

IPv6 Neighbor Discovery

SLAAC, NEIGHBOR DISCOVERY, MULTICAST LISTENER DISCOVERY

NALINI ELKINS

INSIDE PRODUCTS, INC.

NALINI.ELKINS@INSIDETHESTACK.COM



Collaborative Project

- India Internet Engineering Society (IIEsoc) and Industry Network Technology Council (INTC)
- Funding: Grant from ISIF Asia
- Thank you!



<https://www.iiesoc.in/>

<https://industry.netcouncil.org/>

Vision

Multi-year project: IPv6 deployment at enterprises.

Collaboration with American Registry for Internet Numbers (ARIN)

- Provide training,
- Analysis of security and application conversion,
- Help enterprises plan their IPv6 deployment.

Classes

- Introduction to IPv6 : Feb 4, 2021 ✓
- Lab: IPv6 basics : Feb 11, 2021 ✓
- Neighbor Discovery: March 4, 2021
- Lab: Neighbor Discovery: March 18, 2021
- IPv6 Address Planning: April 8, 2021
- Lab: IPv6 Address Planning: April 15, 2021
- IPv6 Transition Mechanisms: May 6, 2021
- Lab: IPv6 Transition Mechanisms: May 13, 2021

- DHCPv6: June 3, 2021
 - Lab: DHCPv6: June 10, 2021
 - IPv6 and Cloud: June 17, 2021
 - Lab: IPv6 and Cloud: June 24, 2021
 - Introduction to IPv6 Security July 8, 2021
- The next sessions are sponsored by a generous grant from ARIN.
- Trace Reading: August 12, 2021
 - Troubleshooting: August 19, 2021

A few words about me

- President: Industry Network Technology Council
- Founder & CEO: Inside Products, Inc.
- Advisory Board: India Internet Engineering Society
- RFCs: RFC8250 (Embedded performance and diagnostics for IPv6) and others
- Product developer (OEMed by IBM and others)
- Working with IPv6 for 15 years
- Working with network management, diagnostic, performance issues at large brick-and-mortar enterprises for over 30 years



Agenda

- Stateless autoconfiguration
- ICMPv6
- Neighbor Discovery
 - ✓ Neighbor Solicitation / Advertisement
 - ✓ Router Solicitation / Advertisement
- Multicast Listener Discovery

Why talk about this?

- Stateless address autoconfiguration (SLAAC) new in IPv6
- Stateful : DHCPv6 (upcoming session!)
- Address planning: host address (64 bits)
- Security planning
 - Protect privacy
 - Protect topology,
 - Potential new attacks)

Link-Local Unicast Address

- IPv6 devices always have a link-local address
- IPv6 devices use link-local to communicate with 'on-link' devices
- IPv6 routers must not forward link-local packets

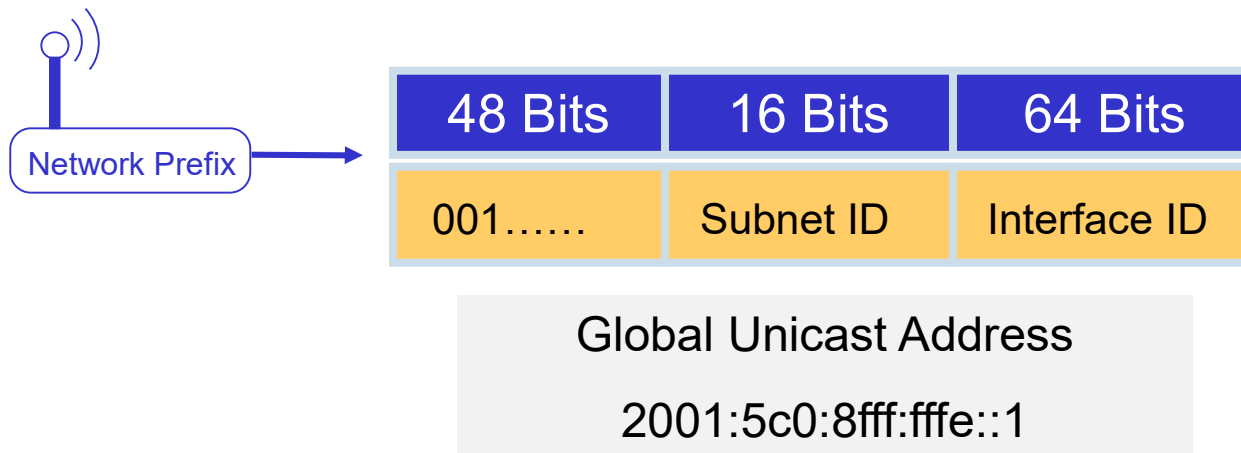
10 Bits	54 Bits	64 Bits
1111111010	zeroes	Interface ID

Sample Link-Local Address

fe80::211:d8ff:fe39:292b

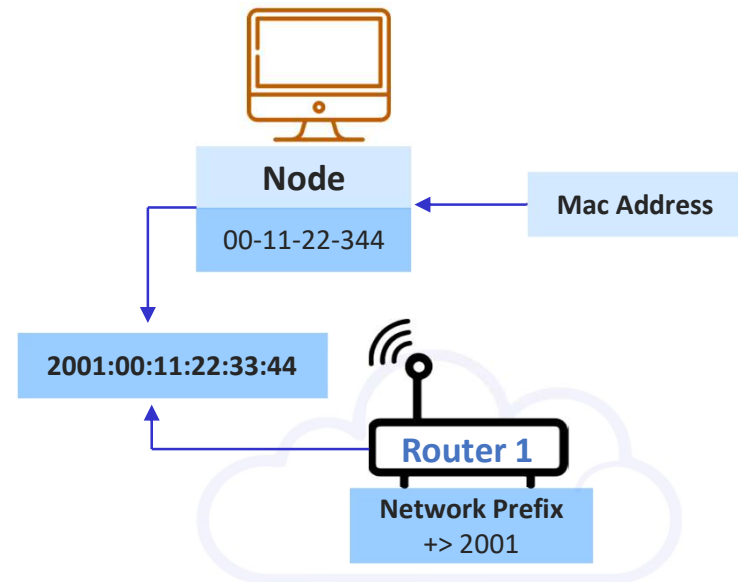
Global Unicast Interface ID (IID)

- IID is for an interface
- IID must be unique
- IID: standard is 64 bits



Stateless Autoconfiguration

- Stateless autoconfiguration (SLAAC)
- Link-local and global unicast address
- How?
- Use MAC address of adapter (Original)
- Talk with connected IPv6 router
- Join multicast groups



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

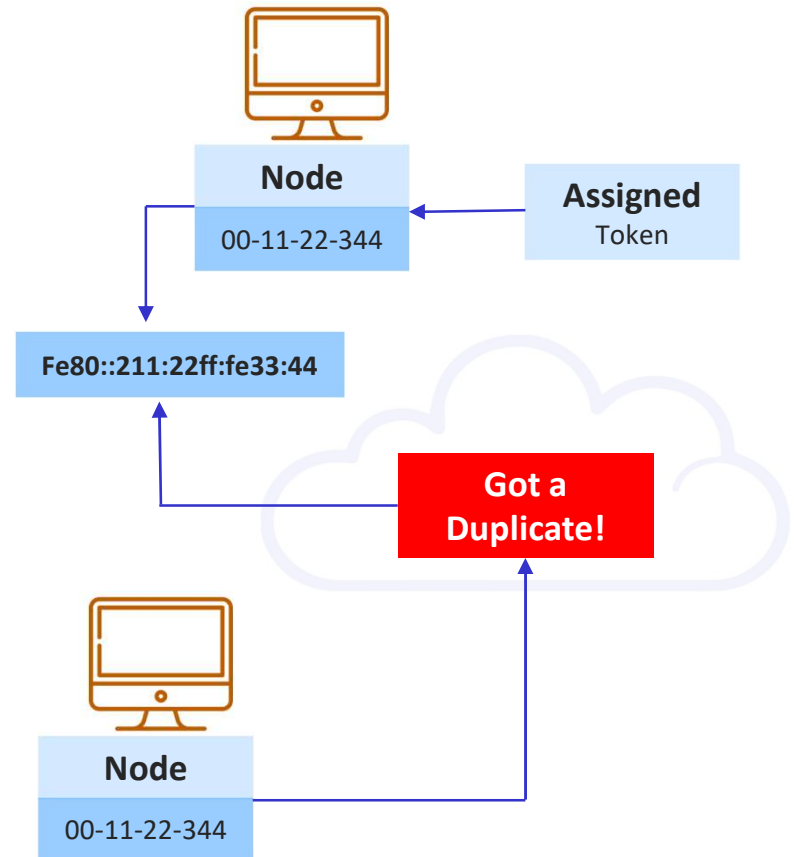
Link-Local Address Generate

Link-Local Address

- TCP/IP stack
- FE80....

Link-Local Address Uniqueness

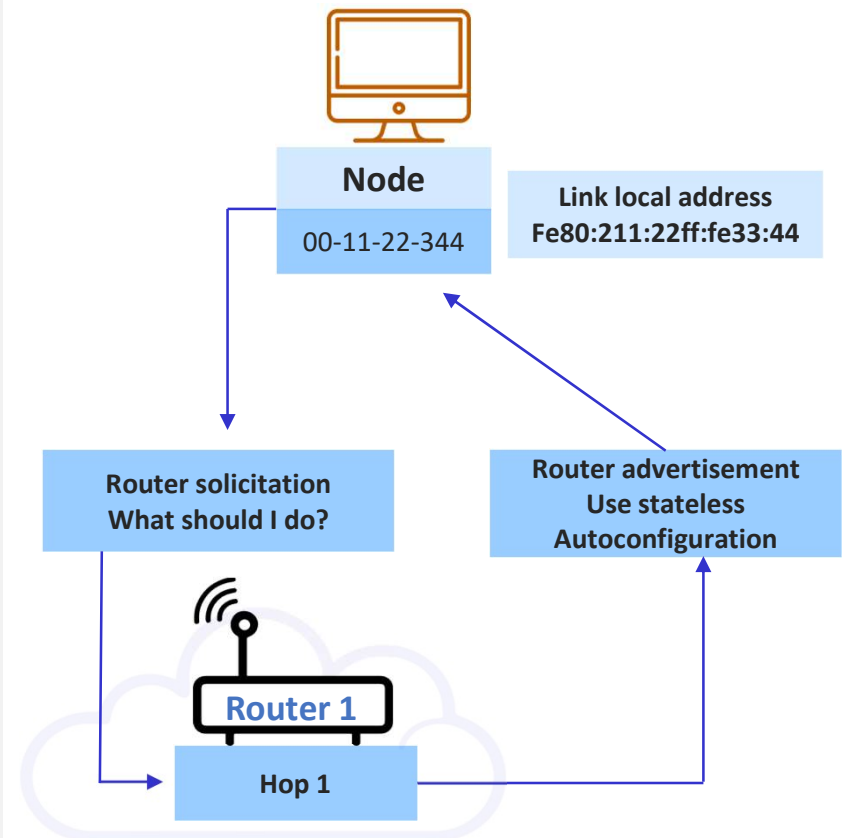
- Duplicate Address Detection (DAD)
- ICMPv6 *Neighbor Solicitation* message
- ICMPv6 *Neighbor Advertisement*



Generate Global Unicast Address

Global unicast

- Communicate over internet
- Other side of local router
- Have IPv6-enabled router?
- ICMPv6 *Router Advertisement / Router Solicitation* messages



Global Unicast Assigned

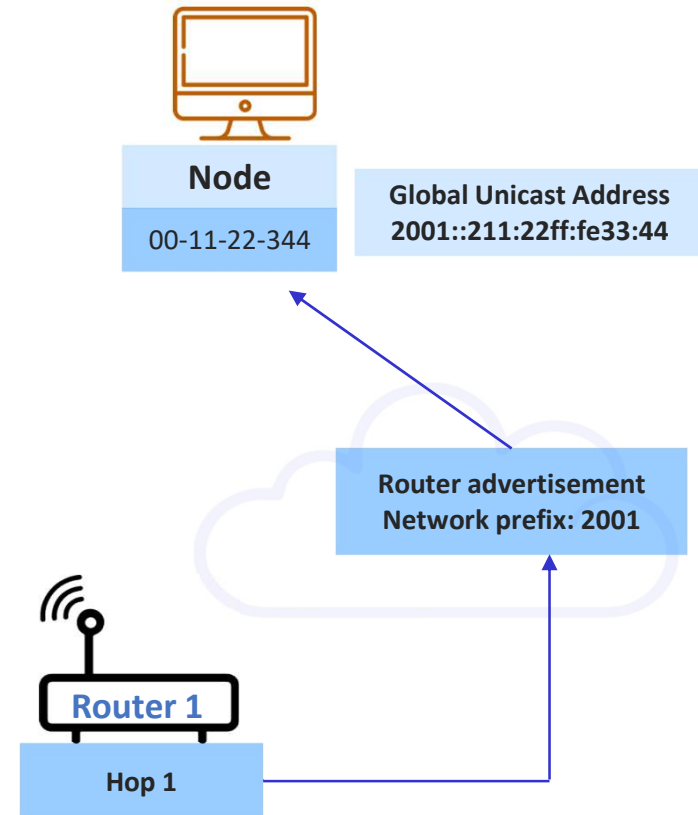
Global unicast address

- Network prefix
- Device identifier (IID)

SLAAC advantages:

- Easy to manage
- No DHCP server required
- Mobile / sensors

Server addresses



The Interface ID is the Issue

- Interface ID (IID): based on the link-layer (MAC) address
- EUI-64 format: OUI field + FFFE + Serial Number

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

Why not use MAC based IID?

- “Since the resulting Interface Identifiers are constant across networks, the resulting IPv6 addresses can be leveraged to track and correlate the activity of a host across multiple networks (e.g., track and correlate the activities of a typical client connecting to the public Internet from different locations), thus negatively affecting the privacy of users.” [RFC7217]
- This means: you can tell what device is going to what web site.



Specific Address Patterns

- “The IPv6 addresses of all hosts manufactured by the same vendor (within a given time frame) will likely contain the same IEEE Organizationally Unique Identifier (OUI) in the Interface Identifier.” [RFC7217]
- So attacker in network can do scans easier.

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

Sample OUIs

08-61-95 (hex)	Rockwell
F4-BD-9E (hex)	Cisco
40-55-82 (hex)	Nokia

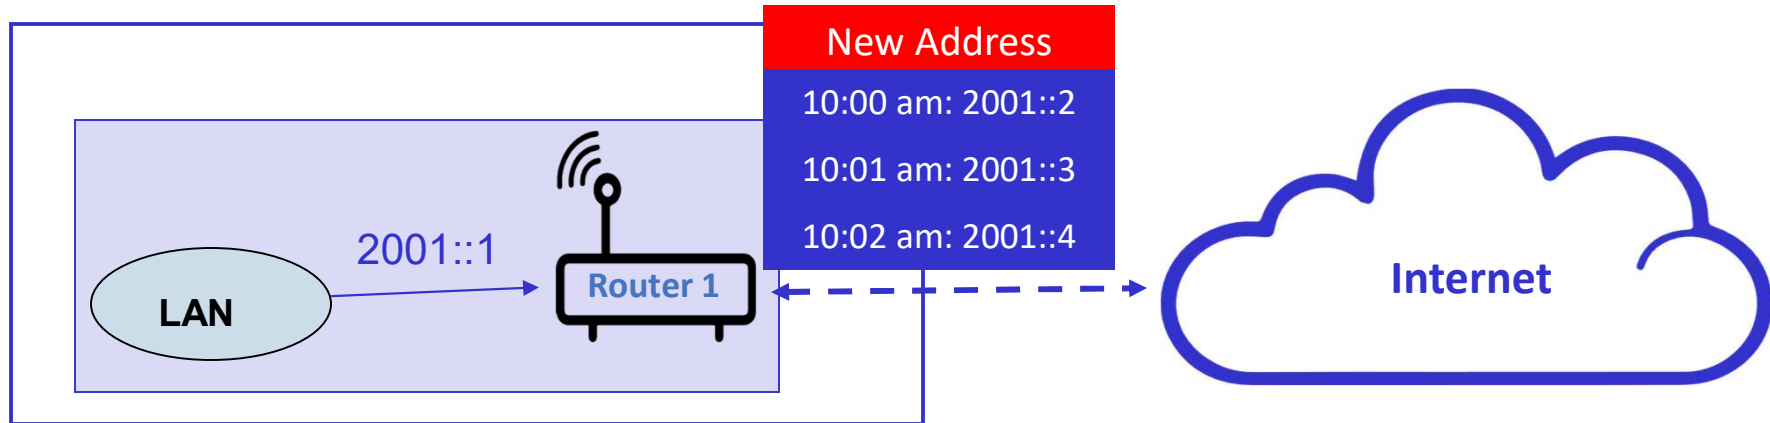
<http://standards-oui.ieee.org/oui/oui.txt>

Device Specific Attacks

“Embedding the underlying hardware address in the Interface Identifier leaks device-specific information that could be leveraged to launch device-specific attacks.” [RFC7217]

- You can get a pretty good idea from the hardware address what kind of device it is.
- On my home network, if the OUI is for HP, then probably it is my printer.
- If the OUI is for Apple, it may be my daughter’s iPhone.

IPv6 Privacy Addresses



- Anonymous addressing
- Change address frequently
- How to diagnose problems?
- Implementation differences

RFC4941 "Privacy Extensions for Stateless Address Autoconfiguration in IPv6"

Network Management

➤ But, you have to be able to do network management!

➤ If addresses change all the time, how do you do that?

“In a variety of scenarios, addresses that remain stable for the lifetime of a host's connection to a single subnet are viewed as desirable. For example, stable addresses may be viewed as beneficial for network management, event logging, enforcement of access control, provision of quality of service, or for server or router interfaces. Similarly, stable addresses (as opposed to temporary addresses [RFC4941]) allow for long-lived TCP connections and are also usually desirable when performing server-like functions (i.e., receiving incoming connections).”
[RFC8064: Recommendation on Stable IPv6 Interface Identifiers]

RFC8064

Gotta love it!

“This document changes the recommended default Interface Identifier (IID) generation scheme for cases where Stateless Address Autoconfiguration (SLAAC) is used to generate a stable IPv6 address. It recommends using the mechanism specified in RFC 7217 in such cases, and recommends against embedding stable link-layer addresses in IPv6 IIDs. It formally updates RFC 2464, RFC 2467, RFC 2470, RFC2491, RFC 2492, RFC 2497, RFC 2590, RFC 3146, RFC 3572, RFC 4291, RFC4338, RFC 4391, RFC 5072, and RFC 5121. This document does not change any existing recommendations concerning the use of temporary addresses as specified in RFC 4941.”

Stateless Autoconfig on Windows

```
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
```

- Windows XP - no IPv6 (remember those days!)
- IPconfig : shows interfaces / addresses
- Only IPv4
- IPv6 install starts SLAAC
- Of course, now IPv6 is preconfigured



After IPv6 Installed Successfully

```
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>
```

- New addresses –new interfaces
- Only link-local addresses
- No IPv6 router
- SLAAC assigned addresses
- Tunneling automatically created

IPConfig with Global Unicast Addresses

```
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 . . . . . :
```

- Global unicast '2001...'
- Used to get Teredo and Tunneling addresses



Ethernet adapter Ethernet:

```
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
```

Ethernet adapter Npcap Loopback Adapter:

```
Connection-specific DNS Suffix . :
Link-local IPv6 Address . . . . . : fe80::ec34:cb9b:bc2b:26d0%6
Autoconfiguration IPv4 Address. . : 169.254.38.208
Subnet Mask . . . . . : 255.255.0.0
Default Gateway . . . . . :
```

Wireless LAN adapter Local Area Connection* 2:

```
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
```

Wireless LAN adapter Local Area Connection* 13:

```
Media State . . . . . : Media disconnected
Connection-specific DNS Suffix . :
```

Wireless LAN adapter Wi-Fi:

```
Connection-specific DNS Suffix . : hsd1.ca.comcast.net
IPv6 Address. . . . . : 2601:642:c202:9550::58b9
IPv6 Address. . . . . : 2601:642:c202:9550:fced:f576:4c8d:11f7
Temporary IPv6 Address. . . . . : 2601:642:c202:9550:1ce4:8f9c:c401:20b5
Link-local IPv6 Address . . . . . : fe80::fced:f576:4c8d:11f7%7
IPv4 Address. . . . . : 10.0.0.118
Subnet Mask . . . . . : 255.255.255.0
Default Gateway . . . . . : fe80::1256:11ff:fe99:e3d7%7
                            10.0.0.1
```

- Today's world
- Npcap (loopback for Wireshark)
- No ethernet
- Using WiFi to go to Internet
- Dual stacked

What is ICMPv6?

- Used by the Internet Protocol (IP)
- Used by SLAAC
- ICMPv4 == > ICMPv6 -- Many changes!
- ICMP has:
 - Error messages
 - Informational messages

Some important error messages

- Destination unreachable
- Packet too big
- Time exceeded
- Parameter problem

Some important informational messages:

- Echo request/reply
- Multicasting messages
 - Group membership query, report, done
- Neighbor discovery
 - Router solicitation and advertisement
 - Neighbor solicitation and advertisement
- Redirect

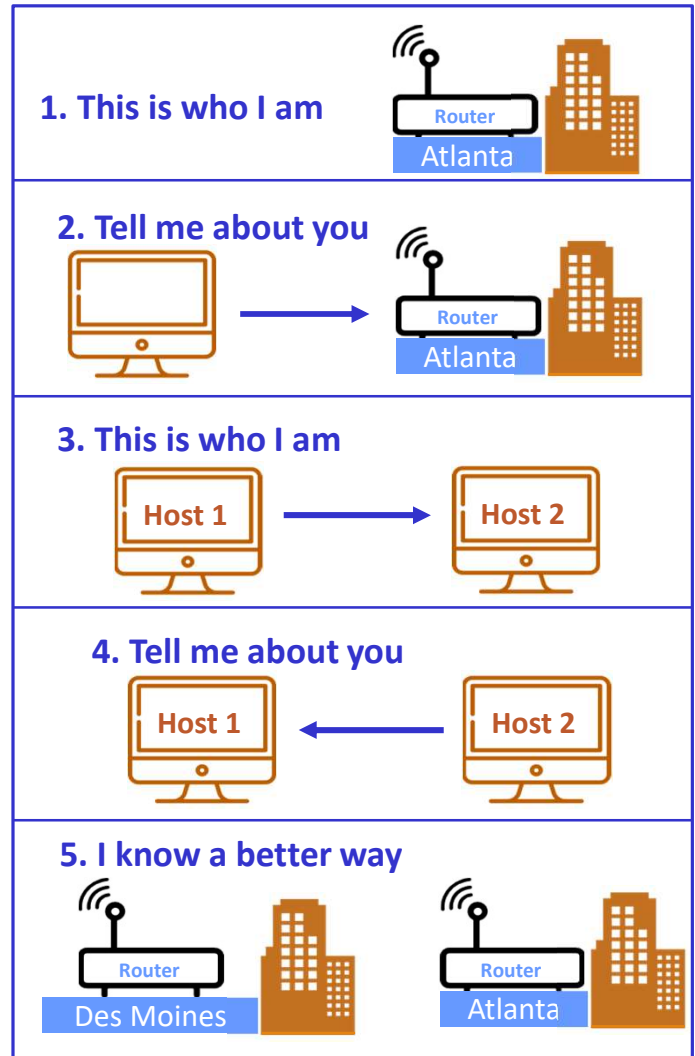
ICMPv6 Informational Messages

Type	Name	Type	Name
128	Echo Request	141	Inverse Neighbor Discovery Solicitation Message
129	Echo Reply	142	Inverse Neighbor Discovery Advertisement Message
130	Multicast Listener Query	143	Version 2 Multicast Listener Report
131	Multicast Listener Report	144	Home Agent Address Discovery Request Message
132	Multicast Listener Done	145	Home Agent Address Discovery Reply Message
133	Router Solicitation	146	Mobile Prefix Solicitation
134	Router Advertisement	147	Mobile Prefix Advertisement
135	Neighbor Solicitation	148	Certification Path Solicitation
136	Neighbor Advertisement	149	Certification Path Advertisement
137	Redirect Message	150	Experimental mobility protocols
138	Router Renumbering	151	Multicast Router Advertisement
139	ICMP Node Info. Query	152	Multicast Router Solicitation
140	ICMP Node Info. Response	153	Multicast Router Termination

Neighbor Discovery

- Neighbor Discovery (ND) replaces ARP
- RFC4861: Neighbor Discovery for IP version 6 (IPv6)
- Used in SLAAC
- Five ICMPv6 message types:

1. Router Advertisement
2. Router Solicitation
3. Neighbor Advertisement
4. Neighbor Solicitation
5. Redirect



Neighbor Discovery

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)

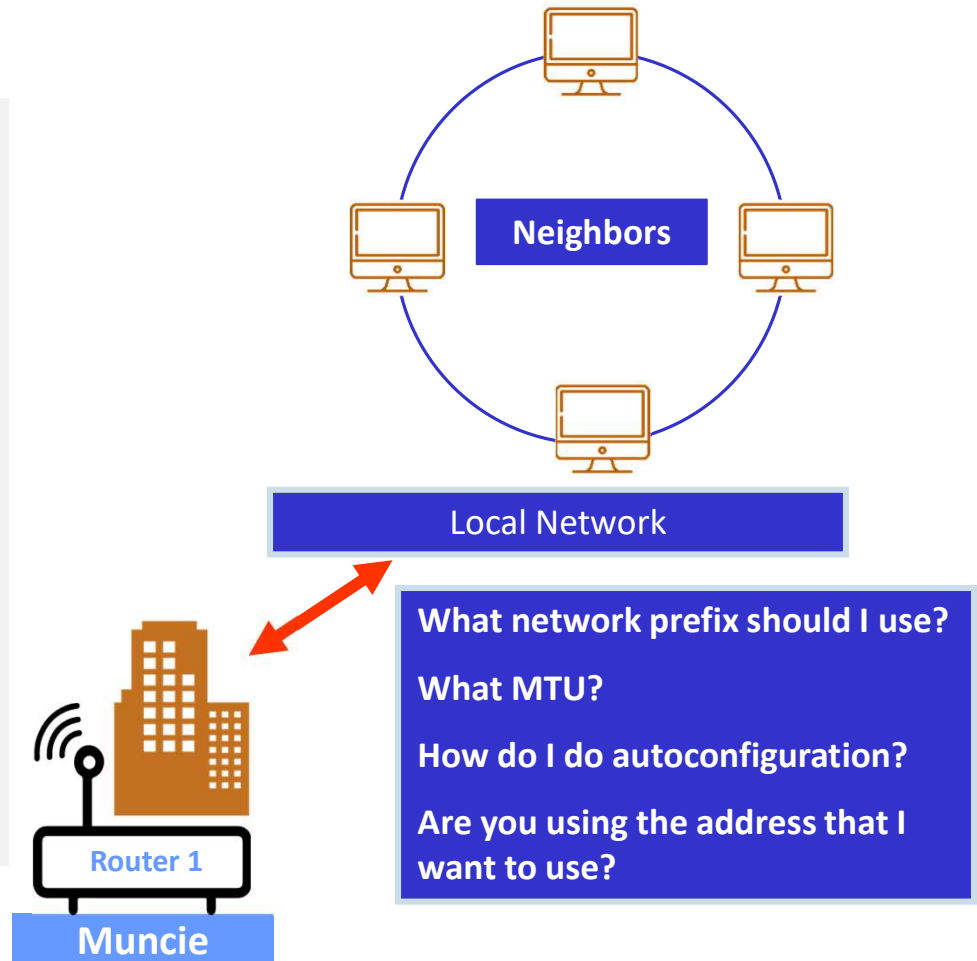
Ethernet II, Src: AsustekC_39:29:2b (00:11:d8:39:29:2b), Dst: IPv6-Neighbor-Discovery_ff:39:29:2b
Destination: IPv6-Neighbor-Discovery_ff:39:29:2b (33:33:ff:39:29:2b)
Source: AsustekC_39:29:2b (00:11:d8:39:29:2b)
Type: IPv6 (0x86dd)

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

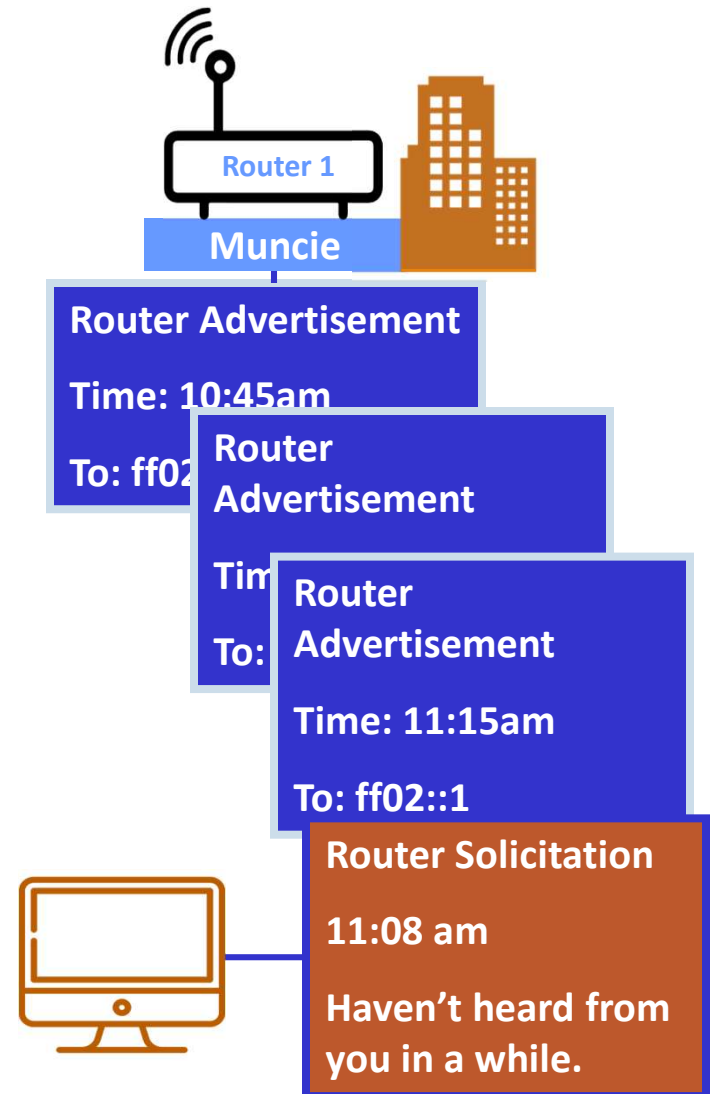
Neighbor and Discovery

- **What is a neighbor?**
- **What is discovery?**
 - Address resolution,
 - Parameter communication,
 - autoconfiguration,
 - local network connectivity,
 - routing and
 - configuration.



Router Advertisement (RA)


- Router Advertisement important for SLAAC.
- Sent at intervals
- Unsolicited RA sent to FF02 ::1
- Receiving hosts update configuration
- RA also responds to *Router Solicitation* (RS)
- Solicited RA sent to address of RS sender



Router Advertisement Contents

Router advertisements contain:

- Stateless / stateful (DHCPv6)
- Network prefix
- Default router
- Hop limit
- MTU



Router Advertisement

Time: 10:45am

To: ff02::1

- Use AutoConfiguration
- Stateless
- Network Prefix: 2001:: /64
- I am default router
- For 200 seconds
- Hop limit: 126
- MTU: 4096


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 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: #0056b3; color: white; 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
- Destination address
- ICMP type
- Hop limit
- Prefix length
- Prefix

Router Solicitation (RS)


- Sent during SLAAC
- Immediate response needed
- Sent 3 times total if no response



Router Solicitation

I need an address.

Please send a router advertisement



Router Solicitation Packet

```
⊟ Frame 2206 (70 bytes on wire, 70 bytes captured)
⊟ Ethernet II, Src: 192.168.1.100 (00:11:d8:39:29:2b),
  Destination: IPv6-Neighbor-Discovery_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
  Flow label: 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
- Destination address
- ICMPv6 type

Neighbor Advertisement (NA)

Neighbor Advertisements sent:

- In response to *Neighbor Solicitation*
- Or if own NIC changes
- Contains link-layer address



Neighbor Advertisement

To: fe80::1:2:3:4

• My link-local address is:
fe80::5:6:7:8

Neighbor Advertisement Packet

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

Neighbor Solicitation (NS)

- *Neighbor Solicitations* request information
- *Neighbor Advertisement* response
- Sent during SLAAC (DAD)
- Sent to verify reachability



Neighbor Solicitation

To: ff02::1

Are you using:

fe80::1:2:3:4?

Neighbor Solicitation Packet

```
⊕ Frame 25 (78 bytes on wire, 78 bytes captured)
⊖ Ethernet II, Src: AsustekC_39:29:2b (00:11:d8:39:29:2b), Dst: IPv6-Neig
  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
```


NS Packet (Reachability)

No. -	Time	Source	Destination	Protocol	Info
6	9.865886	fe80::2ff:8cff:fe10:3976	2001:5c0:8fff:fffe::3f52	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.




Multicast Groups

- Multicast: frequently used
 - All-nodes
 - All-routers
 - All-OSPF-routers
- Dynamic membership
- Multicast Listener Discovery (MLD) protocol used



Multicast Group at 10:00 am



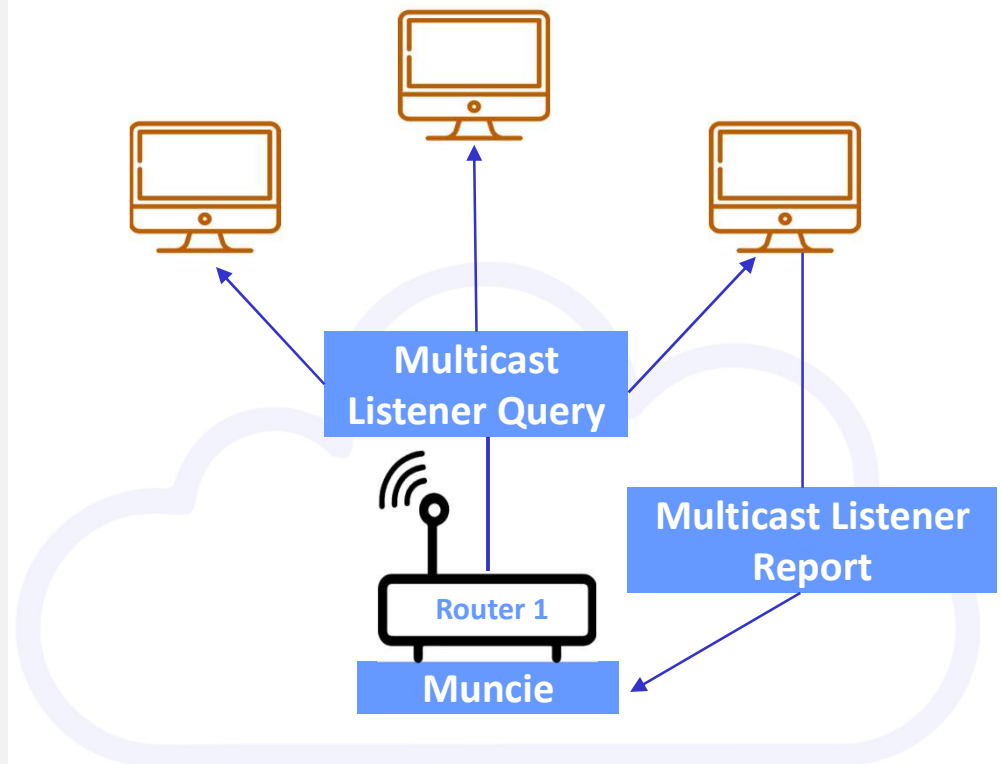
Multicast Group at 11:00 am



Multicast group at 2:00 pm

Multicast Listener Discovery

- RFC2710: Multicast Listener Discovery (MLD) for IPv6
- RFC3590: Source Address Selection for the Multicast Listener Discovery (MLD) Protocol
- RFC3810: Multicast Listener Discovery Version 2 (MLDv2) for IPv6



MLD Message Types

MLD message type	Description
Multicast Listener Query	General Query, used to learn which multicast addresses have listeners on an attached link. Multicast-Address-Specific Query, used to learn if a particular multicast address has any listeners on an attached link.
Multicast Listener Report	Sent by a host when it joins a multicast group, or in response to a 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.

Multicast Listener Report

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					

Router Advertisements for DNS (RDNSS)

RFC8106: IPv6 Router Advertisement Options for DNS Configuration

- This document specifies IPv6 Router Advertisement (RA) options (called "DNS RA options") to allow IPv6 routers to advertise a list of DNS Recursive Server Addresses and a DNS Search List to IPv6 hosts.
- Why do we need this? (From RFC8106)
- “It is infeasible to manually configure nomadic hosts each time they connect to a different network. While a one-time static configuration is possible, it is generally not desirable on general-purpose hosts such as laptops. For instance, locally defined namespaces would not be available to the host if it were to run its own recursive name server directly connected to the global DNS.”

RDNSS: Why?

- From RFC8106: “It is infeasible to manually configure nomadic hosts each time they connect to a different network. While a one-time static configuration is possible, it is generally not desirable on general-purpose hosts such as laptops. For instance, locally defined namespaces would not be available to the host if it were to run its own recursive name server directly connected to the global DNS.”
- What does this mean?
- If you are at Starbucks with your laptop, you don't want to have to reconfigure so that you can do a google search.
- If you just used one of the big global DNS servers like 1.1.1.1 or 6.6.6.6 and you are inside your company, you wouldn't have your company resources (DNS names) available to you

Why in RA?

- From RFC8106: “The DNS information can also be provided through DHCPv6 [RFC3315] [RFC3736] [RFC3646]. However, access to DNS is a fundamental requirement for almost all hosts, so IPv6 SLAAC cannot stand on its own as an alternative deployment model in any practical network without any support for DNS configuration.”
- What does this mean?
- You could get DNS information using DHCPv6.
- But, if you ONLY have DHCPv6, then you would force networks to run DHCPv6. So, you have to provide another way.
- You can use both DHCPv6 AND RDNSS (take a look at RFC8106)

IPv6-Only Networks

- From RFC8106: “These issues are not pressing in dual-stack networks as long as a DNS server is available on the IPv4 side, but they become more critical with the deployment of IPv6-only networks. As a result, this document defines a mechanism based on DNS RA options to allow IPv6 hosts to perform automatic DNS configuration.”
- What does this mean?
- You are OK as long as you still have IPv4 DNS but once you go to IPv6-only, then what are you going to do?

Recursive DNS Server Option

No.	Time	Source	Destination	Protocol	Info
2...	1922...	fe80::a5b:eff:fea1:835e	ff02::1	ICMPv6	Router Ad


```
> Frame 21933: 182 bytes on wire (1456 bits), 182 bytes captured (1456 bits) on interface 30
> Ethernet II, Src: Fortinet_a1:83:5e (08:5b:0e:a1:83:5e), Dst: IPv6mcast_01 (33:33:00:00:00:01)
> Internet Protocol Version 6, Src: fe80::a5b:eff:fea1:835e, Dst: ff02::1
v Internet Control Message Protocol v6
  Type: Router Advertisement (134)
  Code: 0
  Checksum: 0xa142 [correct]
  [Checksum Status: Good]
  Cur hop limit: 0
  > Flags: 0x00, Prf (Default Router Preference): Medium
  Router lifetime (s): 1800
  Reachable time (ms): 0
  Retrans timer (ms): 0
  > ICMPv6 Option (Prefix information : 2001:470:1f0b:16b0::/64)
v ICMPv6 Option (Recursive DNS Server 2606:4700:4700::1111 2620:fe::fe)
  Type: Recursive DNS Server (25)
  Length: 5 (40 bytes)
  Reserved
  Lifetime: 120
  Recursive DNS Servers: 2606:4700:4700::1111
  Recursive DNS Servers: 2620:fe::fe
  > ICMPv6 Option (DNS Search List Option weberlab.de)
  > ICMPv6 Option (MTU : 1500)
  > ICMPv6 Option (Source link-layer address : 08:5b:0e:a1:83:5e)
```



Router Advertisement
To show DNS servers



No.	Time	Source	Destination	Protocol	Info
2...	1922...	fe80::a5b:eff:fea1:835e	ff02::1	ICMPv6	Router Advert

- > Frame 21933: 182 bytes on wire (1456 bits), 182 bytes captured (1456 bits) on
- > Ethernet II, Src: Fortinet_a1:83:5e (08:5b:0e:a1:83:5e), Dst: IPv6mcast_01 (:
- > Internet Protocol Version 6, Src: fe80::a5b:eff:fea1:835e, Dst: ff02::1
- ∨ Internet Control Message Protocol v6

Type: Router Advertisement (134) 

Code: 0

Checksum: 0xa142 [correct]

[Checksum Status: Good]

Cur hop limit: 0

Router Advertisement

To show DNS servers

> Flags: 0x00, Prf (Default Router Preference): Medium

Router lifetime (s): 1800

Reachable time (ms): 0

Retrans timer (ms): 0

> ICMPv6 Option (Prefix information : 2001:470:1f0b:16b0::/64)


∨ ICMPv6 Option (Recursive DNS Server 2606:4700:4700::1111 2620:fe::fe)

Type: Recursive DNS Server (25)

Length: 5 (40 bytes)

Reserved

Lifetime: 120

Recursive DNS Servers: 2606:4700:4700::1111 

Recursive DNS Servers: 2620:fe::fe

> ICMPv6 Option (DNS Search List Option weberlab.de)

> ICMPv6 Option (MTU : 1500)

> ICMPv6 Option (Source link-layer address : 08:5b:0e:a1:83:5e)

No.	Time	Source	Destination	Protocol	Info
2...	1922...	fe80::a5b:eff:fea1:835e	ff02::1	ICMPv6	Router Advertisement
<ul style="list-style-type: none"> > Frame 21933: 182 bytes on wire (1456 bits), 182 bytes captured (1456 bits) on interf > Ethernet II, Src: Fortinet_a1:83:5e (08:5b:0e:a1:83:5e), Dst: IPv6mcast_01 (33:33:00 > Internet Protocol Version 6, Src: fe80::a5b:eff:fea1:835e, Dst: ff02::1 <ul style="list-style-type: none"> <ul style="list-style-type: none"> Type: Router Advertisement (134) Code: 0 Checksum: 0xa142 [correct] [Checksum Status: Good] Cur hop limit: 0 > Flags: 0x00, Prf (Default Router Preference): Medium Router lifetime (s): 1800 Reachable time (ms): 0 Retrans timer (ms): 0 > ICMPv6 Option (Prefix information : 2001:470:1f0b:16b0::/64) > ICMPv6 Option (Recursive DNS Server 2606:4700:4700::1111 2620:fe::fe) <ul style="list-style-type: none"> Type: DNS Search List Option (31) Length: 3 (24 bytes) Reserved Lifetime: 120 Domain Names: weberlab.de Padding > ICMPv6 Option (MTU : 1500) > ICMPv6 Option (Source link-layer address : 08:5b:0e:a1:83:5e) 					



DNS Search List

The list of DNS suffix domain names used by IPv6 hosts when they perform DNS query searches for short, unqualified domain names.

What can go wrong?

- RFC8106 “For the RDNSS option, an attacker could send an RA with a fraudulent RDNSS address, misleading IPv6 hosts into contacting an unintended DNS server for DNS name resolution. Also, for the DNSSL option, an attacker can let IPv6 hosts resolve a hostname without a DNS suffix into an unintended host's IP address with a fraudulent DNSSL. These attacks are similar to ND attacks specified in [RFC4861] that use Redirect or Neighbor Advertisement messages to redirect traffic to individual addresses of malicious parties.”
- What does this mean?
- Someone could spoof a bad DNS name or address and send you to the wrong place. (But this can happen with regular DNS also.)

Summary

- Stateless autoconfiguration has benefits
- Creating an IPv6 address automatically is complicated
- Neighbor Discovery and Multicast Listener Discovery are very important to understand

Classes

- Introduction to IPv6 : Feb 4, 2021 ✓
- Lab: IPv6 basics : Feb 11, 2021 ✓
- Neighbor Discovery: March 4, 2021 ✓
- Lab: Neighbor Discovery: March 18, 2021
- IPv6 Address Planning: April 8, 2021
- Lab: IPv6 Address Planning: April 15, 2021
- IPv6 Transition Mechanisms: May 6, 2021
- Lab: IPv6 Transition Mechanisms: May 13, 2021

- DHCPv6: June 3, 2021
 - Lab: DHCPv6: June 10, 2021
 - IPv6 and Cloud: June 17, 2021
 - Lab: IPv6 and Cloud: June 24, 2021
 - Introduction to IPv6 Security July 8, 2021
- The next sessions are sponsored by a generous grant from ARIN.
- Trace Reading: August 12, 2021
 - Troubleshooting: August 19, 2021

Next Steps

- Will publicize initial results of discussions with enterprises and universities
- Initial discussions with folks signed up in Call for Participation
- Stay tuned!
- IETF is next week. You can attend remotely.

Questions?

Contact:

nalini.elkins@insidestack.com
president@industry.netcouncil.org

More webinars at:

industry.netcouncil.org