv4** **!** interface FastEthernet0/1 ip address 10.1.1.2 255.255.255.252 ip router isis duplex auto speed auto ! router isis net 49.0001.1921.6810.0002.00 is-type level-2-only metric-style wide**!** **!** **! All that is needed on R2 is an IPv6 static route to direct** **! 6to4 2002:: prefixes via Tunnel 0** **ipv6 route 2002::/16 Tunnel0**

Table 5-2: 6to4 Configurations for R1 and R2