

# SHARKFEST '12

Wireshark Developer and User Conference

## ICMPv6

Nalini Elkins

CEO

Inside Products, Inc.

[Nalini.elkins@insidestack.com](mailto:Nalini.elkins@insidestack.com)

# Agenda

- Changes ICMPv4 / ICMPv6
- New ICMPv6 functions
  - Router discovery,
  - Prefix discovery,
  - Parameter discovery,
  - Address resolution,
  - Neighbor unreachability,
  - Duplicate Address Detection,
  - Redirect

# Why ICMP?

- IP uses ICMP to convey error information.
- Like what?

# Why ICMP?

- IP uses ICMP to convey error information.
- Like what?
  - Host unreachable
  - Port unreachable
  - Firewall stopped the packet
  - There is a better way to get from here to there

# ICMPv4 Messages

- Some ICMPv4 packets are 'functional'
- Like what?

# ICMPv4 Messages

- Some ICMPv4 packets are ‘functional’
- Like what?
  - Ping
  - Redirect
  - Sometimes called ‘informational’

# ICMP Header

Octet	Len	Name	Notes
0	1	ICMP Type	ICMP Message Type 0 = Echo Reply(PING) 3 = Destination Unreachable 4 = Source Quench 5 = Redirect (Route Change) 8 = Echo Request(Ping) 11 = Time Exceeded 12 = Parameter Problem 13 = Timestamp Request 14 = Timestamp Reply 17 = Address Mask Request 18 = Address Mask Reply
1	1	Code	Code values are message specific.
2-3	2	Checksum	-

- ICMP messages are transferred through the network as the data portion of an IP datagram.
- This means that ICMP messages themselves can be lost.
- To avoid generation of error messages about error messages, new error messages about ICMP errors are not generated.
- Each ICMP message has a slightly different format but the first 4 bytes are ALWAYS the same.

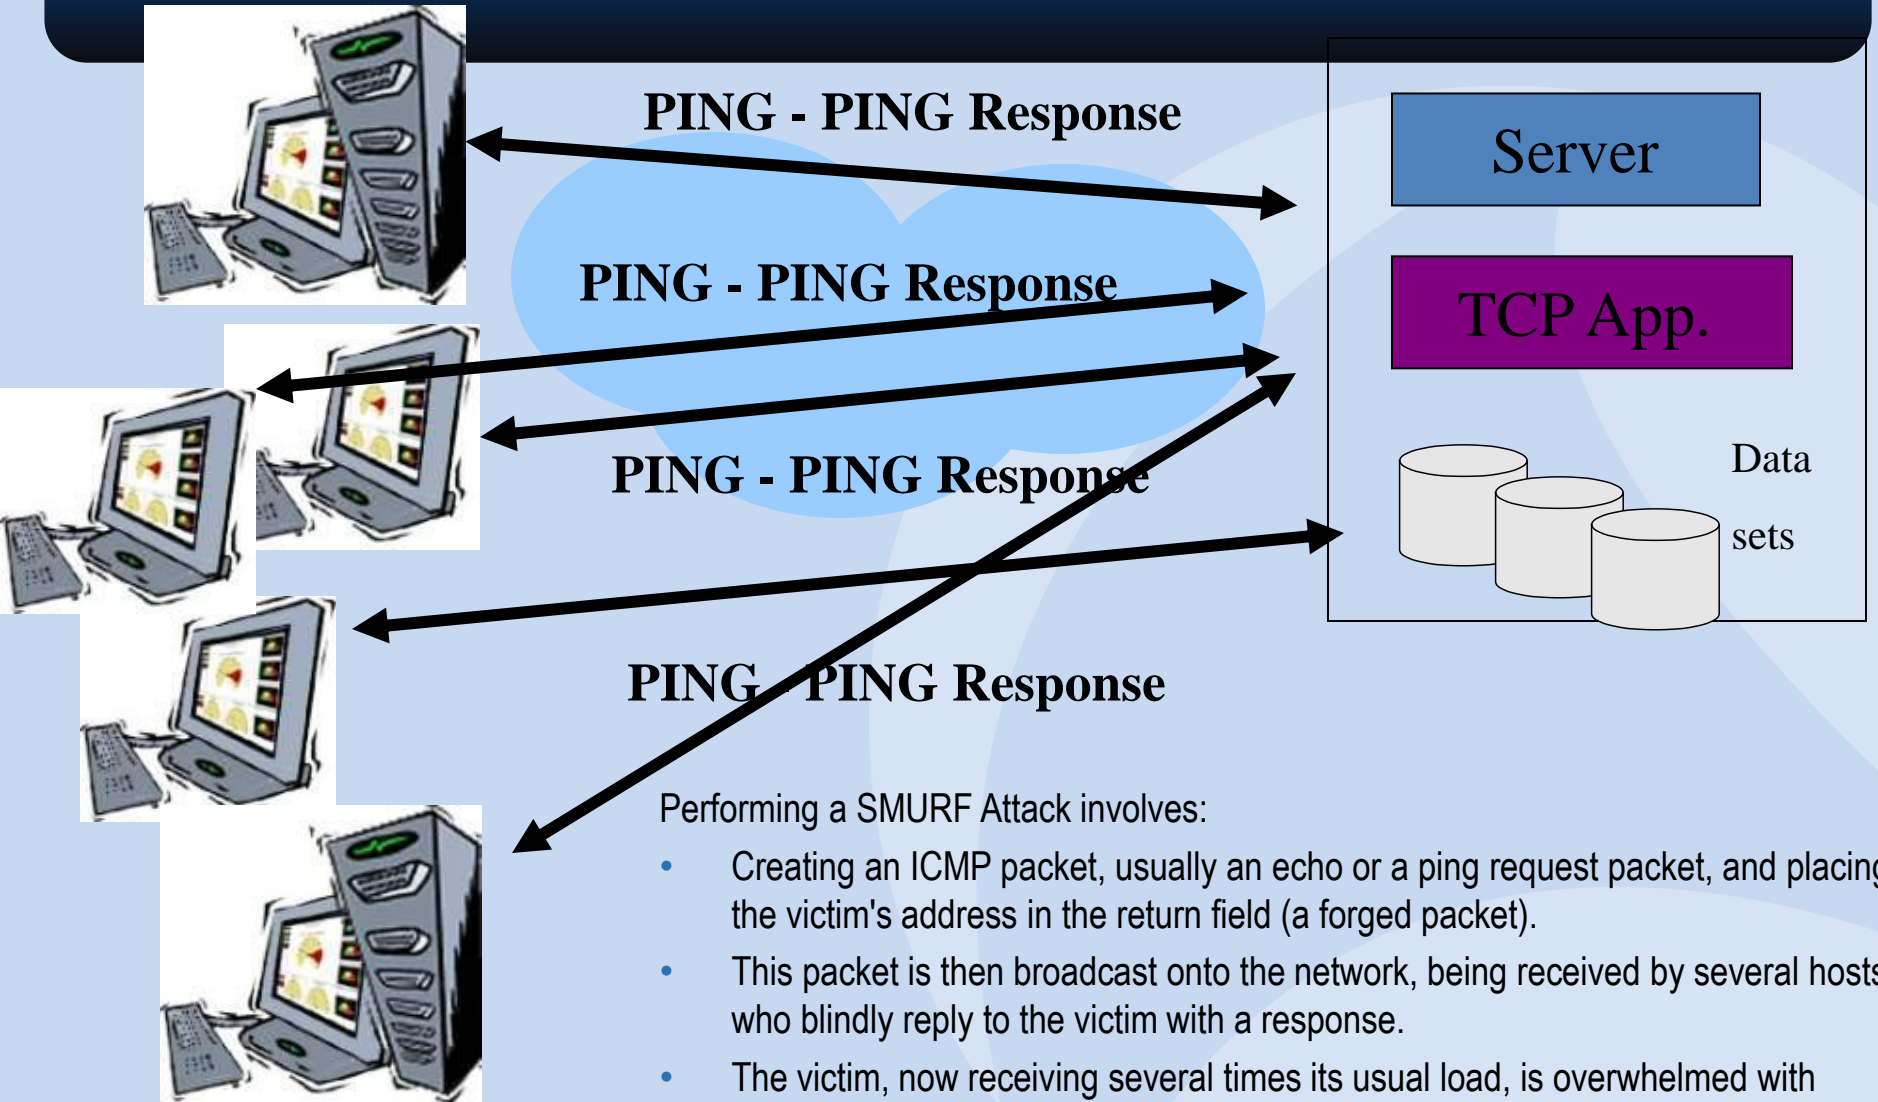
# Hack Start!



- Hack – SMURF : ICMP Protocol
- A SMURF Attack is a denial-of-service network attack (DoS) that is directed towards some pre-determined target, usually a server.
- Any server that is plugged into a network and can receive IP packets is vulnerable.
- These attacks come very quickly and present themselves as very hard to trace.



# ICMP SMURF



Performing a SMURF Attack involves:

- Creating an ICMP packet, usually an echo or a ping request packet, and placing the victim's address in the return field (a forged packet).
- This packet is then broadcast onto the network, being received by several hosts who blindly reply to the victim with a response.
- The victim, now receiving several times its usual load, is overwhelmed with response packets.

# Reflector Attacks

- Reflectors: All Web or DNS servers, and routers are potential reflectors, since they will return
  - SYN acks or RSTs in response to SYN or other TCP packets;
  - Query replies in response to query requests; or
  - ICMP Time Exceeded or
  - Host Unreachable in response to particular IP packets.
- By spoofing IP addresses from slaves — a massive distributed Denial of Service (dDoS) attack can be arranged.

# What has changed?

## ICMPv4 Messages

-----	-----
0	Echo Reply
3	Destination Unreachable
4	Source Quench
5	Redirect Message
8	Echo Request
11	Time Exceeded
12	Parameter Problem
13	Timestamp Request
14	Timestamp Reply
17	Address Mask Request
18	Address Mask Reply

# What has changed?

## ICMPv4 Messages

----      -----

- 0    Echo Reply
- 3    Destination Unreachable
- ~~4    Source Quench~~
- 5    Redirect Message
- 8    Echo Request
- 11   Time Exceeded
- 12   Parameter Problem
- ~~13   Timestamp Request~~
- ~~14   Timestamp Reply~~
- ~~17   Address Mask Request~~
- ~~18   Address Mask Reply~~

# ICMPv6 Error Messages

Type	Name	Reference
1	Destination Unreachable	[RFC2463]
2	Packet Too Big	[RFC2463]
3	Time Exceeded	[RFC2463]
4	Parameter Problem	[RFC2463]

Error messages have message types from 0 to 127.

# ICMPv4 Error – Info Ratio

- Error messages : 90%
- Informational : 10%

# ICMPv6 Error – Info Ratio

- Error messages : 20%
- Informational : 80%

# ICMPv6 Info Messages

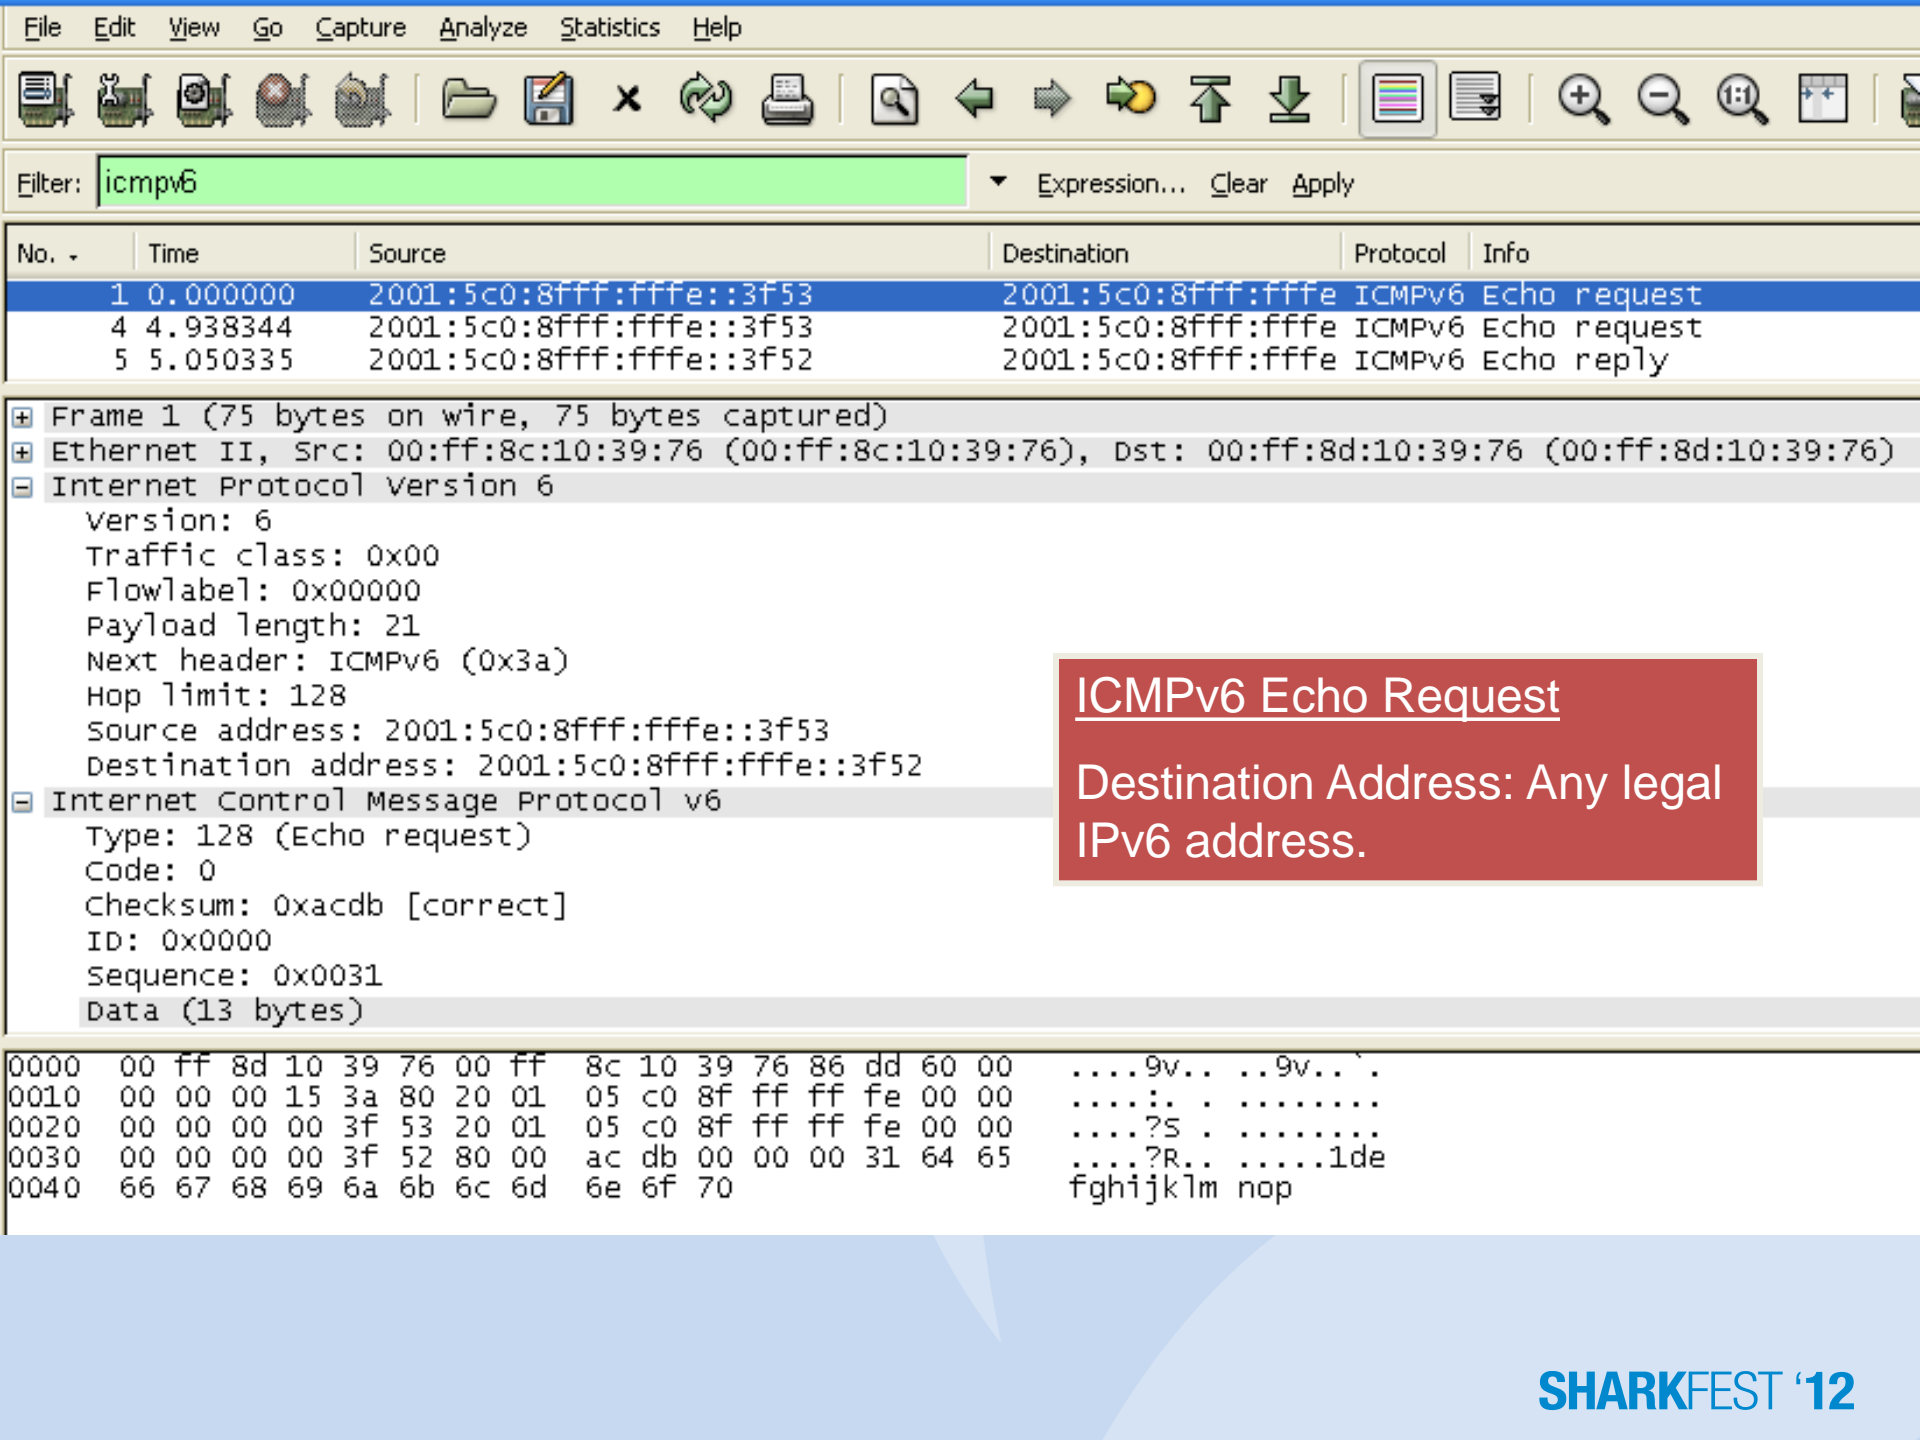
- Why????
- Informational : 80%
  - ARP gone!
  - Replaced by Neighbor discovery / Router discovery, Multicast Listener Discovery
  - Mobile IP



# ICMPv6 Informational Messages

Type	Name
128	Echo Request
129	Echo Reply
130	Multicast Listener Query
131	Multicast Listener Report
132	Multicast Listener Done
133	Router Solicitation
134	Router Advertisement
135	Neighbor Solicitation
136	Neighbor Advertisement
137	Redirect Message
138	Router Renumbering
139	ICMP Node Info. Query
140	ICMP Node Info. Response
141	Inverse Neighbor Discovery Solicitation Message

Type	Name
142	Inverse Neighbor Discovery Advertisement Message
143	Version 2 Multicast Listener Report
144	Home Agent Address Discovery Request Message
145	Home Agent Address Discovery Reply Message
146	Mobile Prefix Solicitation
147	Mobile Prefix Advertisement
148	Certification Path Solicitation
149	Certification Path Advertisement
150	Experimental mobility protocols
151	Multicast Router Advertisement
152	Multicast Router Solicitation
153	Multicast Router Termination



Filter: icmpv6 Expression... Clear Apply

No.	Time	Source	Destination	Protocol	Info
1	0.000000	2001:5c0:8fff:fffe::3f53	2001:5c0:8fff:fffe::3f52	ICMPv6	Echo request
4	4.938344	2001:5c0:8fff:fffe::3f53	2001:5c0:8fff:fffe::3f52	ICMPv6	Echo request
5	5.050335	2001:5c0:8fff:fffe::3f52	2001:5c0:8fff:fffe::3f52	ICMPv6	Echo reply

```

+ Frame 1 (75 bytes on wire, 75 bytes captured)
+ Ethernet II, Src: 00:ff:8c:10:39:76 (00:ff:8c:10:39:76), Dst: 00:ff:8d:10:39:76 (00:ff:8d:10:39:76)
- Internet Protocol Version 6
  Version: 6
  Traffic class: 0x00
  Flowlabel: 0x00000
  Payload length: 21
  Next header: ICMPv6 (0x3a)
  Hop limit: 128
  Source address: 2001:5c0:8fff:fffe::3f53
  Destination address: 2001:5c0:8fff:fffe::3f52
- Internet Control Message Protocol v6
  Type: 128 (Echo request)
  Code: 0
  Checksum: 0xacdb [correct]
  ID: 0x0000
  Sequence: 0x0031
  Data (13 bytes)

```

ICMPv6 Echo Request  
 Destination Address: Any legal IPv6 address.

```

0000  00 ff 8d 10 39 76 00 ff 8c 10 39 76 86 dd 60 00  ....9v.. ..9v...
0010  00 00 00 15 3a 80 20 01 05 c0 8f ff ff fe 00 00  ....:.. .
0020  00 00 00 00 3f 53 20 01 05 c0 8f ff ff fe 00 00  ....?S .
0030  00 00 00 00 3f 52 80 00 ac db 00 00 00 31 64 65  ....?R.. ....1de
0040  66 67 68 69 6a 6b 6c 6d 6e 6f 70                fghijklm nop

```

No. ↓	Time	Source	Destination	Protocol	Info
1	0.000000	2001:5c0:8fff:fffe::3f53	2001:5c0:8fff:fffe	ICMPv6	Echo request
4	4.938344	2001:5c0:8fff:fffe::3f53	2001:5c0:8fff:fffe	ICMPv6	Echo request
5	5.050335	2001:5c0:8fff:fffe::3f52	2001:5c0:8fff:fffe	ICMPv6	Echo reply

⊕ Frame 5 (75 bytes on wire, 75 bytes captured)

⊕ Ethernet II, Src: 00:ff:8d:10:39:76 (00:ff:8d:10:39:76), Dst: 00:ff:8c:10:39:76 (00:ff:8c:10:39:76)

⊖ Internet Protocol Version 6

Version: 6

Traffic class: 0x00

Flowlabel: 0x00000

Payload length: 21

Next header: ICMPv6 (0x3a)

Hop limit: 64

Source address: 2001:5c0:8fff:fffe::3f52

Destination address: 2001:5c0:8fff:fffe::3f53

⊖ Internet Control Message Protocol v6

Type: 129 (Echo reply)

Code: 0

Checksum: 0xabda [correct]

ID: 0x0000

Sequence: 0x0032

Data (13 bytes)

## ICMPv6 Echo Reply

- An Echo Reply SHOULD be sent in response to an Echo Request message sent to an IPv6 multicast address.

- The source address of the reply MUST be a unicast address belonging to the interface on which the multicast Echo Request message was received.

```

0000  00 ff 8c 10 39 76 00 ff 8d 10 39 76 86 dd 60 00  ....9v.. ..9v..`
0010  00 00 00 15 3a 40 20 01 05 c0 8f ff ff fe 00 00  ....:@ . ....
0020  00 00 00 00 3f 52 20 01 05 c0 8f ff ff fe 00 00  ....?R . ....
0030  00 00 00 00 3f 53 81 00 ab da 00 00 00 32 64 65  ....?S.. .....2de
0040  66 67 68 69 6a 6b 6c 6d 6e 6f 70                fghijklm nop

```

# Ping to Multicast Addresses

Pinging ff02::1 with 32 bytes of data:

Reply from ff02::1: time<1ms

Reply from ff02::1: time<1ms

Reply from ff02::1: time<1ms

Reply from ff02::1: time<1ms

Ping statistics for ff02::1:

Packets: Sent = 4, Received = 4,

Lost = 0 (0% loss),

Approximate round trip times in  
milliseconds:

Minimum = 0ms, Maximum = 0ms,

Average = 0ms

Did a Ping for Multicast address:

FF02:0:0:0:0:0:0:1 All Nodes Address

Pinging ff02::2 with 32 bytes of  
data:

Request timed out.

Request timed out.

Request timed out.

Request timed out.

Ping statistics for ff02::2:

Packets: Sent = 4, Received = 0,  
Lost = 4 (100% loss),

Did a Ping for Multicast address:

FF02:0:0:0:0:0:0:2 All Routers Address

Does this mean my router is down?

# Ping to www.kame.net

```
Pinging www.kame.net
```

```
2001:200:0:8002:203:47ff:fea5:3085] with 32 bytes of data:
```

```
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=227ms
```

```
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=228ms
```

```
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=250ms
```

```
Reply from 2001:200:0:8002:203:47ff:fea5:3085: time=349ms
```

```
Ping statistics for 2001:200:0:8002:203:47ff:fea5:3085:
```

```
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
```

```
Approximate round trip times in milli-seconds:
```

```
    Minimum = 227ms, Maximum = 349ms, Average = 263ms
```

The router stack SHOULD implement an echo reply but there is no MUST in the RFC! Do not have to implement echo reply for multicast address.

# IPv6 Destination Unreachable

Code	Meaning	Description
0	<i>No Route To Destination</i>	The datagram was not delivered because it could not be routed to the destination. Since this means the datagram could not be sent to the destination device's local network, this is basically equivalent to the "Network Unreachable" message subtype in ICMPv4.
1	<i>Communication With Destination Administratively Prohibited</i>	The datagram could not be forwarded due to filtering that blocks the message based on its contents. Equivalent to the message subtype with the same name (and Code value 13) in ICMPv4.
3	<i>Address Unreachable</i>	There was a problem attempting to deliver the datagram to the host specified in the destination address. This code is equivalent to the ICMPv4 "Host Unreachable" code and usually means the destination address was bad or there was a problem with resolving it into a layer two address.
4	<i>Port Unreachable</i>	The destination port specified in the UDP or TCP header was invalid or does not exist on the destination host.

## ICMPv4 Dest Unreach Subcodes

0:Network Unreachable

1: Host Unreachable

2: Protocol Unreachable

4:Fragmentation Needed and DF Set

5:Source Route Failed

6: Destination Network Unknown

7:Destination Host Unknown

8:Source Host Isolated

9:Communication with Destination Network is Administratively Prohibited

10:Communication with Destination Host is Administratively Prohibited

11:Destination Network Unreachable for Type of Service

12:Destination Host Unreachable for Type of Service

13:Communication Administratively Prohibited

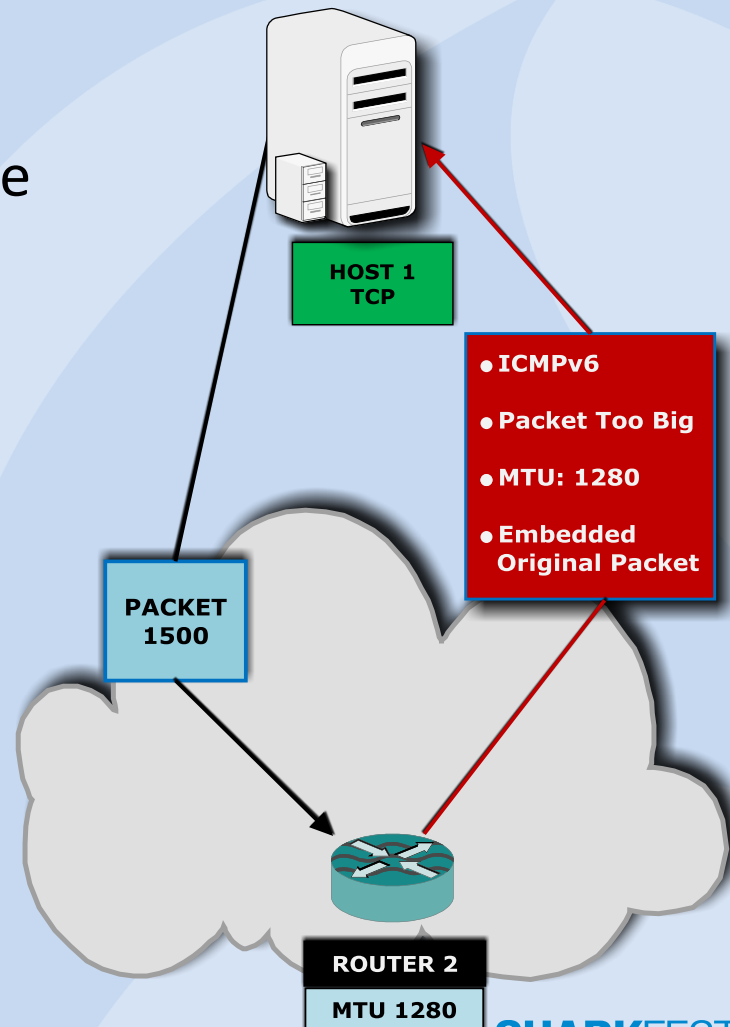
14:Host Precedence Violation

15Precedence Cutoff In Effect

No.	Time	Source	Destination	Protocol	Info
136	59.819289	3ffe:501:4819::42	3ffe:507:0:1:200:8	DNS	Standard query response AAAA 3ffe:
137	59.820360	3ffe:507:0:1:200:8	3ffe:501:4819::42	ICMPv6	Unreachable (Port unreachable)
+ Frame 137 (300 bytes on wire, 300 bytes captured)					
Ethernet II, Src: Megahert_05:80:da (00:00:86:05:80:da), Dst: 3com_07:69:ea (00:60:97:07:69:ea)					
Destination: 3com_07:69:ea (00:60:97:07:69:ea)					
Source: Megahert_05:80:da (00:00:86:05:80:da)					
Type: IPv6 (0x86dd)					
Internet Protocol Version 6					
Version: 6					
Traffic class: 0x00					
Flowlabel: 0x00000					
Payload length: 246					
Next header: ICMPv6 (0x3a)					
Hop limit: 64					
Source address: 3ffe:507:0:1:200:86ff:fe05:80da					
Destination address: 3ffe:501:4819::42					
Internet Control Message Protocol v6					
Type: 1 (Unreachable)					
Code: 4 (Port unreachable) ←					
Checksum: 0xb1b5 [correct]					
Internet Protocol Version 6					
Version: 6					
Traffic class: 0x00					
Flowlabel: 0x00000					
Payload length: 198					
Next header: UDP (0x11)					
Hop limit: 230					
Source address: 3ffe:501:4819::42					
Destination address: 3ffe:507:0:1:200:86ff:fe05:80da					
User Datagram Protocol, Src Port: domain (53), Dst Port: 2410 (2410)					
Source port: domain (53)					
Destination port: 2410 (2410)					
Length: 198					
Checksum: 0x1e36 [correct]					
Domain Name System (response)					
Transaction ID: 0x5c74					
Flags: 0x8580 (Standard query response, No error)					

# ICMPv6 Packet Too Big

- In IPv6, routers are not allowed to fragment datagrams that are too large to send over a physical link are connected.
- Packet is dropped, and an ICMPv6 *Packet Too Big* message sent. (minimum IPv6 MTU 1280 bytes)
- Used in Path MTU Discovery






# Now, the more complicated ones!

- Neighbor discovery,
- Router discovery,
- Multicast Listener Discovery

# Stateless Autoconfiguration

- Stateless autoconfiguration allows a node to be configured without any configuration server.
- How? A node configures its own globally routable addresses in cooperation with a local IPv6 router.
- The address combines the 48- or 64-bit MAC address of the adapter with network prefixes that are learned from the neighboring router.
- In the case of multi-homed devices, autoconfiguration is performed for each interface separately.
- Stateless autoconfiguration uses the Neighbor Discovery protocol.

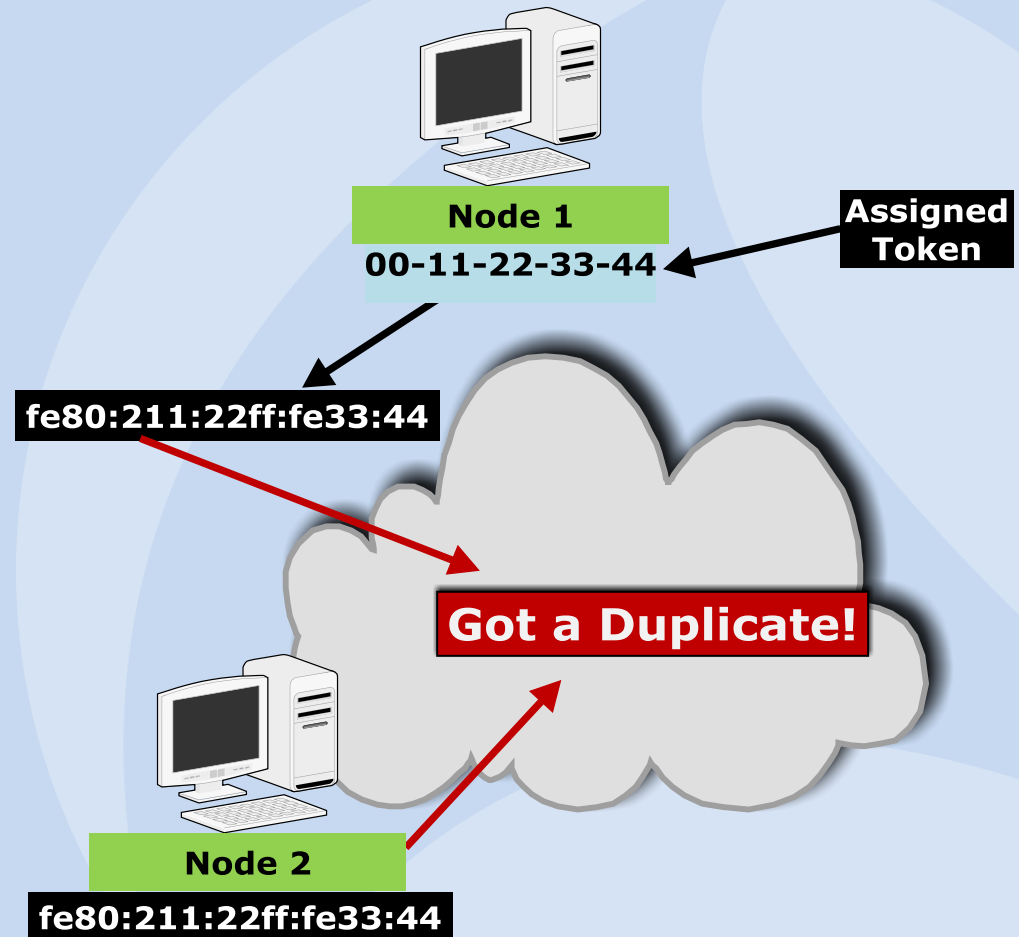


Example on Windows PC: result of IPConfig

```
Ethernet adapter Local Area Connection:  
Description : Realtek Family Fast Ethernet NIC  
Physical Address : 00-11-D8-39-29-2B  
Autoconfiguration Enabled . : Yes  
IP Address : fe80::211:d8ff:fe39:292b%4
```

# Stateless Autoconfiguration Steps 1 - 2

- **Link-Local Address Generation:**  
The device generates a link-local address.
- **Link-Local Address Uniqueness Test:**
  - Is someone using my address?
  - Sends *Neighbor Solicitation* message
  - Listens for a *Neighbor Advertisement*



# Stateless Autoconfiguration Steps 3 - 4

- **Link-Local Address Assignment:**

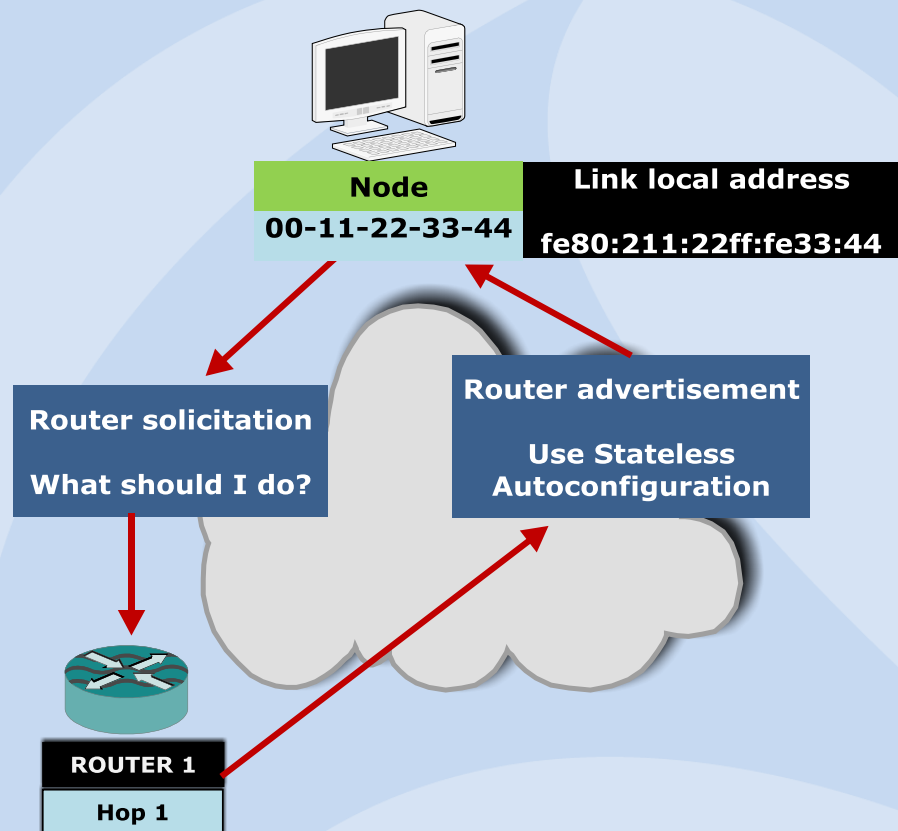
- Can be used for communication on the local network, but not on internet or intranet.

- **Router Contact:**

- Asks local router what to do
- Sends *Router Solicitation*
- Listens for *Router Advertisement*

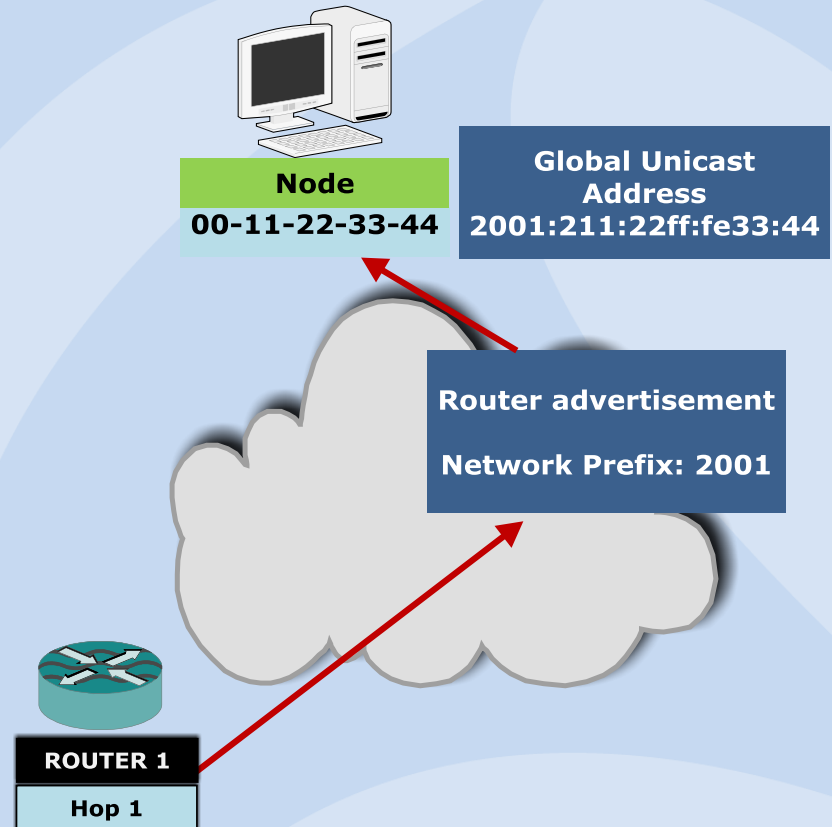
- **Router Direction:**

- Are we stateful / stateless
- What prefix do we use?

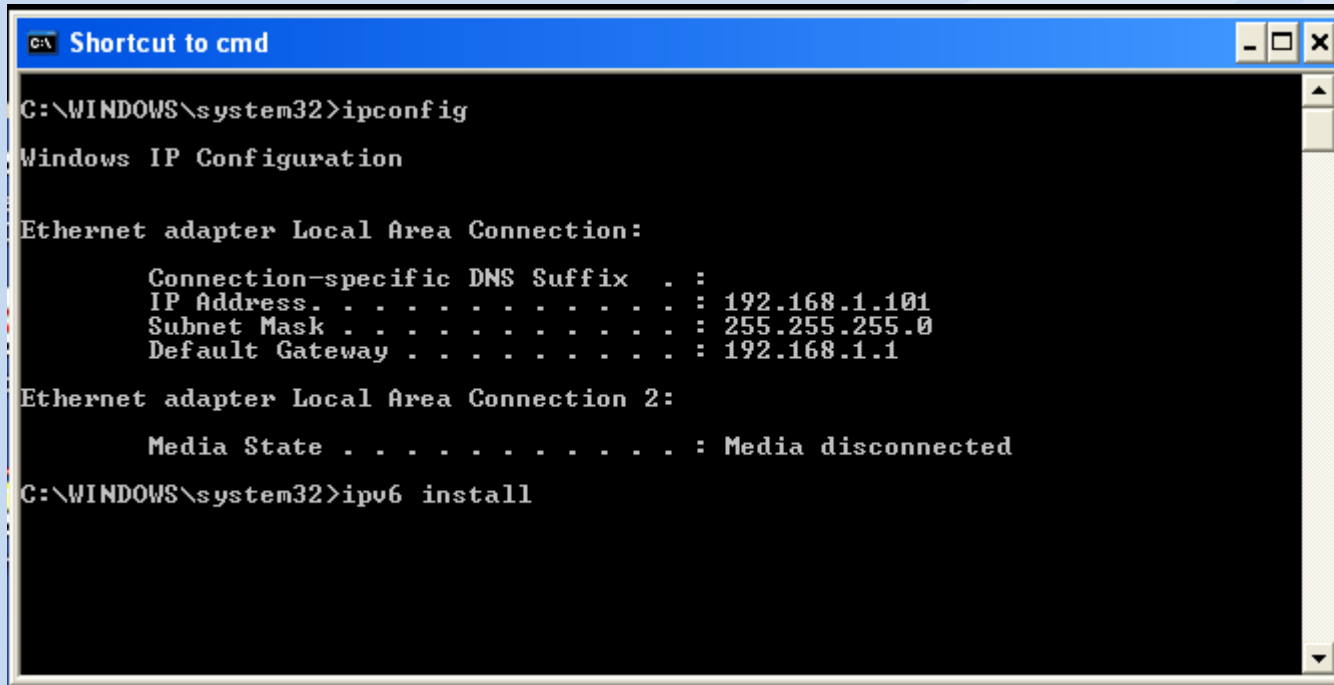


# Stateless Autoconfiguration Step 5

- **Global Address Configuration:**
  - If using stateless autoconfiguration , form global unicast address combining network prefix and MAC address (IID).
- Advantages:
  - Low administrative costs
- Disadvantages
  - Low administrative costs



# Stateless Autoconfig on Windows



```
C:\> Shortcut to cmd
C:\WINDOWS\system32>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : 192.168.1.101
    Subnet Mask . . . . .             : 255.255.255.0
    Default Gateway . . . . .         : 192.168.1.1

Ethernet adapter Local Area Connection 2:

    Media State . . . . .             : Media disconnected

C:\WINDOWS\system32>ipv6 install
```

- To see stateless autoconfiguration at work, start with a Windows PC with no IPv6 enabled.
- Look at the IPconfig above.
- You see only IPv4 connections
- Let's install IPv6.

# After IPv6 Installed Successfully

```
C:\> Shortcut to cmd
C:\WINDOWS\system32>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : 192.168.1.101
    Subnet Mask . . . . .             : 255.255.255.0
    IP Address. . . . .               : fe80::211:d8ff:fe39:292b%5
    Default Gateway . . . . .         : 192.168.1.1

Ethernet adapter Local Area Connection 2:

    Media State . . . . .             : Media disconnected

Tunnel adapter Teredo Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : fe80::5445:5245:444f%4
    Default Gateway . . . . .         :

Tunnel adapter Automatic Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : fe80::5efe:192.168.1.101%2
    Default Gateway . . . . .         :

C:\WINDOWS\system32>
```

- Notice what addresses are assigned.
- Will we be able to go out over the internet?
- What do you think is the MAC address?
- Why did this happen?

# IPConfig with Global Unicast Addresses

```
c:\ Shortcut to cmd
C:\WINDOWS\system32>ipconfig

Windows IP Configuration

Ethernet adapter Local Area Connection:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : 192.168.1.100
    Subnet Mask . . . . .             : 255.255.255.0
    IP Address. . . . .               : 2001:4840:ffff:c012:5d8c:c7f:6d5:1047
    IP Address. . . . .               : 2001:4840:ffff:c012:211:d8ff:fe39:292b
    IP Address. . . . .               : fe80::211:d8ff:fe39:292b%5
    Default Gateway . . . . .         : 192.168.1.1
                                         fe80::214:bfff:feba:45f9%5

Ethernet adapter Local Area Connection 2:

    Media State . . . . .             : Media disconnected

Tunnel adapter Teredo Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : fe80::5445:5245:444f%4
    Default Gateway . . . . .         :

Tunnel adapter Automatic Tunneling Pseudo-Interface:

    Connection-specific DNS Suffix  . : 
    IP Address. . . . .               : fe80::5efe:192.168.1.100%2
    Default Gateway . . . . .         :

C:\WINDOWS\system32>_
```

- Will we be able to go out over the internet?
- Why did this happen?
- Notice default IPv6 gateway.



start | Shortcut to cmd | untitled - Paint | SecurityAndIPv6Blue... | Microsoft PowerPoint ... | (Untitled) - Ethereal

File Edit View Go Capture Analyze Statistics Help

Filter: icmpv6 Expression... Clear Apply


No. -	Time	Source	Destination	Protocol	Info
23	13.642801	::	ff02::1:ff39:292b	ICMPv6	Multicast listener report
24	13.642826	::	ff02::2	ICMPv6	Router solicitation
25	13.642847	::	ff02::1:ff39:292b	ICMPv6	Neighbor solicitation
31	17.642731	fe80::211:d8ff:fe39:292b	ff02::2	ICMPv6	Router solicitation
46	21.642662	fe80::211:d8ff:fe39:292b	ff02::2	ICMPv6	Router solicitation
47	22.642644	fe80::211:d8ff:fe39:292b	ff02::1:ff39:292b	ICMPv6	Multicast listener report

Frame 25 (78 bytes on wire, 78 bytes captured)

- Ethernet II, Src: AsustekC\_39:29:2b (00:11:d8:39:29:2b), Dst: IPv6-Neighbor-Discovery\_ff:39:29:2b (33:33:ff:39:29:2b)
  - Destination: IPv6-Neighbor-Discovery\_ff:39:29:2b (33:33:ff:39:29:2b)
  - Source: AsustekC\_39:29:2b (00:11:d8:39:29:2b)
  - Type: IPv6 (0x86dd)
- Internet Protocol Version 6
  - Version: 6
  - Traffic class: 0x00
  - Flowlabel: 0x00000
  - Payload length: 24
  - Next header: ICMPv6 (0x3a)
  - Hop limit: 255
  - Source address: ::
  - Destination address: ff02::1:ff39:292b
- Internet Control Message Protocol v6
  - Type: 135 (Neighbor solicitation)
  - Code: 0
  - Checksum: 0x504d [correct]
  - Target: fe80::211:d8ff:fe39:292b

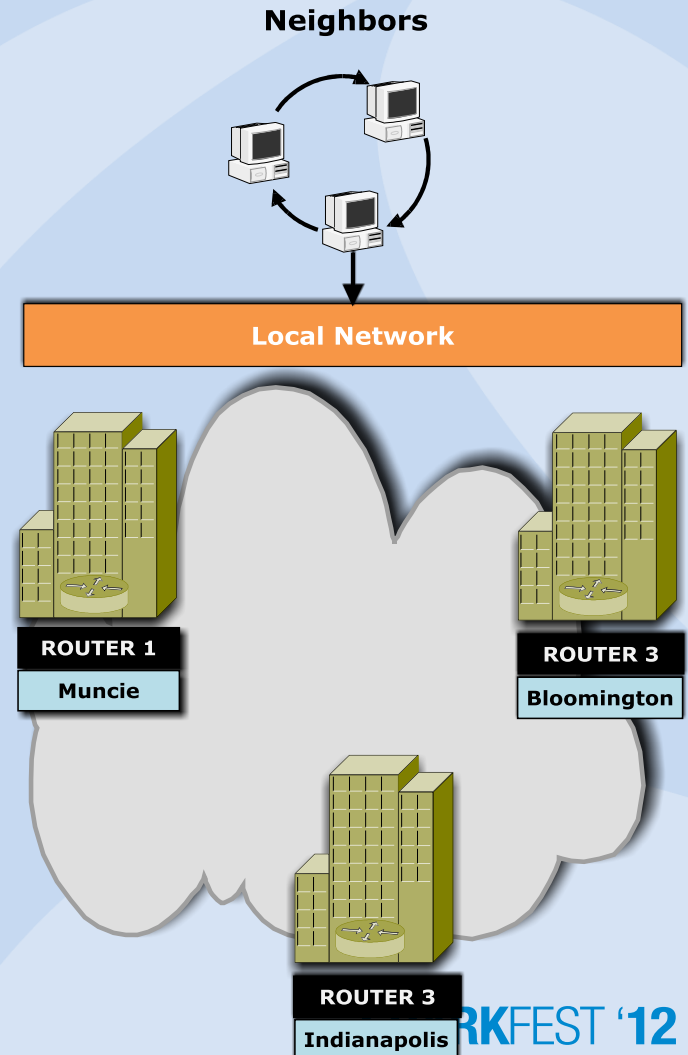
- Notice the sequence of events.
- Where is the MAC address?
- What is the Next Header field?
- What address do you think will be assigned?

- What kind of an address is ::?
- How about ff02::2?
- How about ff02::1:ff39:292b?
- And fe80::211:d8ff:fe39:292b?



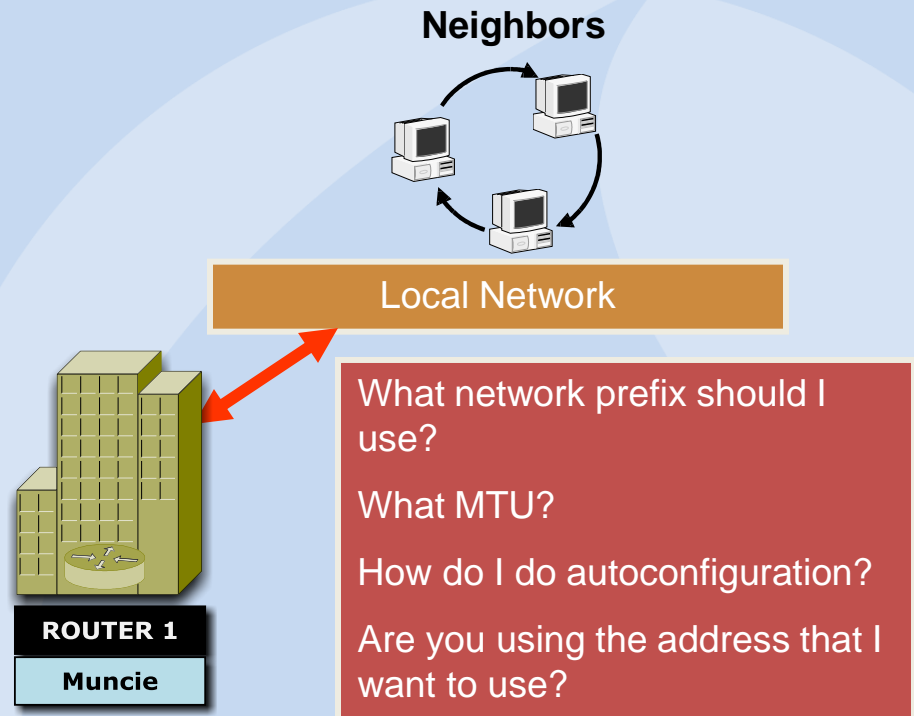
# What is a Neighbor?

- Two devices are *neighbors* if they are on the same local network
- Either a host or a router.



# What is Discovery?

- Not just who our neighbors are but also important information about them.
- Such as:
  - address resolution,
  - parameter communication,
  - autoconfiguration,
  - local network connectivity,
  - datagram routing and
  - configuration.



# Neighbor Discovery Standards

- The Neighbor Discovery protocol originally defined in RFC 1970 (1996) revised in RFC 2461 (1998) and ongoing....
- Most of the functions of the ND protocol are implemented using a set of four ICMPv6 control messages.
- ND can use of the authentication and encryption with IPSec

## Neighbor Discovery Messages - ICMPv6

**Router Advertisement**

**Router Solicitation**

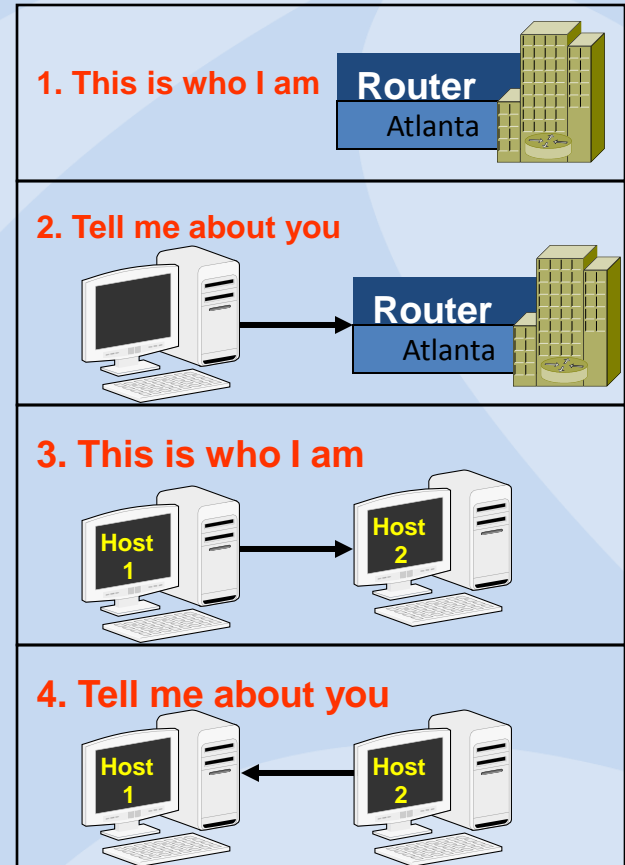
**Neighbor Advertisement**

**Neighbor Solicitation**

# ND Implementation – ICMPv6

- ND implements its functions using ICMPv6 messages.

- 1. Router Advertisement Messages:** Sent regularly by routers to tell hosts that they exist and provide important prefix and parameter information to them.
- 2. Router Solicitation Messages:** Sent by hosts to request that any local routers send a *Router Advertisement* message so they don't have to wait for the next regular advertisement message.
- 3. Neighbor Advertisement Messages:** Sent by hosts to indicate the existence of the host and provide information about it.
- 4. Neighbor Solicitation Messages:** Sent to verify the existence of another host and to ask it to transmit a *Neighbor Advertisement*.



No. -	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::214:bfff:feba:45f9	ff02::1	ICMPv6	Router advertisement
<div style="background-color: #f0f0f0; padding: 2px;">           Frame 1 (110 bytes on wire, 110 bytes captured)         </div> <div style="background-color: #f0f0f0; padding: 2px;">           Ethernet II, Src: 192.168.1.1 (00:14:bf:ba:45:f9), Dst: IPv6-Neighbor-Discovery_00:00:00:01 (33:33:00:00:00:01)                       Destination: IPv6-Neighbor-Discovery_00:00:00:01 (33:33:00:00:00:01)                       Source: 192.168.1.1 (00:14:bf:ba:45:f9)                       Type: IPv6 (0x86dd)         </div> <div style="background-color: #f0f0f0; padding: 2px;">           Internet Protocol Version 6                       Version: 6                       Traffic class: 0x00                       Flowlabel: 0x00000                       Payload length: 56                       Next header: ICMPv6 (0x3a)                       Hop limit: 255                       Source address: fe80::214:bfff:feba:45f9                       Destination address: ff02::1         </div> <div style="background-color: #f0f0f0; padding: 2px;">           Internet Control Message Protocol v6                       Type: 134 (Router advertisement)                       Code: 0                       Checksum: 0xecdd [correct]                       Cur hop limit: 64                       Flags: 0x00                         0... .. = Not managed                         .0.. .. = Not other                         ..0. .. = Not Home Agent                         ...0 0... = Router preference: Medium                       Router lifetime: 1800                       Reachable time: 0                       Retrans time: 0         </div> <div style="background-color: #f0f0f0; padding: 2px;">           ICMPv6 options                       Type: 3 (Prefix information)                       Length: 32 bytes (4)                       Prefix length: 64         </div> <div style="background-color: #f0f0f0; padding: 2px;">           Flags: 0xc0                         1... .. = Onlink                         .1.. .. = Auto                         ..0. .... = Not router address                         ...0 .... = Not site prefix                       valid lifetime: 0x00278d00                       Preferred lifetime: 0x00093a80                       Prefix: 2001:4840:ffff:c012:214:bfff:feba:45f9         </div> <div style="background-color: #f0f0f0; padding: 2px;">           ICMPv6 options                       Type: 1 (Source link-layer address)                       Length: 8 bytes (1)                       Link-layer address: 00:14:bf:ba:45:f9         </div>					

## Router Advertisement Packet

Source Address : MUST be the link-local address assigned to the interface from which this message is sent.

Destination Address: Typically the Source Address of an invoking Router Solicitation or the all-nodes multicast address.

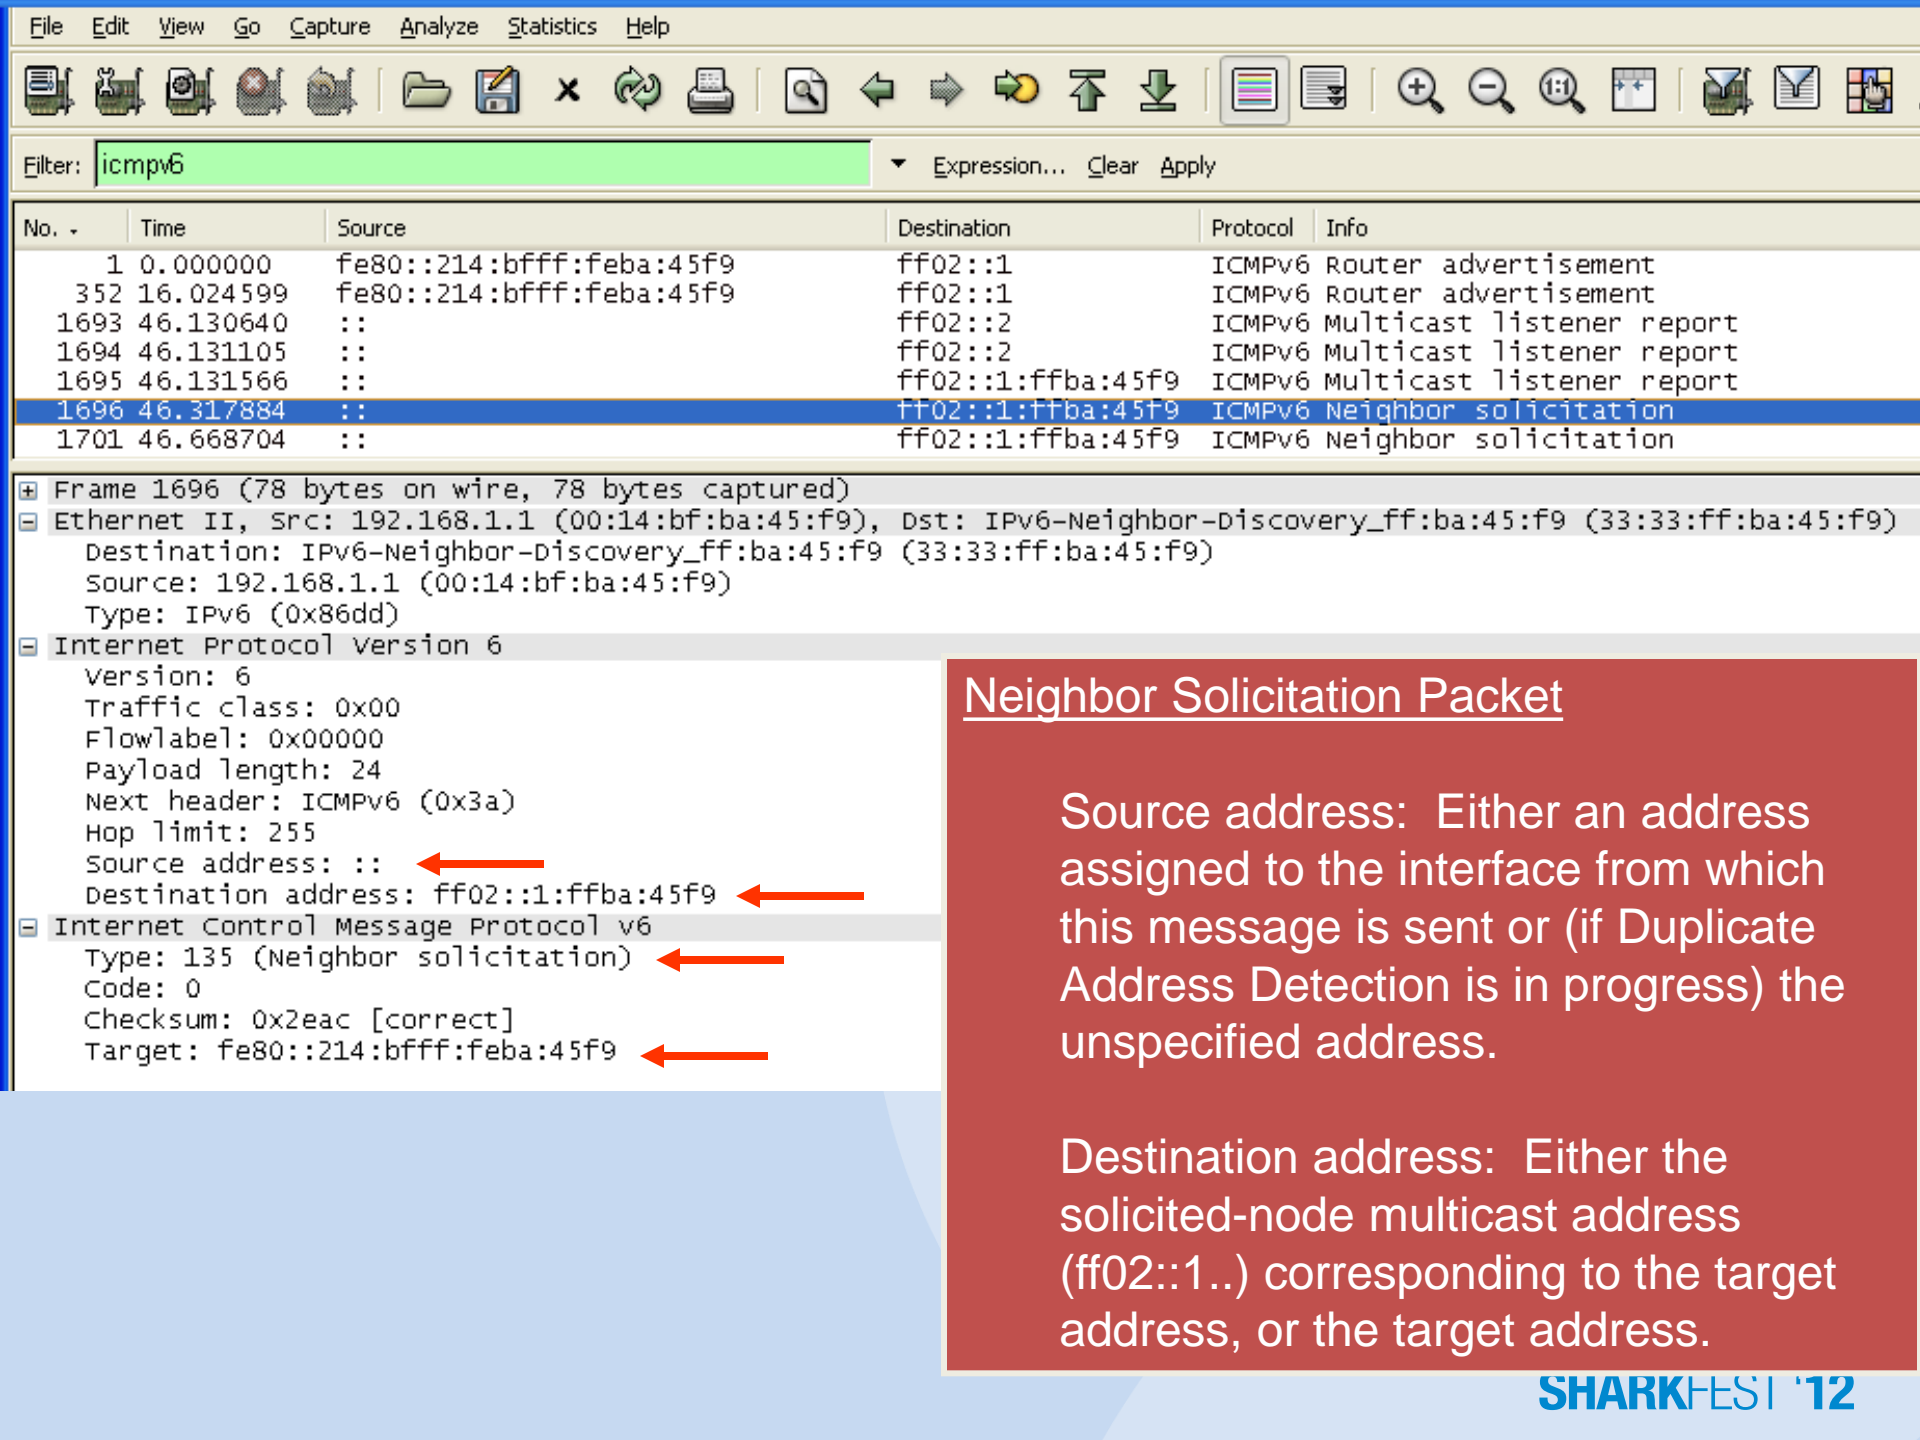
No.	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::214:bfff:feba:45f9	ff02::1	ICMPv6	Router advertisement
352	16.024599	fe80::214:bfff:feba:45f9	ff02::1	ICMPv6	Router advertisement
1693	46.130640	::	ff02::2	ICMPv6	Multicast listener report
1694	46.131105	::	ff02::2	ICMPv6	Multicast listener report
1695	46.131566	::	ff02::1:ffba:45f9	ICMPv6	Multicast listener report
1696	46.317884	::	ff02::1:ffba:45f9	ICMPv6	Neighbor solicitation
1701	46.668704	::	ff02::1:ffba:45f9	ICMPv6	Neighbor solicitation
2172	48.981759	fe80::214:bfff:feba:45f9	ff02::2	ICMPv6	Multicast listener report
2173	51.164256	fe80::214:bfff:feba:45f9	ff02::1:ffba:45f9	ICMPv6	Multicast listener report
2183	53.755395	fe80::214:bfff:feba:45f9	ff02::2	ICMPv6	Multicast listener report
2184	54.695543	fe80::214:bfff:feba:45f9	ff02::1:ffba:45f9	ICMPv6	Multicast listener report
2204	75.487400	fe80::211:d8ff:fe39:292b	ff02::1:ff39:292b	ICMPv6	Multicast listener report
2205	75.487419	fe80::211:d8ff:fe39:292b	ff02::1:ff3e:9113	ICMPv6	Multicast listener report
2206	75.487435	fe80::211:d8ff:fe39:292b	ff02::2	ICMPv6	Router solicitation

- Frame 2206 (70 bytes on wire, 70 bytes captured)
- Ethernet II, Src: 192.168.1.100 (00:11:d8:39:29:2b), Dst: IPv6-Neighbor-Discovery\_00:00:00:02 (33:33:00:00:00:02)
  - Destination: IPv6-Neighbor-Discovery\_00:00:00:02 (33:33:00:00:00:02)
  - Source: 192.168.1.100 (00:11:d8:39:29:2b)
  - Type: IPv6 (0x86dd)
- Internet Protocol Version 6
  - Version: 6
  - Traffic class: 0x00
  - Flowlabel: 0x00000
  - Payload length: 16
  - Next header: ICMPv6 (0x3a)
  - Hop limit: 255
  - Source address: fe80::211:d8ff:fe39:292b ←
  - Destination address: ff02::2 ←
- Internet Control Message Protocol v6
  - Type: 133 (Router solicitation) ←
  - Code: 0
  - Checksum: 0x7842 [correct]
- ICMPv6 options
  - Type: 1 (Source link-layer address)
  - Length: 8 bytes (1)
  - Link-layer address: 00:11:d8:39:29:2b

## Router Solicitation Packet

Source address: usually the unspecified IPv6 address (0:0:0:0:0:0:0:0) or configured unicast address of the interface.

Destination address: the all-routers multicast address (FF02::2) with the link-local scope.



Filter: icmpv6 Expression... Clear Apply

No. -	Time	Source	Destination	Protocol	Info
1	0.000000	fe80::214:bfff:feba:45f9	ff02::1	ICMPv6	Router advertisement
352	16.024599	fe80::214:bfff:feba:45f9	ff02::1	ICMPv6	Router advertisement
1693	46.130640	::	ff02::2	ICMPv6	Multicast listener report
1694	46.131105	::	ff02::2	ICMPv6	Multicast listener report
1695	46.131566	::	ff02::1:ffba:45f9	ICMPv6	Multicast listener report
1696	46.317884	::	ff02::1:ffba:45f9	ICMPv6	Neighbor solicitation
1701	46.668704	::	ff02::1:ffba:45f9	ICMPv6	Neighbor solicitation

- Frame 1696 (78 bytes on wire, 78 bytes captured)
- Ethernet II, Src: 192.168.1.1 (00:14:bf:ba:45:f9), Dst: IPv6-Neighbor-Discovery\_ff:ba:45:f9 (33:33:ff:ba:45:f9)
  - Destination: IPv6-Neighbor-Discovery\_ff:ba:45:f9 (33:33:ff:ba:45:f9)
  - Source: 192.168.1.1 (00:14:bf:ba:45:f9)
  - Type: IPv6 (0x86dd)
- Internet Protocol Version 6
  - Version: 6
  - Traffic class: 0x00
  - Flowlabel: 0x00000
  - Payload length: 24
  - Next header: ICMPv6 (0x3a)
  - Hop limit: 255
  - Source address: :: ←
  - Destination address: ff02::1:ffba:45f9 ←
- Internet Control Message Protocol v6
  - Type: 135 (Neighbor solicitation) ←
  - Code: 0
  - Checksum: 0x2eac [correct]
  - Target: fe80::214:bfff:feba:45f9 ←

### Neighbor Solicitation Packet

Source address: Either an address assigned to the interface from which this message is sent or (if Duplicate Address Detection is in progress) the unspecified address.

Destination address: Either the solicited-node multicast address (ff02::1..) corresponding to the target address, or the target address.



No.	Time	Source	Destination	Protocol	Info
6	9.865886	fe80::2ff:8cff:fe10:3976	2001:5c0:8fff:fffe	ICMPv6	Neighbor solicitation
7	9.865895	2001:5c0:8fff:fffe::3f52	fe80::2ff:8cff:fe1	ICMPv6	Neighbor advertisement

```

+ Frame 7 (86 bytes on wire, 86 bytes captured)
+ Ethernet II, Src: 00:ff:8d:10:39:76 (00:ff:8d:10:39:76), Dst: 00:ff:8c:10:39:76 (00:ff:8c:10:39:76)
- Internet Protocol Version 6
  Version: 6
  Traffic class: 0x00
  Flowlabel: 0x00000
  Payload length: 32
  Next header: ICMPv6 (0x3a)
  Hop limit: 255
  Source address: 2001:5c0:8fff:fffe::3f52
  Destination address: fe80::2ff:8cff:fe10:3976
- Internet Control Message Protocol v6
  Type: 136 (Neighbor advertisement) ←
  Code: 0
  Checksum: 0xbdf3 [correct]
+ Flags: 0x40000000
  0... .. = Not router
  .1.. .. = solicited
  ..0. .... = Not override
  Target: 2001:5c0:8fff:fffe::3f52
- ICMPv6 options
  Type: 2 (Target link-layer address)
  Length: 8 bytes (1)
  Link-layer address: 00:ff:8d:10:39:76

```

Neighbor Advertisement

- ICMP type 136

- From RFC2461: A node sends Neighbor Advertisements in response to Neighbor Solicitations and sends unsolicited Neighbor Advertisements in order to (unreliably) propagate new information quickly.

No. -	Time	Source	Destination	Protocol	Info
6	9.865886	fe80::2ff:8cff:fe10:3976	2001:5c0:8fff:fffe	ICMPv6	Neighbor solicitation
7	9.865895	2001:5c0:8fff:fffe::3f52	fe80::2ff:8cff:fe1	ICMPv6	Neighbor advertisement


  

+ Frame 6 (86 bytes on wire, 86 bytes captured)  
 + Ethernet II, Src: 00:ff:8c:10:39:76 (00:ff:8c:10:39:76), Dst: 00:ff:8d:10:39:76 (00:ff:8d:10:39:76)  
 - Internet Protocol Version 6  
   Version: 6  
   Traffic class: 0x00  
   Flowlabel: 0x00000  
   Payload length: 32  
   Next header: ICMPv6 (0x3a)  
   Hop limit: 255  
   Source address: fe80::2ff:8cff:fe10:3976 ←  
   Destination address: 2001:5c0:8fff:fffe::3f52 ←  
 - Internet Control Message Protocol v6  
   Type: 135 (Neighbor solicitation) ←  
   Code: 0  
   Checksum: 0x00f4 [correct]  
   Target: 2001:5c0:8fff:fffe::3f52  
 - ICMPv6 options  
   Type: 1 (source link-layer address)  
   Length: 8 bytes (1)  
   Link-layer address: 00:ff:8c:10:39:76

Neighbor Solicitation Packet  
 To a specific unicast address.  
 Duplicate Address Detection

# Multicast Group Membership

- Group membership is dynamic, allowing hosts to join and leave the group at any time.
- The joining of multicast groups is performed through the sending of group membership messages.
- In IPv6, Multicast Listener Discovery (MLD) messages are used to determine group membership on a network segment.



Multicast Group at 10:00 am



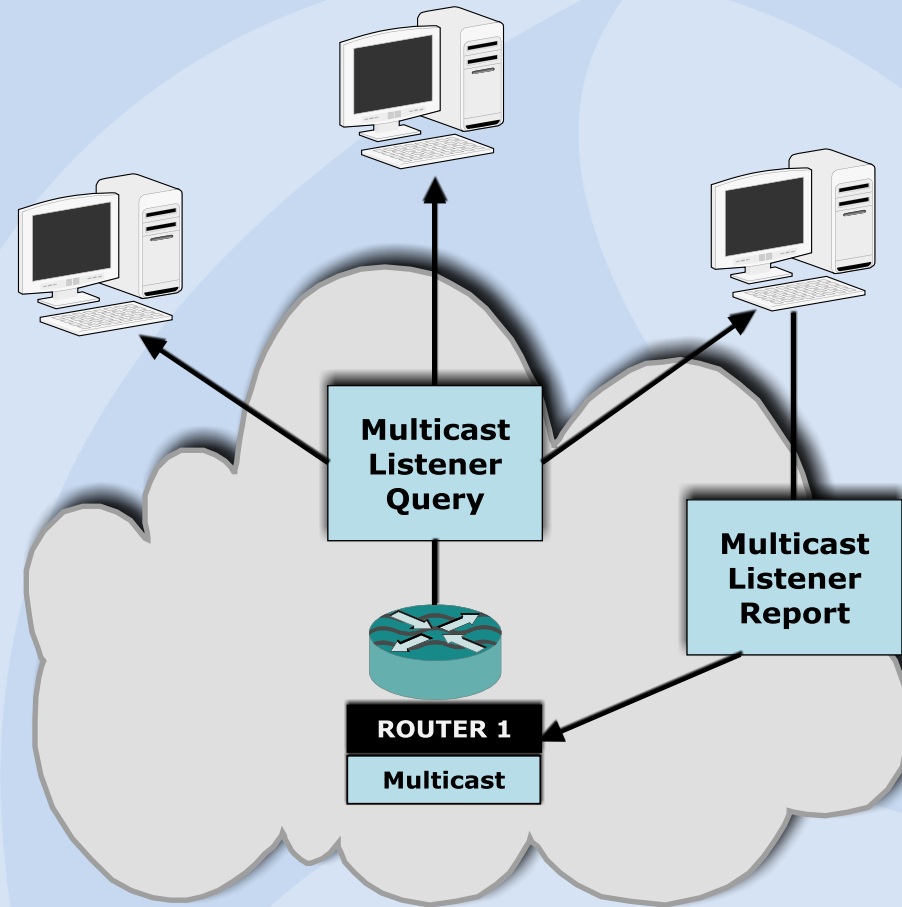
Multicast Group at 11:00 am



Multicast group at 2:00 pm

# Multicast Listener Discovery

- MLD is used to exchange membership status information between IPv6 routers that support multicasting and members of multicast groups on a network segment.
- Host membership in a multicast group is reported by individual member hosts, and membership status is periodically polled by multicast routers.
- MLD is defined in RFC 2710, "Multicast Listener Discovery (MLD) for IPv6."



# MLD Message Types

MLD message type	Description
Multicast Listener Query	Sent by a multicast router to poll a network segment for group members. Queries can be general (requesting group membership for all groups), or specific (requesting group membership for a specific group).
Multicast Listener Report	Sent by a host when it joins a multicast group, or in response to a MLD Multicast Listener Query sent by a router.
Multicast Listener Done	Sent by a host when it leaves a host group and might be the last member of that group on the network segment.

No. -	Time	Source	Destination	Protocol	Info
1693	46.130640	::	ff02::2	ICMPv6	Multicast listener report
+ Frame 1693 (86 bytes on wire, 86 bytes captured)					
- Ethernet II, Src: 192.168.1.1 (00:14:bf:ba:45:f9), Dst: IPv6-Neighbor-Discovery_00:00:00:02 Destination: IPv6-Neighbor-Discovery_00:00:00:02 (33:33:00:00:00:02) Source: 192.168.1.1 (00:14:bf:ba:45:f9) Type: IPv6 (0x86dd)					
- Internet Protocol version 6 ←					
Version: 6 Traffic class: 0x00 Flowlabel: 0x00000 Payload length: 32 Next header: IPv6 hop-by-hop option (0x00) Hop limit: 1 Source address: :: Destination address: ff02::2					
- Hop-by-hop Option Header Next header: ICMPv6 (0x3a) Length: 0 (8 bytes) Router alert: MLD (4 bytes) PadN: 2 bytes					
- Internet Control Message Protocol v6 Type: 131 (Multicast listener report) ← Code: 0 Checksum: 0x7ea3 [correct] Maximum response delay: 0 Multicast Address: ff02::2 ←					

# RFC3971 SEcure Neighbor Discovery

To secure the various functions in NDP, a set of new Neighbor Discovery options is introduced. The components of the solution are:

- Certification paths, anchored on trusted parties, are expected to certify the authority of routers.
- A host must be configured with a trust anchor to which the router has a certification path before the host can adopt the router as its default router.
- Certification Path Solicitation and Advertisement messages are used to discover a certification path to the trust anchor without requiring the actual Router Discovery messages to carry lengthy certification paths.
- The receipt of a protected Router Advertisement message for which no certification path is available triggers the authorization delegation discovery process.
- Cryptographically Generated Addresses are used to make sure that the sender of a Neighbor Discovery message is the "owner" of the claimed address.
- A public-private key pair is generated by all nodes before they can claim an address.
- A new NDP option, the CGA option, is used to carry the public key and associated parameters.

# Summary

- I will have a job forever because no one can keep up with all this!
- Email: [nalini.elkins@insidestack.com](mailto:nalini.elkins@insidestack.com)