CEN
SINCE 2000
25 YEARS
CONNECTING CT COMMUNITIES

IPv6 Deep Dive: Won't You Be My Neighbor?

CEN CONNECT
2025 Member Conference

About Me

Jason Healy

Suffield Academy

- 420 Students
- High School (grades 9-12)
- Boarding and Day students

Director of Technology:

- Manage IT department
- Core network and wireless
- Teach Computer Science

IPv4 Stickers

- Limited supply available
- Feel free to stick to vendor products that do not support IPv6
- First person to (legitimately) stick one to a vendor booth outside gets my drink ticket

Review Concepts

IPv6 Micro-review

- 128-bit addresses, represented as 8 groups of 4 hex digits:

  `2001:0db8:1234:5678:0000:90ab:cdef`

- Can "zero-compress" and remove leading zeros:

  `2001:db8:1234:5678::90ab:cdef`

- Overwhelming default of 64 bits for network, 64 bits for address:

  `2001:db8:1234:5678::/64`

- Multiple addresses per node are common
- Link-local (`fe80::/10`) addresses automatically assigned

Layer 2: Ethernet

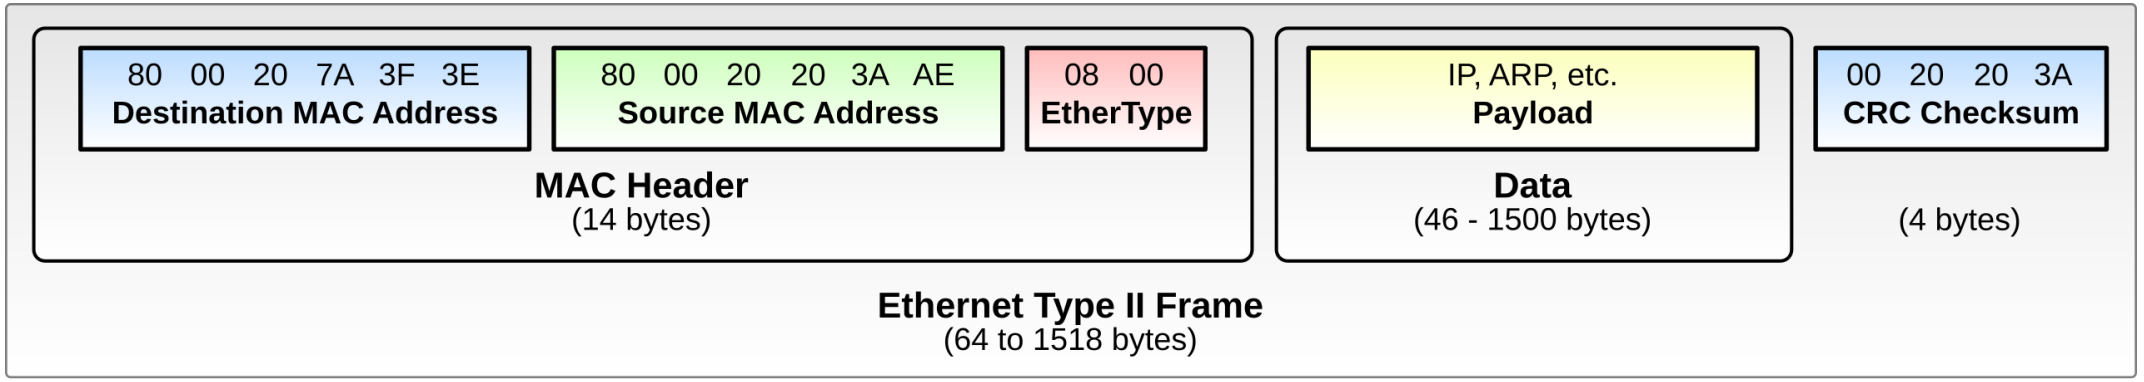

https://en.wikipedia.org/wiki/Ethernet_frame#/media/File:Ethernet_Type_II_Frame_format.svg

- Local networking only, simple protocol
- Nodes must be on same local segment (no routing)
- MAC/LLA addresses are used for src/dst
- LSB of 1 in octet 1 of dst addr is explicitly multicast
- Frames are (potentially) seen by all nodes on segment
- EtherType defines payload type, higher-level protocol data contained inside frame data portion

Layer 3: IPv6

- Local or wide-area networking, as determined by subnet mask

```
2001:db8:1337:dead:0000:0000:0000:0001/64 (node)
2001:db8:1337:dead:cafe:babe:8bad:f00d/64 (on)
2001:db8:d00d:000d:0000:0000:0000:0001/64 (off)
```

- On-subnet packets are delivered directly using lower protocols (e.g., Ethernet)
- Off-subnet packets are forwarded to a router for delivery
- Router must be on the local subnet
- All packets require layer-2 destination!

Neighbor Discovery

Problem Statement

IPv6 packets must be embedded in Ethernet frames before they are sent.

Given an IPv6 destination address, how do we determine the destination MAC address for the corresponding Ethernet packet?

Neighbor Discovery

Neighbor Solicitations – ICMP 135

- Equivalent of ARP "who-has"

- Target Address: <IPv6 Address>
- Option: Source LLA: <MAC Address>

- MAC of the solicited node is unknown, so we use a special **Solicited-Node multicast address** as the destination:
- Use link-local multicast prefix `ff02::ff` and append last 3 bytes (6 nybbles) of target IPv6 address
- Use `33:33:ff` + last 3 bytes for destination MAC

Example: `2001:db8::1234:5678` would map to

IPv6 `ff02::ff34:5678` and MAC `33:33:ff:34:56:78`

Neighbor Advertisements – ICMP 136

- Equivalent of ARP "is-at"

- Target Address: <IPv6 Address> (echoes NS target)
- Option: Target LLA: <MAC Address>
- Flag: Router (1 if target node is a router)
- Flag: Solicited (1 if target node is responding to an NS)
- Flag: Override (1 if target node has changed its LLA)

- Target node knows the MAC address of requestor from the solicitation, so can send traffic via unicast

NS / NA Example

How can I reach
2001:db8:1337::b?

Node: A
MAC: DE:AD:BE:EF:12:0A
IPv6: 2001:db8:1337::a

Neighbor Solicitation: ICMPv6 135
IPv6 dst: ff02::ff00:000b
MAC dst: 33:33:FF:00:00:0B
Target Address: 2001:db8:1337::b
Option: Source LLA: DE:AD:BE:EF:12:0A

Node: B
MAC: DE:AD:BE:EF:34:0B
IPv6: 2001:db8:1337::b

Neighbor Advertisement: ICMPv6 136
IPv6 dst: 2001:db8:1337::a
MAC dst: DE:AD:BE:EF:12:0A
Target Address: 2001:db8:1337::b
Option: Target LLA: DE:AD:BE:EF:34:0B
Router=0 Solicited=1 Override=0

Router Discovery

Router Discovery Summary

- Considered part of Neighbor Discovery (RFC 4861)
- Used to find routers on the subnet that can forward traffic
- Provides information to help nodes auto-assign addresses
- Provides other network options to nodes, replacing (some) functionality of DHCPv4 (e.g., DNS server address, domain name suffix)
- Router Advertisements can be solicited or unsolicited (periodically broadcast)

Router Solicitations – ICMP 133

- Option: Source LLA: <MAC Address>
- Sent to the **All-Routers** multicast address: `ff02::2`

Router Advertisements – ICMP 134

- Option: Source LLA: <MAC Address>
- Flag: M (**managed** address configuration, e.g., DHCPv6)
- Flag: O (**other** configuration provided by DHCPv6, e.g., NTP server, DNS server)
- (Other flags not discussed: Preference, Home Agent, ND Proxy)
- Option: Prefix
  - IPv6 subnet
  - lifetime
  - etc.
- Other Options: rDNS server, DNS suffix, PREF64, etc.

RA in Wireshark

- Example of a basic RA
- No managed config

- Defines rDNS server
- Defines DNS suffix

- Lists single subnet prefix, requests automatic address config (SLAAC)

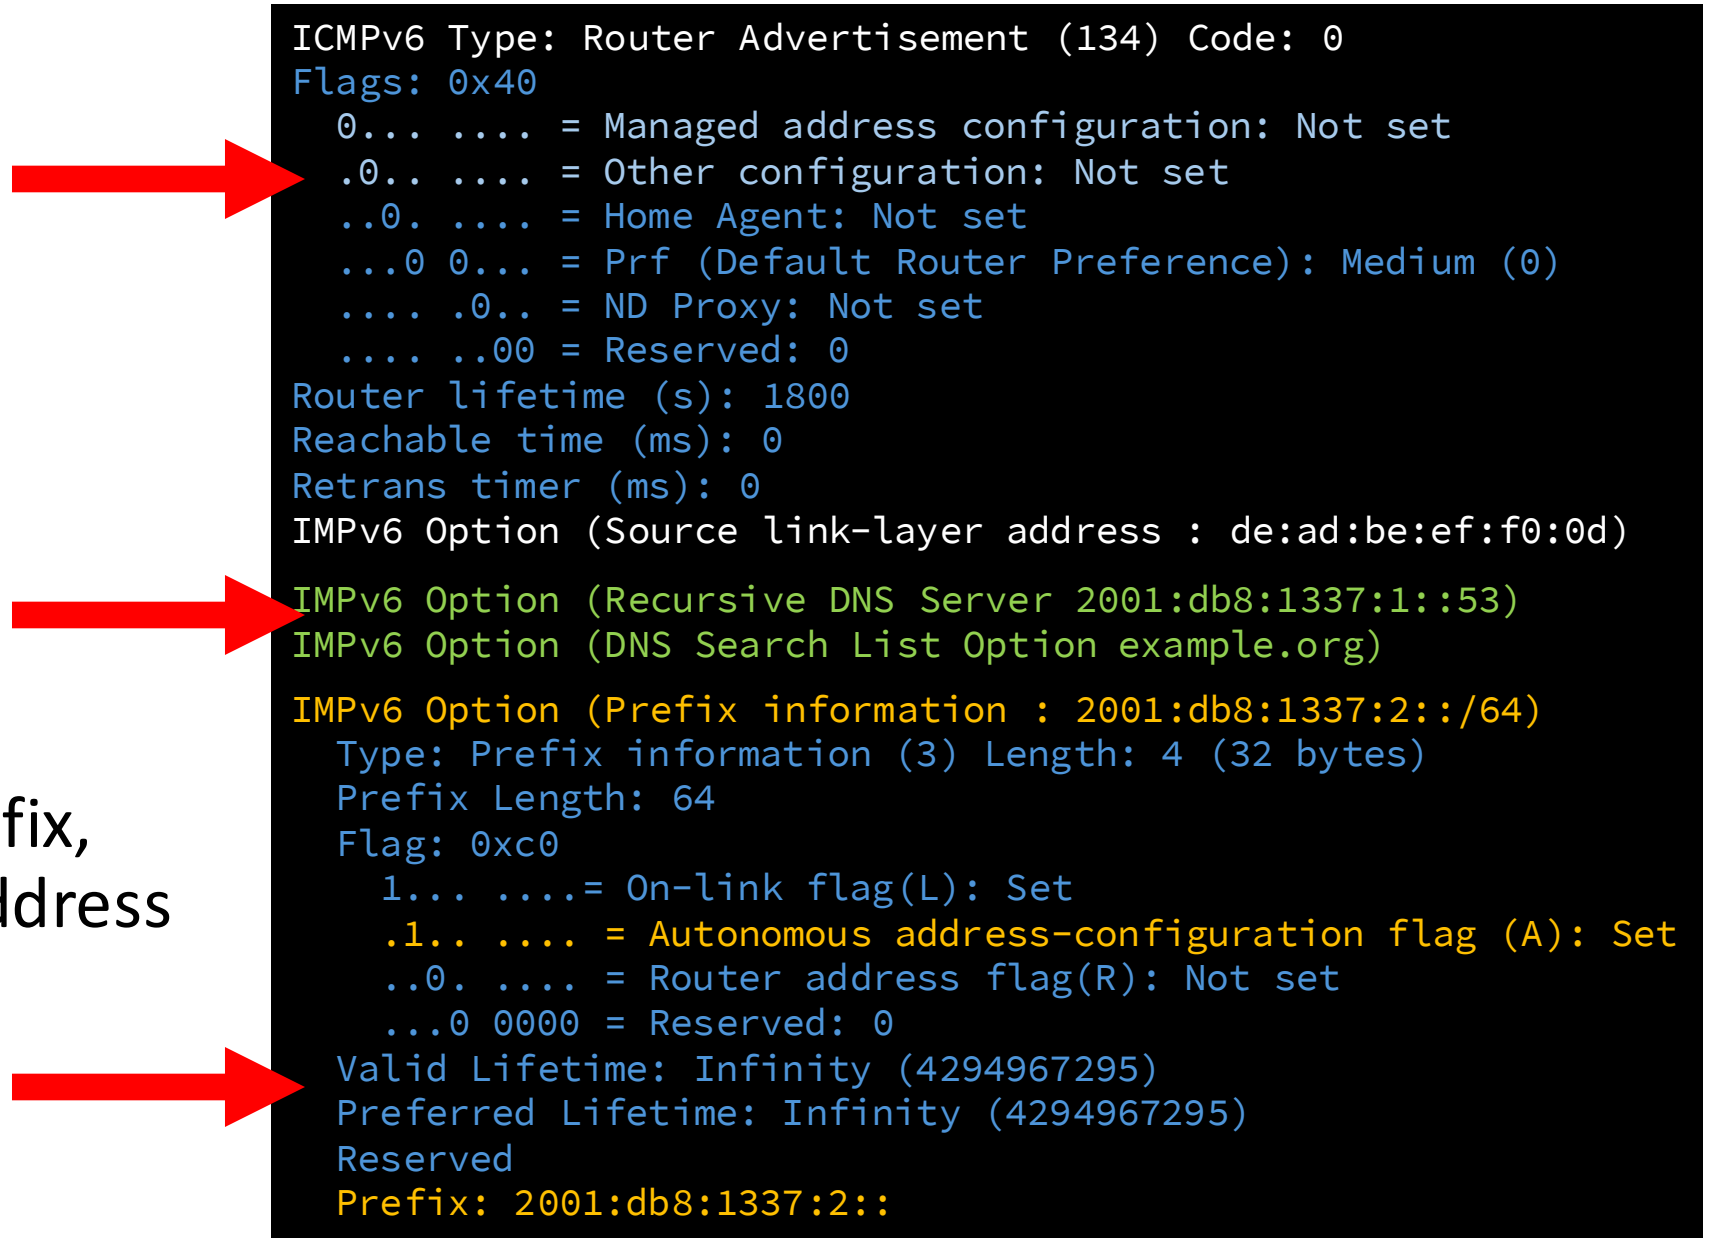

Duplicate Address Detection

DAD

- Straightforward application of Neighbor Discovery
- Node sends NS for the address it wants to use
- If NA is received, node must choose a different address
- If no NA is received, address remains a candidate for use
- The **Unspecified Address** (all-zeros, "**::**") is used as the source address for the NS

SLAAC

Stateless Address Autoconfiguration

State **L**ess **A**ddress **A**uto **C**onfiguration (**SLAAC**)

- Allows a node to generate a Global Unicast Address (GUA) without requiring any stateful protocols (e.g., DHCP)
- Combines all the ND items we've discussed:
  - RA provides next-hop router for non-link-local packets
  - RA gives us a prefix to choose addresses from
  - DAD confirms our desired address isn't already in use

Let's see how SLAAC brings together everything we've learned

SLAAC Typical Host Flow – Link Up

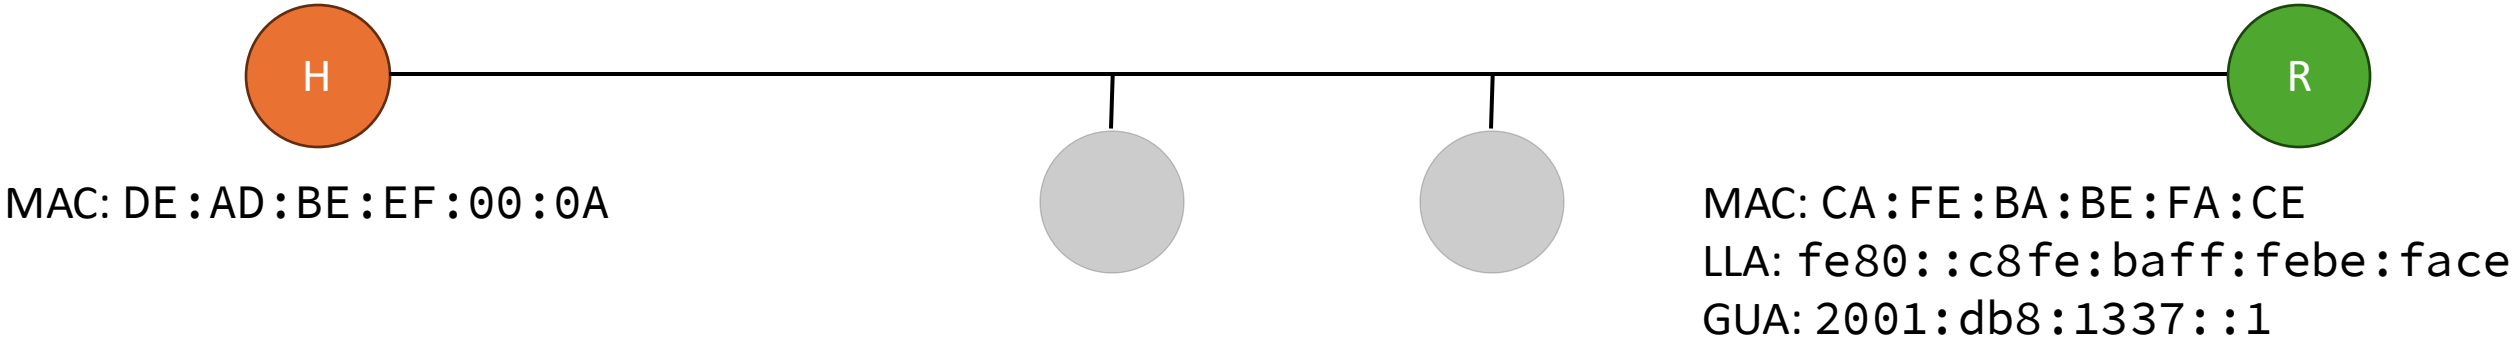

- Our host "H" is joining an IPv6 subnet; link has just become active
- Router "R" is present on the subnet as well
- Other hosts (grey) are present

SLAAC Typical Host Flow – LL ADDR

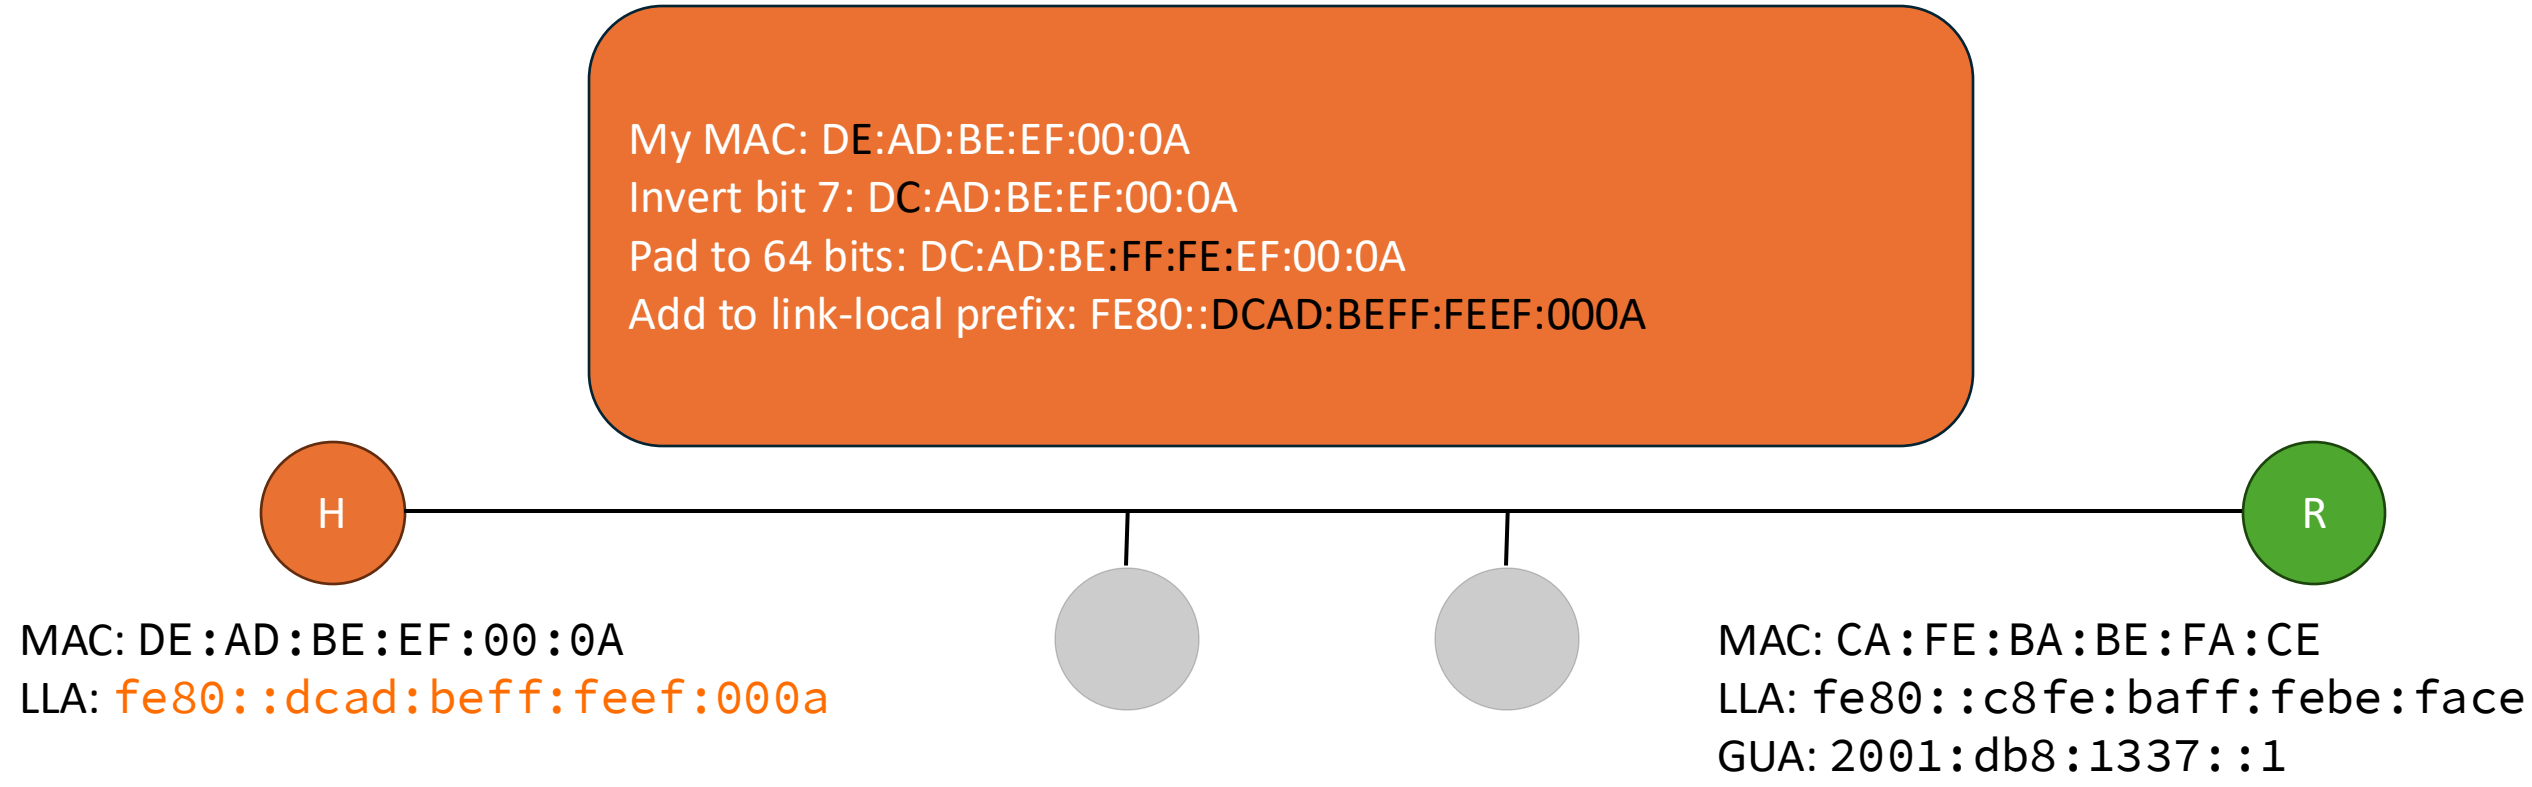

- Host begins by generating an interface ID (IID), e.g., EUI-64
- Appends this to the well-known link-local prefix `fe80::`
- Assigns this as a tentative link-local address

SLAAC Typical Host Flow – LL DAD

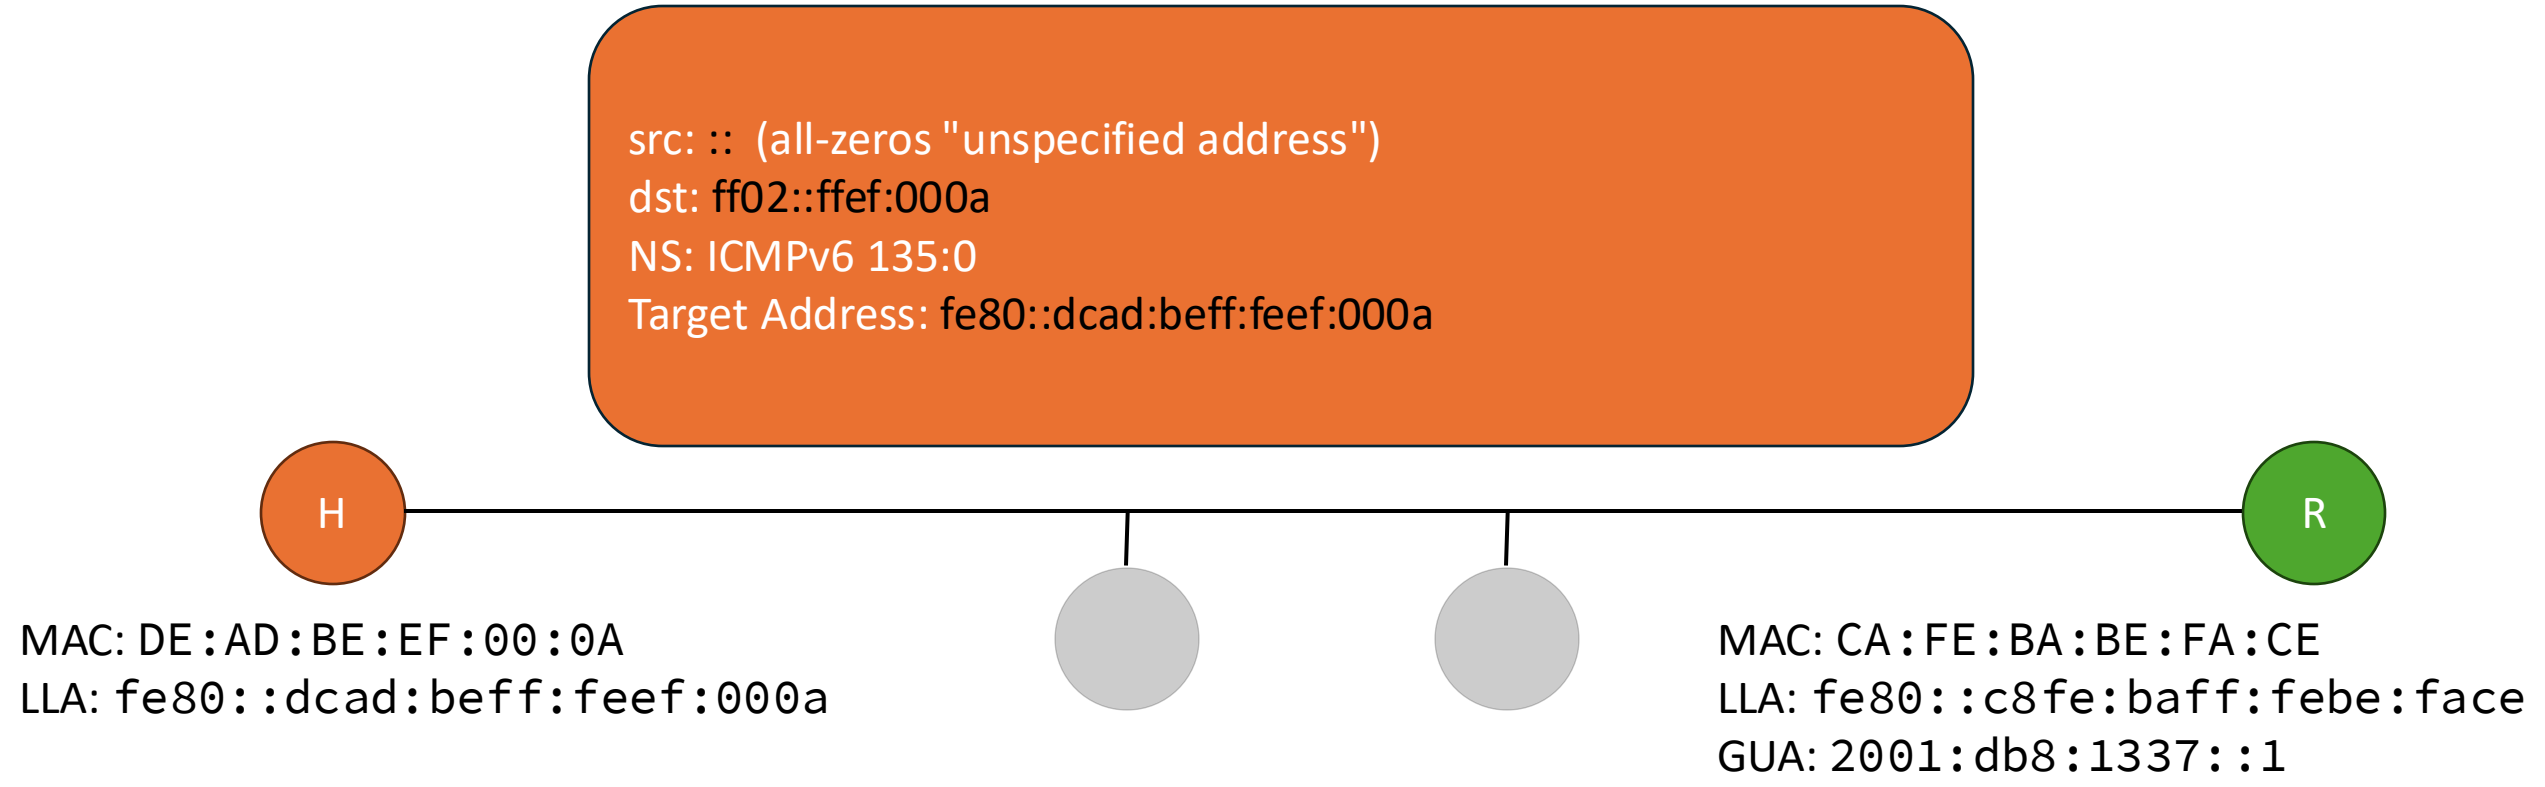

- Host performs DAD on its new link-local address
- Sends multiple NS to see if any other node claims address
- If duplicate detected, error out; otherwise proceed

SLAAC Typical Host Flow – RS

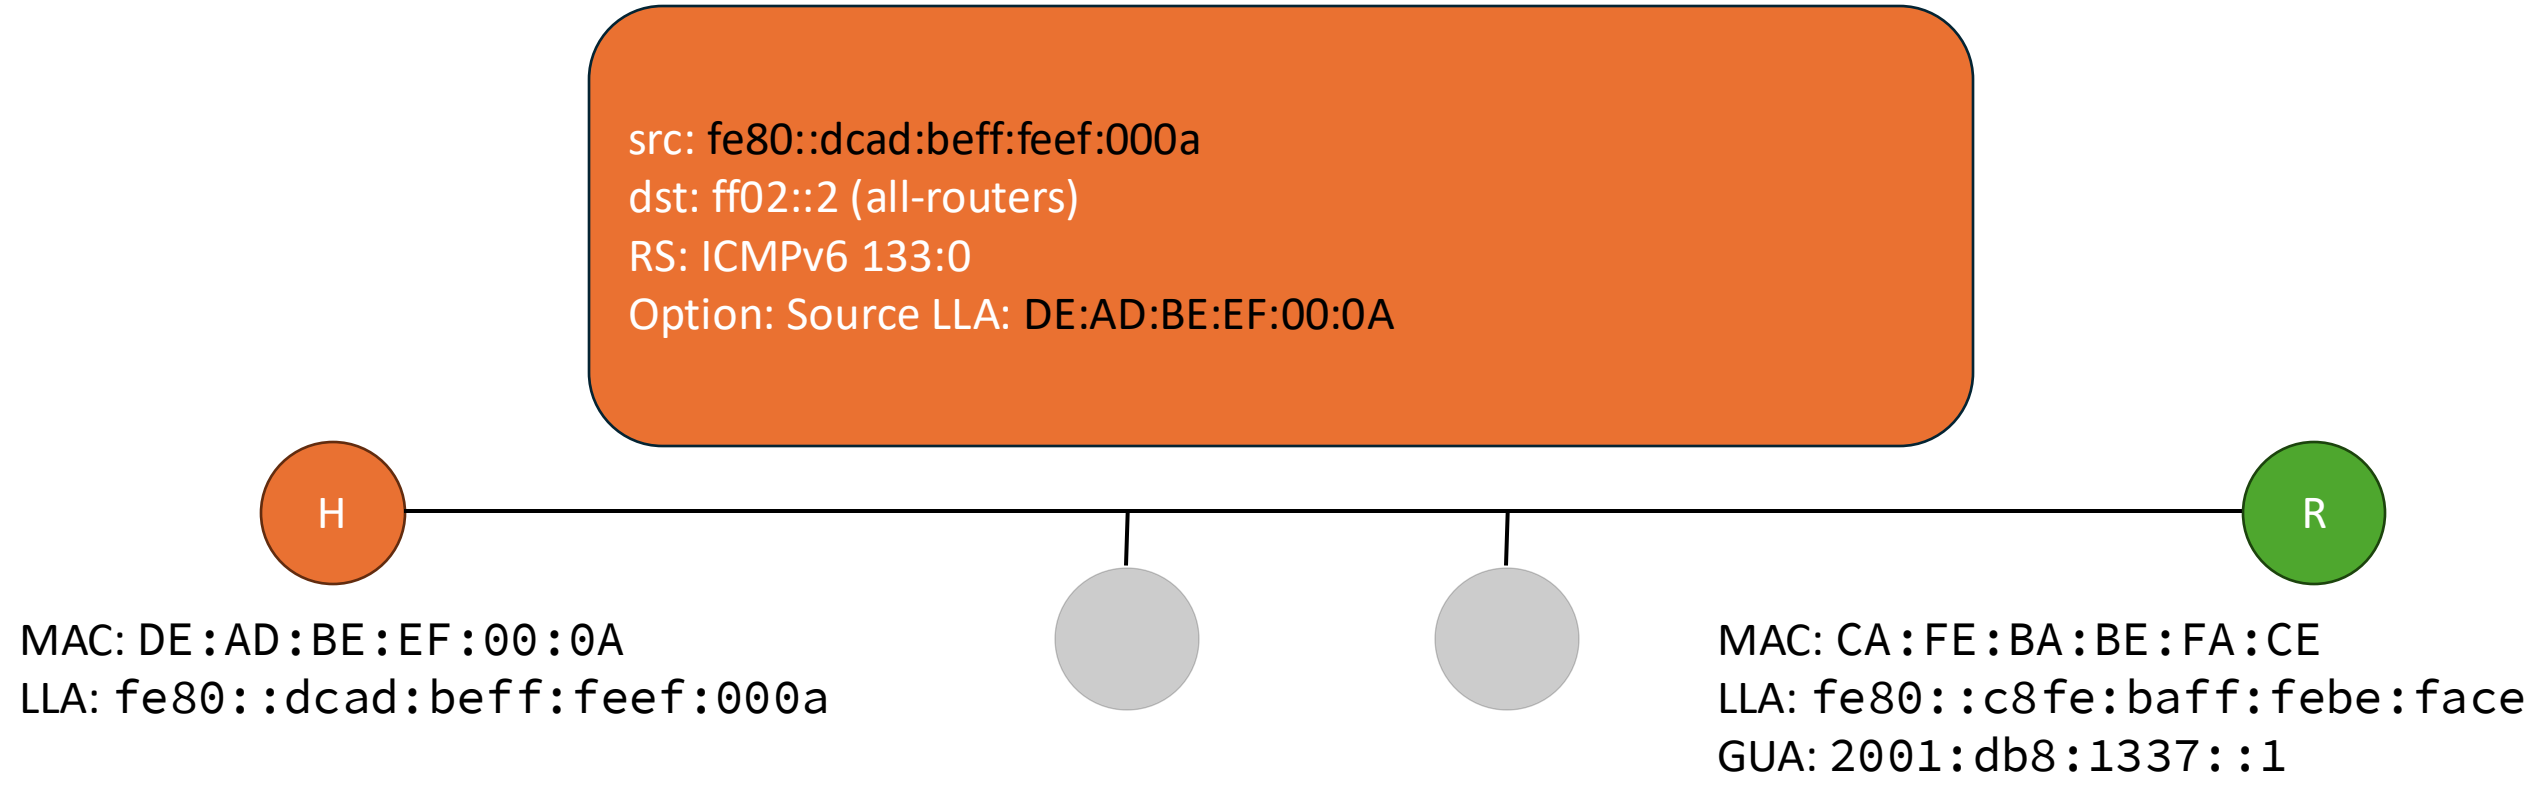

- Host sends a multicast Router Solicitation packet
- Host includes its MAC address to receive a reply
- (This step can be done in parallel with DAD for link-local)

SLAAC Typical Host Flow – RA

src: fe80::c8fe:baff:febe:face
dst: ff02::1 (all-nodes)
RA: ICMPv6 134:0
Prefix: 2001:db8:1337::/64 ; Autonomous auto-config=1
DNS Server: 2001:db8::53
DNS Suffix: example.org
(Other options not shown: preference, lifetime, etc)

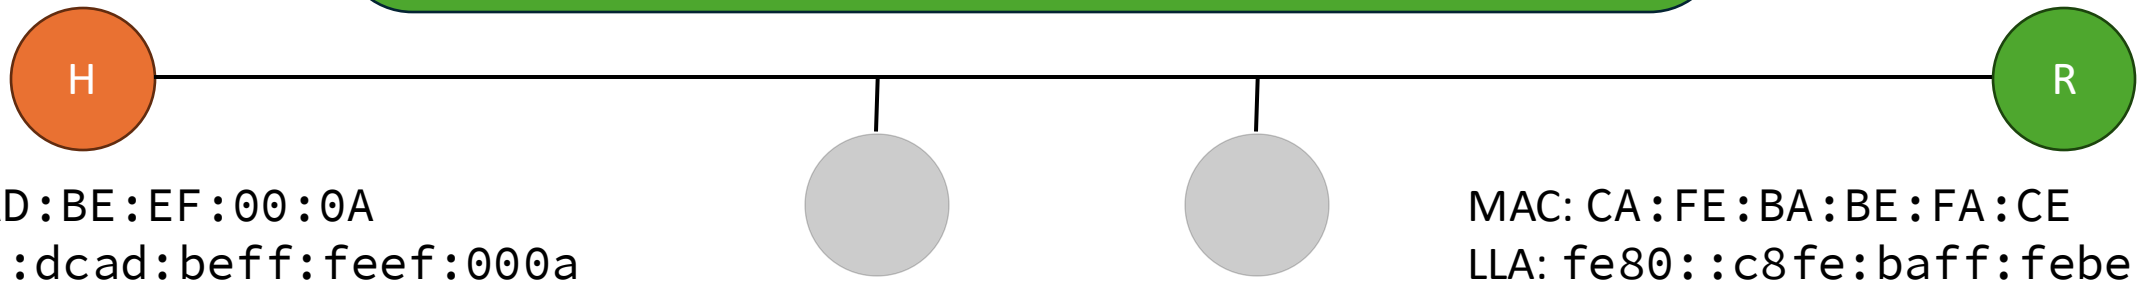

MAC: DE:AD:BE:EF:00:0A
LLA: fe80::dcad:beff:feef:000a

MAC: CA:FE:BA:BE:FA:CE
LLA: fe80::c8fe:baff:febe:face
GUA: 2001:db8:1337::1

- Router gets solicitation, responds with Router Advertisement
- RA can be unicast or multicast
- RA contains global options, plus one or more prefixes

SLAAC Typical Host Flow – RA (host)

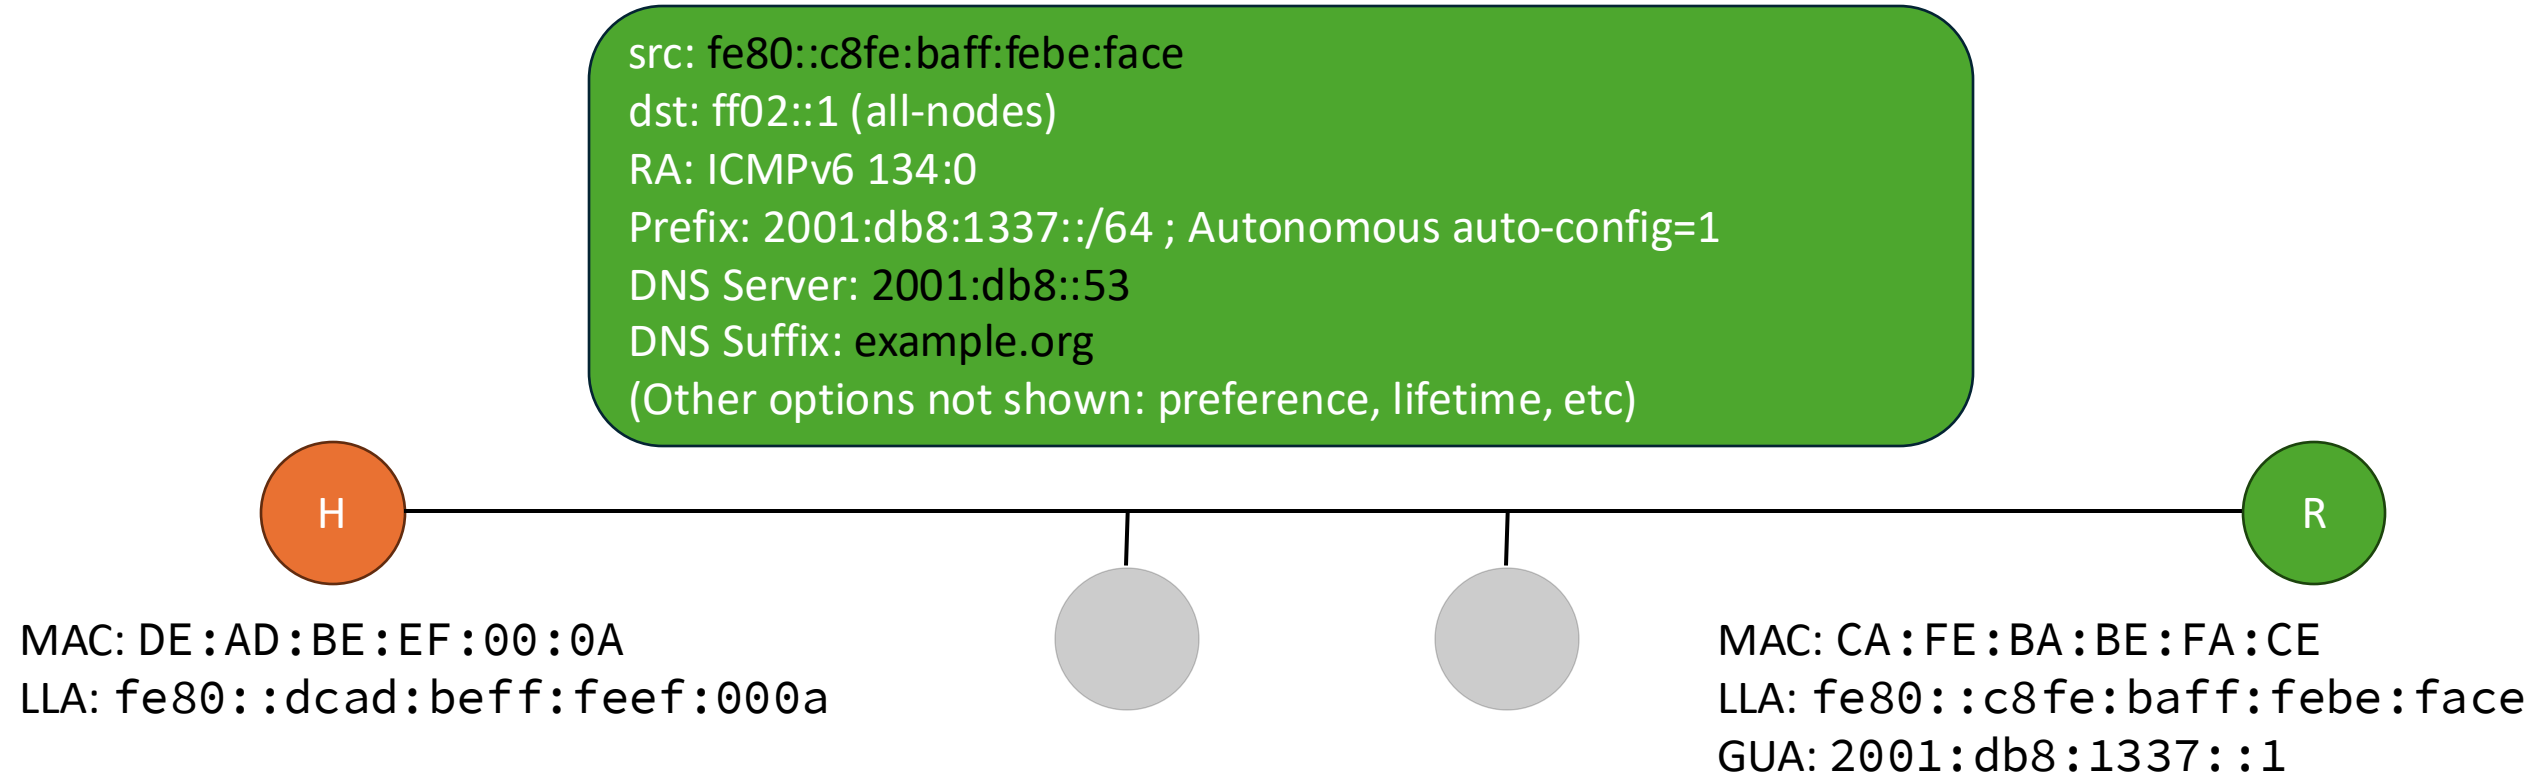

- Host receives RA
- Adds LLA of router to its route table (possibly as default)
- Processes other global options (DNS, etc)

SLAAC Typical Host Flow – RA (prefix)

src: fe80::c8fe:baff:febe:face
dst: ff02::1 (all-nodes)
RA: ICMPv6 134:0
Prefix: 2001:db8:1337::/64 ; Autonomous auto-config=1
DNS Server: 2001:db8::53
DNS Suffix: example.org
(Other options not shown: preference, lifetime, etc)

MAC: DE:AD:BE:EF:00:0A
LLA: fe80::dcad:beff:feef:000a
GUA: 2001:db8:1337:aabb:ccdd:eeff:1234

MAC: CA:FE:BA:BE:FA:CE
LLA: fe80::c8fe:baff:febe:face
GUA: 2001:db8:1337::1

- Host processes each Prefix option
- If autonomous auto-config is set, generate a IID for this subnet and combine it with the prefix to form a candidate GUA

SLAAC Typical Host Flow – GA DAD

src: :: (all-zeros "unspecified address")
dst: ff02::ffff:1234
NS: ICMPv6 135:0
Target Address: 2001:db8:1337:aabb:ccdd:eeff:1234

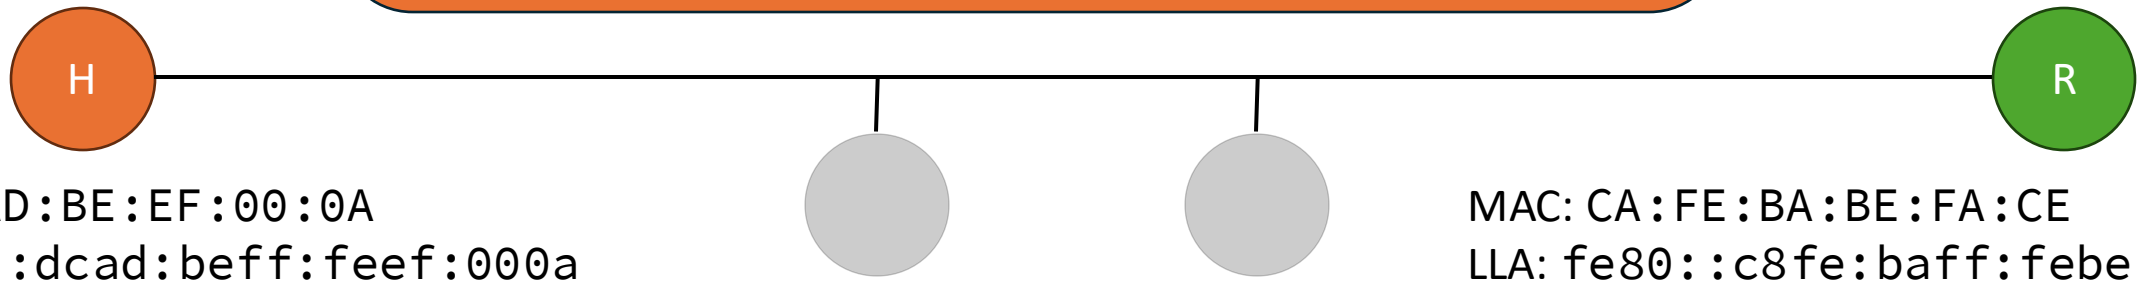

MAC: DE:AD:BE:EF:00:0A
LLA: fe80::dcad:beff:feef:000a
GUA: 2001:db8:1337:aabb:ccdd:eeff:1234

MAC: CA:FE:BA:BE:FA:CE
LLA: fe80::c8fe:baff:febe:face
GUA: 2001:db8:1337::1

- Host performs DAD on this new address
- Multiple NS packets sent to confirm no answer
- If NA received, generate a new address and try again
- If no NA received, address is good

SLAAC Typical Host Flow – Done

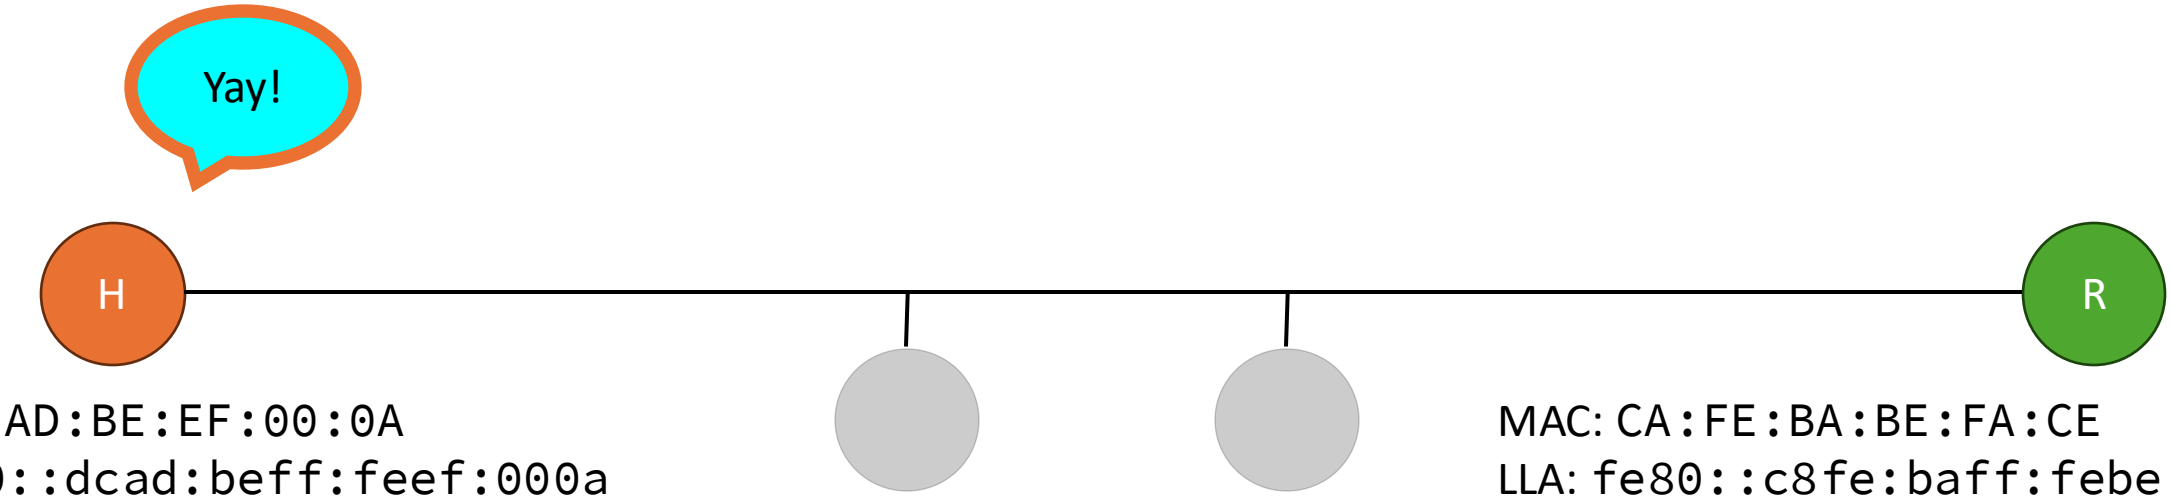

MAC: DE:AD:BE:EF:00:0A
LLA: fe80::dcad:beff:feef:000a
GUA: 2001:db8:1337:aabb:ccdd:eeff:1234

MAC: CA:FE:BA:BE:FA:CE
LLA: fe80::c8fe:baff:febe:face
GUA: 2001:db8:1337::1

- Host is now minimally configured
- Can configure additional (temporary/private) addresses
- Can use NS to discover other nodes on the network

Q & A

Bonus Topics

If we have time...

Firewall Warning

- Not the same as IPv4 + ARP; all traffic is IP now
- Don't block all ICMPv6
- Don't block link-local addresses (`fe80::/10`)
- Don't block subnet multicast addresses (`ff02::/16`)
- You may prevent some/all address resolution
- Connectivity may be spotty (as caches timeout)

Multicast Issues

- IPv6 ND relies heavily on multicast groups
- Degrades gracefully to broadcast on Ethernet

- Does NOT play well with shared network segments:
  - Per-MAC VLAN on single port
  - Multiple VLANs on single PSK Wi-Fi SSID

- Can be issues with low-power devices (see RFC 7772)

ND Security Issues

Just like ARP, ND suffers from too much trust; see:
https://datatracker.ietf.org/doc/draft-ietf-v6ops-nd-considerations/

- Cache Poisoning (spoofing NAs)
- Cache Exhaustion (sending too many NAs)
- DAD DoS (answer every DAD attempt)
- Forged RAs

- Some vendors have mitigations (RA guard), similar to ones for DHCPv4/ARP

Neighbor Cache Entries

Neighbor Cache Entries Overview

- Once we've resolved the LLA of our target, need bind layer 3 (IP) address to layer 2 (MAC)
- Keep bindings in Neighbor Cache (formerly "ARP table")
- Unlike ARP, several states exist for each NCE
- The states determine what to do when a cache entry ages, stops/starts responding, or changes LLA
- State names and terminology are defined in RFC 4681 and followed by vendors in their debugging tools:
  - `ip -6 neighbor show`
  - `ndp -a`
  - `netsh`

Neighbor Cache Entry States

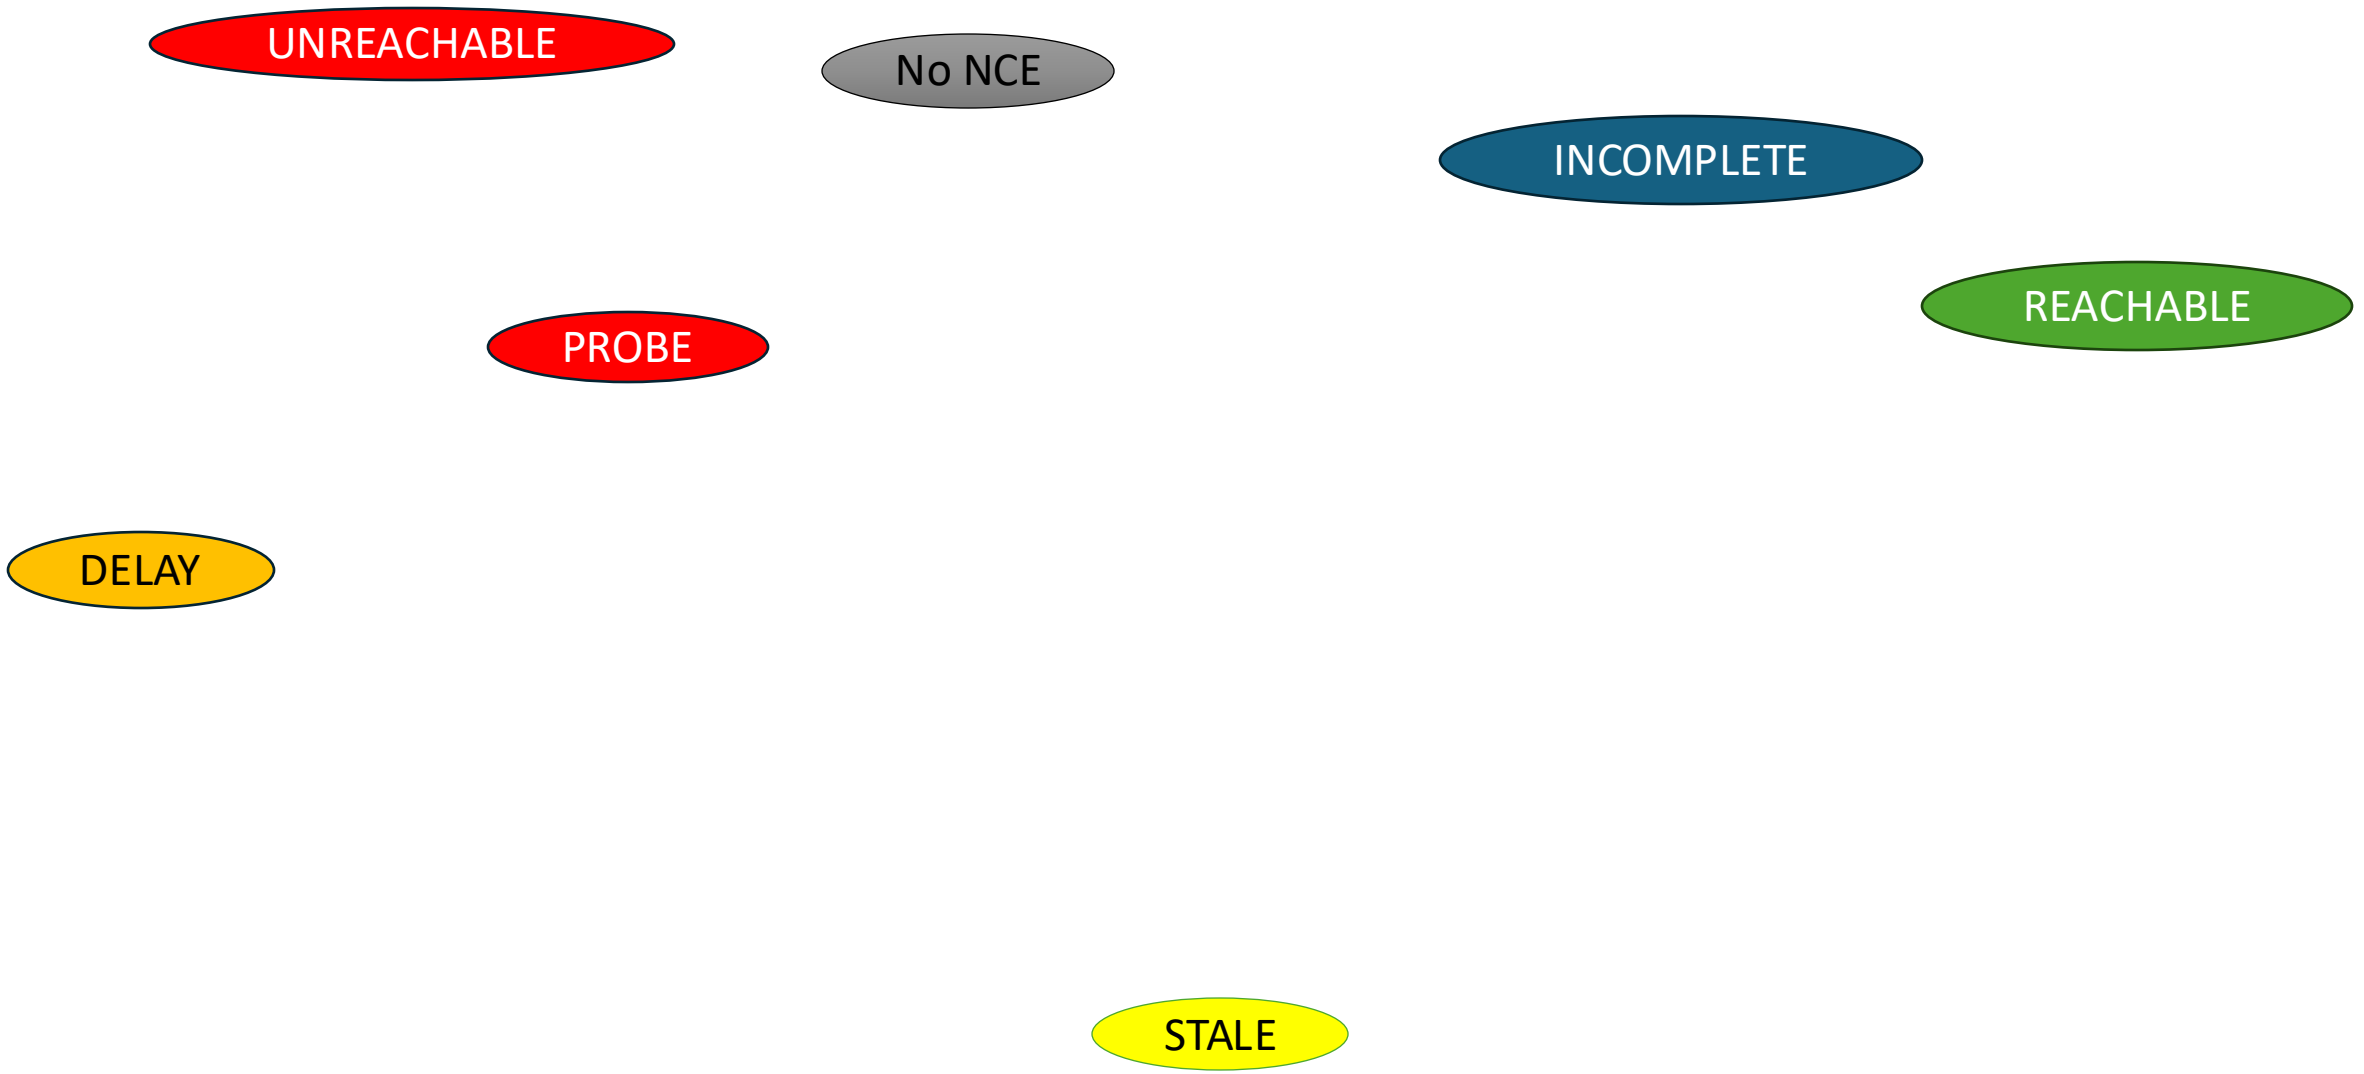

Neighbor Cache Entry States (N)

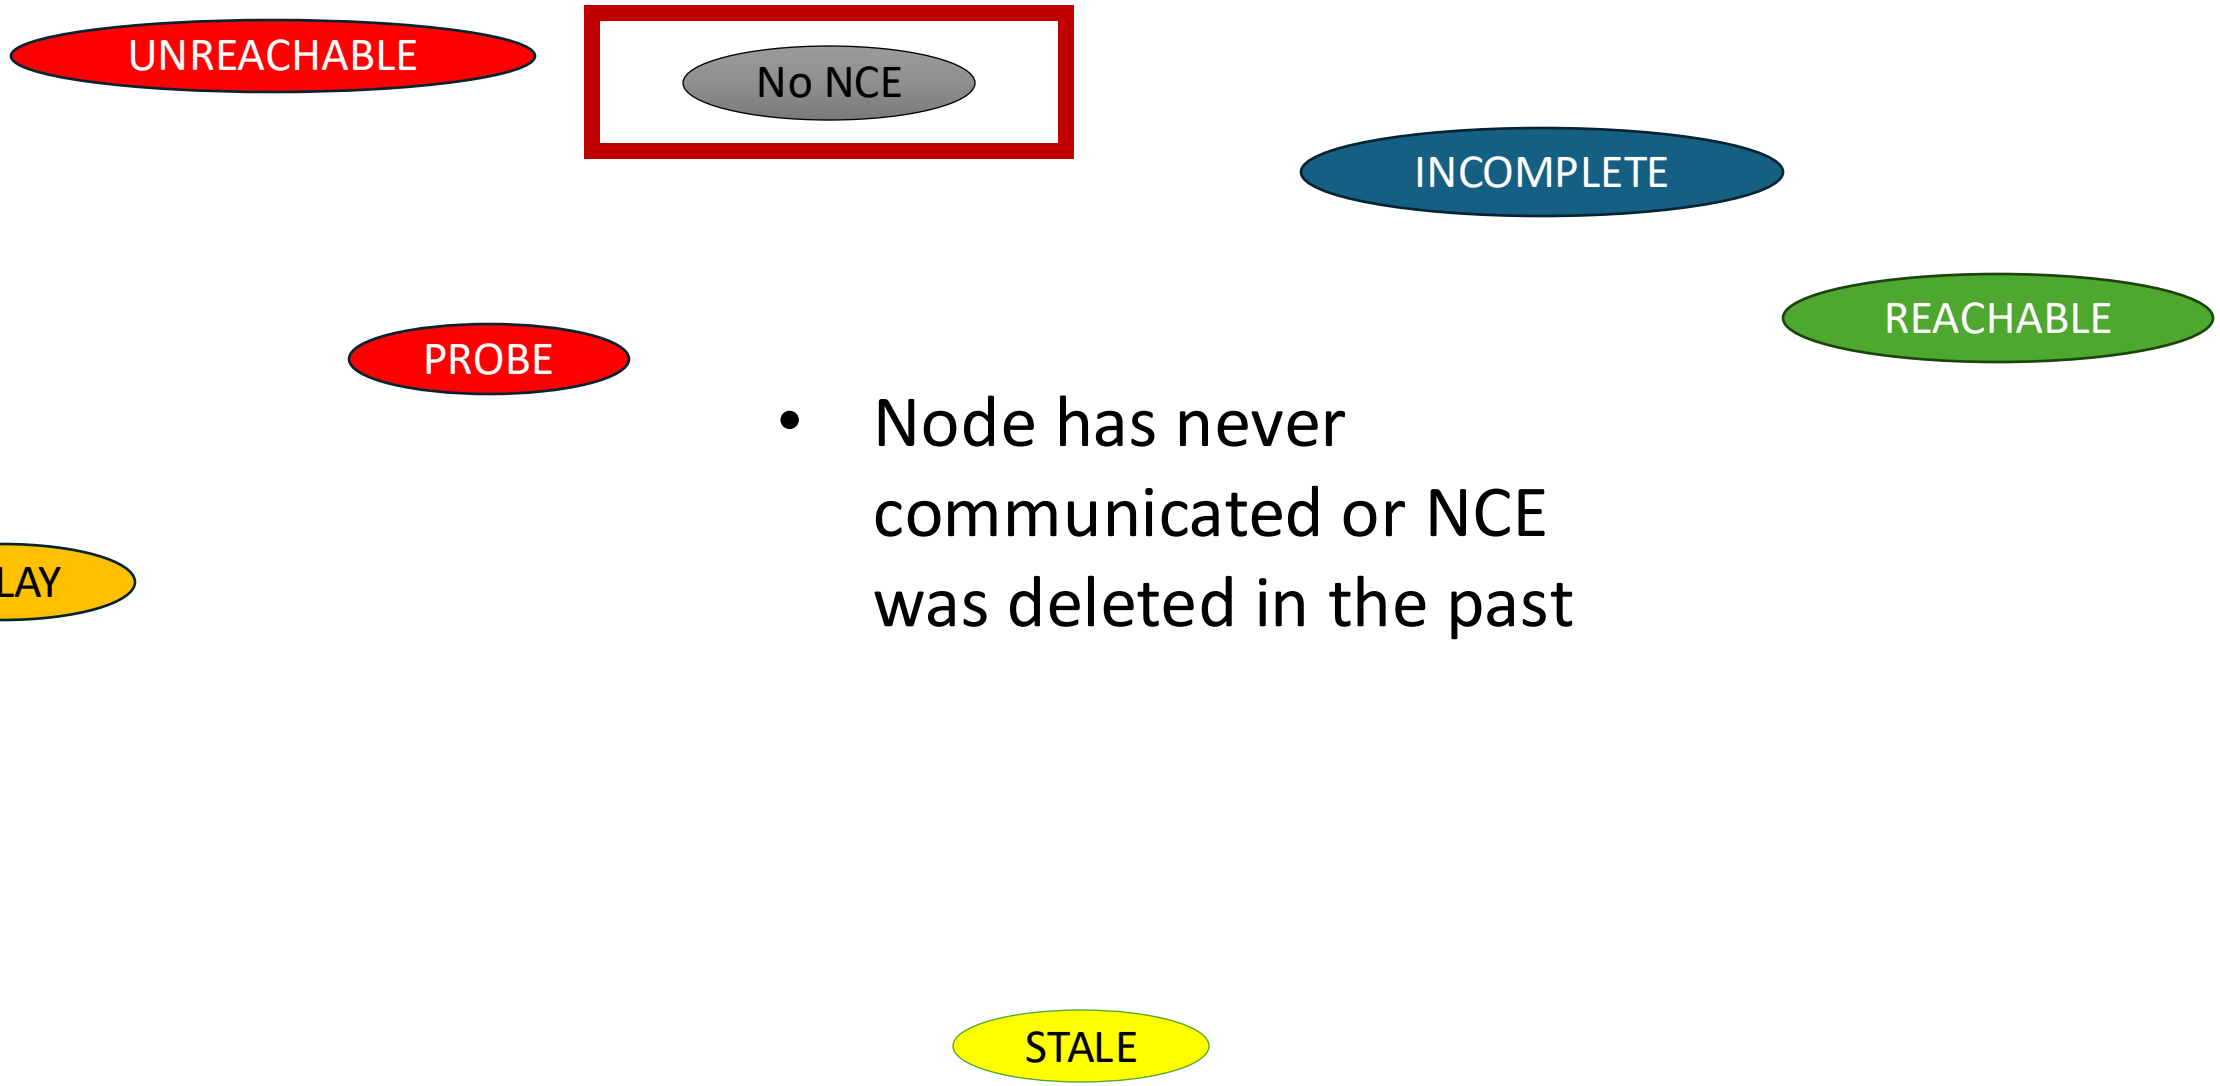

- Node has never communicated or NCE was deleted in the past

Neighbor Cache Entry States (I)

UNREACHABLE

No NCE

INCOMPLETE

REACHABLE

PROBE

- NS has been sent
- Waiting for NA

DELAY

STALE

Neighbor Cache Entry States (R)

UNREACHABLE

No NCE

INCOMPLETE

REACHABLE

PROBE

DELAY

STALE

- Positive confirmation received recently
- No special treatment when sending packets

Neighbor Cache Entry States (S)

UNREACHABLE

No NCE

INCOMPLETE

REACHABLE

PROBE

- No recent packets
  OR
- Change of LLA

DELAY

- *Perfectly normal state!*

STALE

Neighbor Cache Entry States (D)

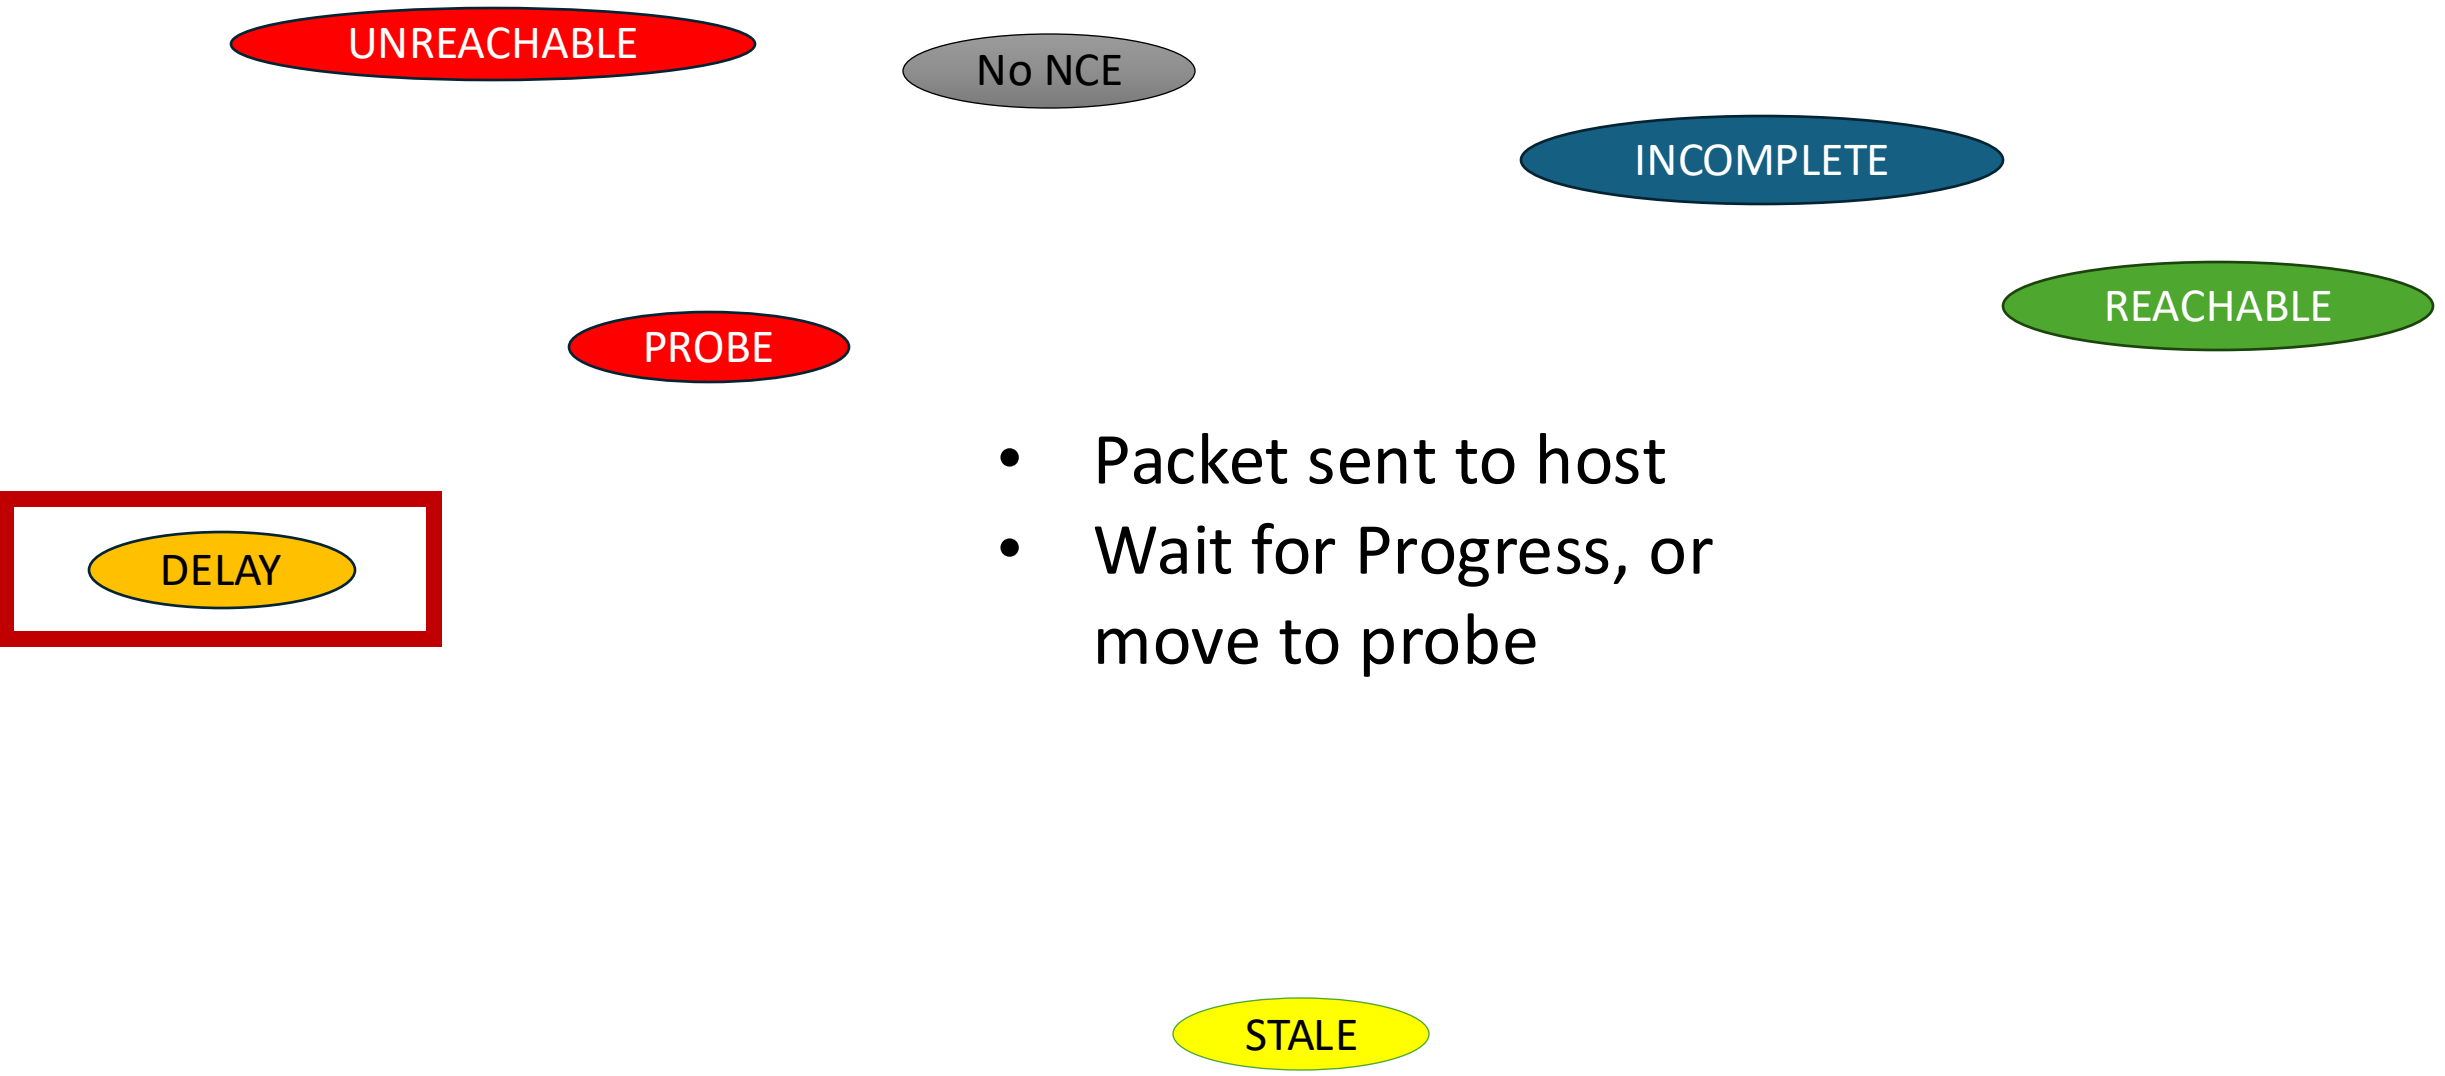

- Packet sent to host
- Wait for Progress, or move to probe

Neighbor Cache Entry States (P)

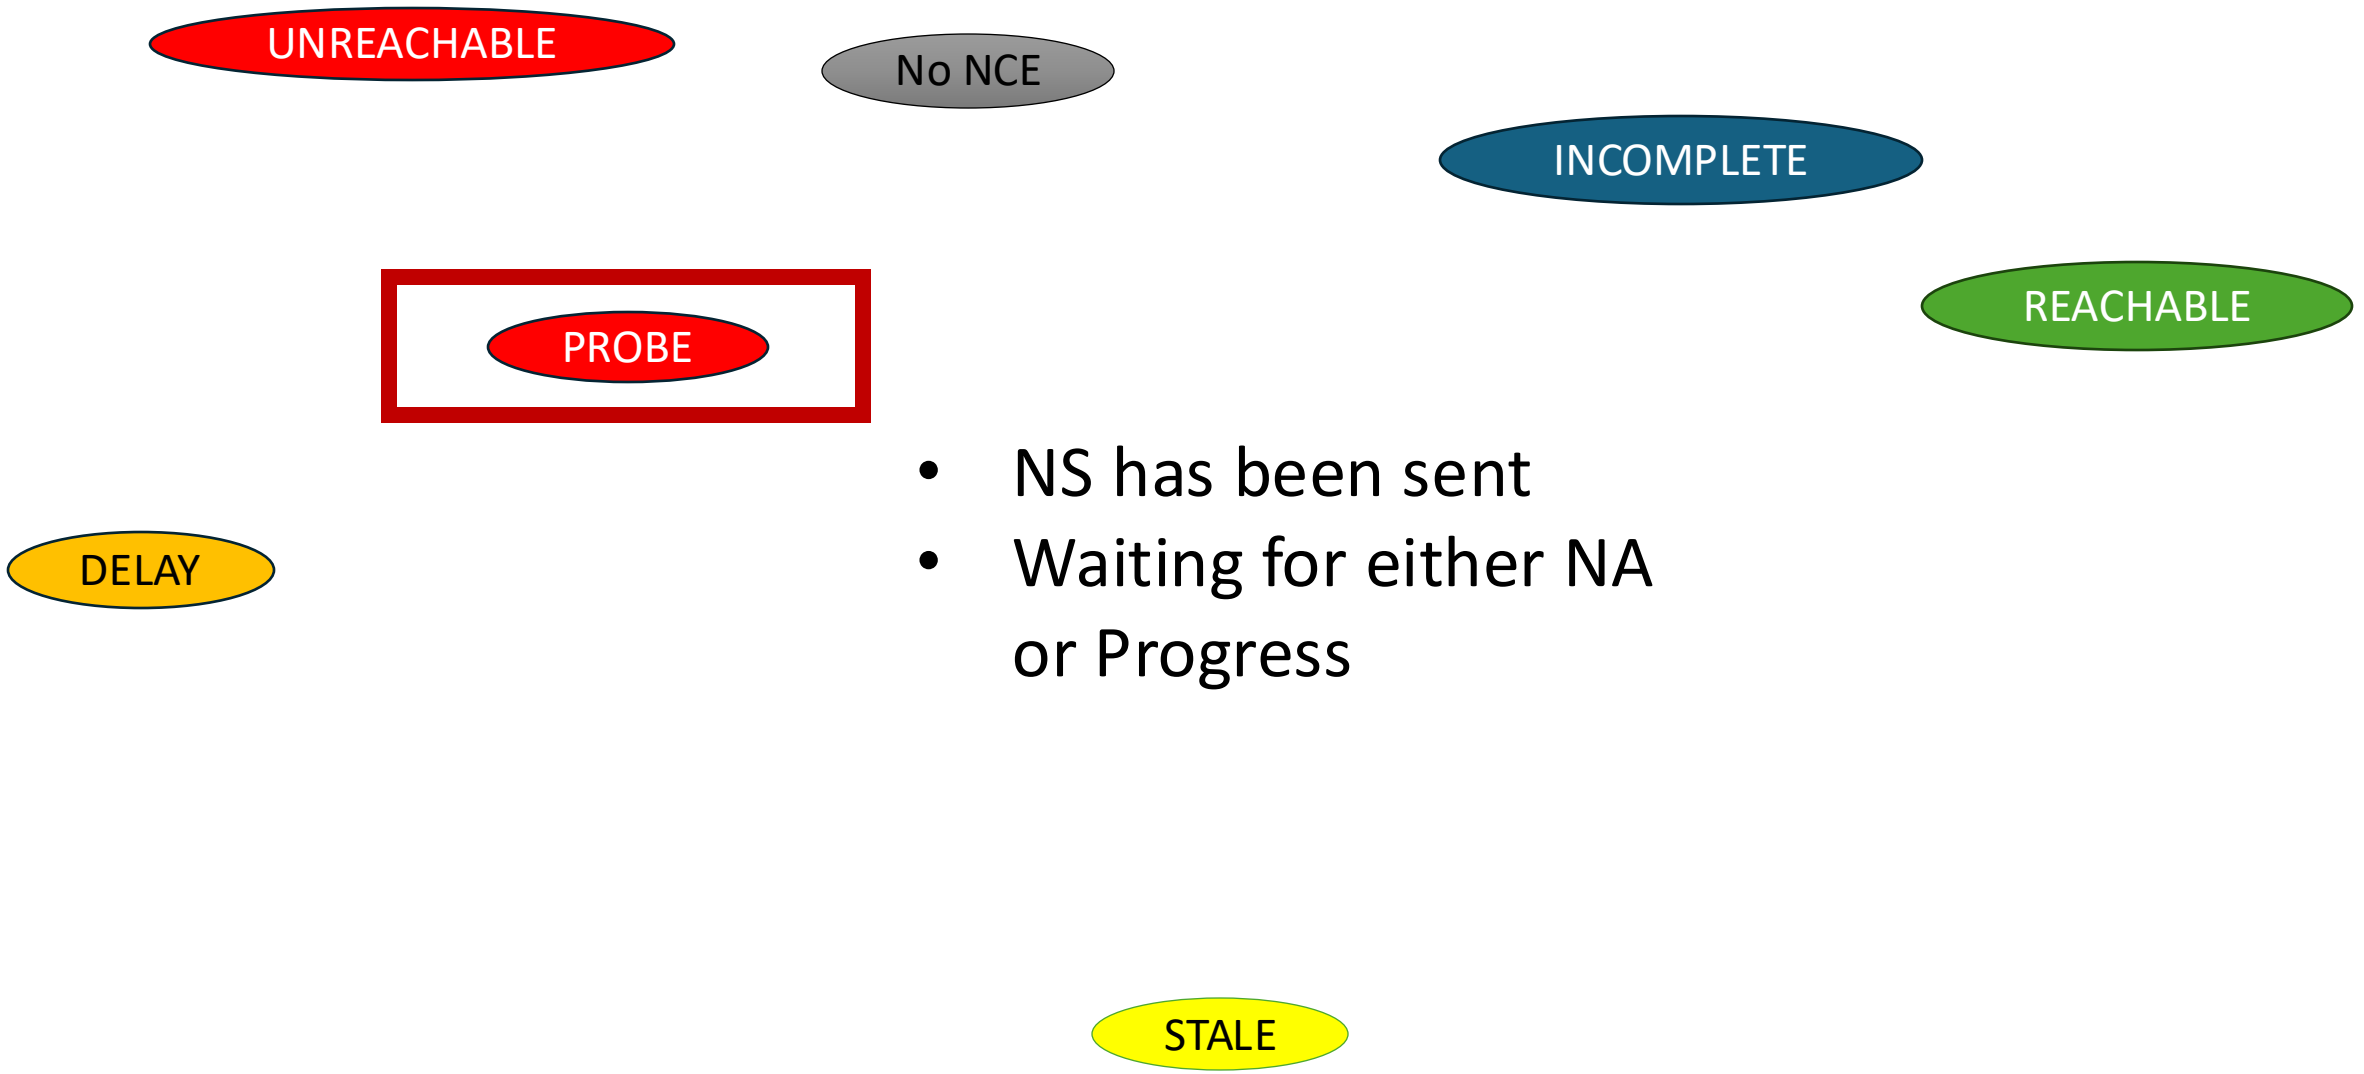

- NS has been sent
- Waiting for either NA or Progress

Neighbor Cache Entry States (U)

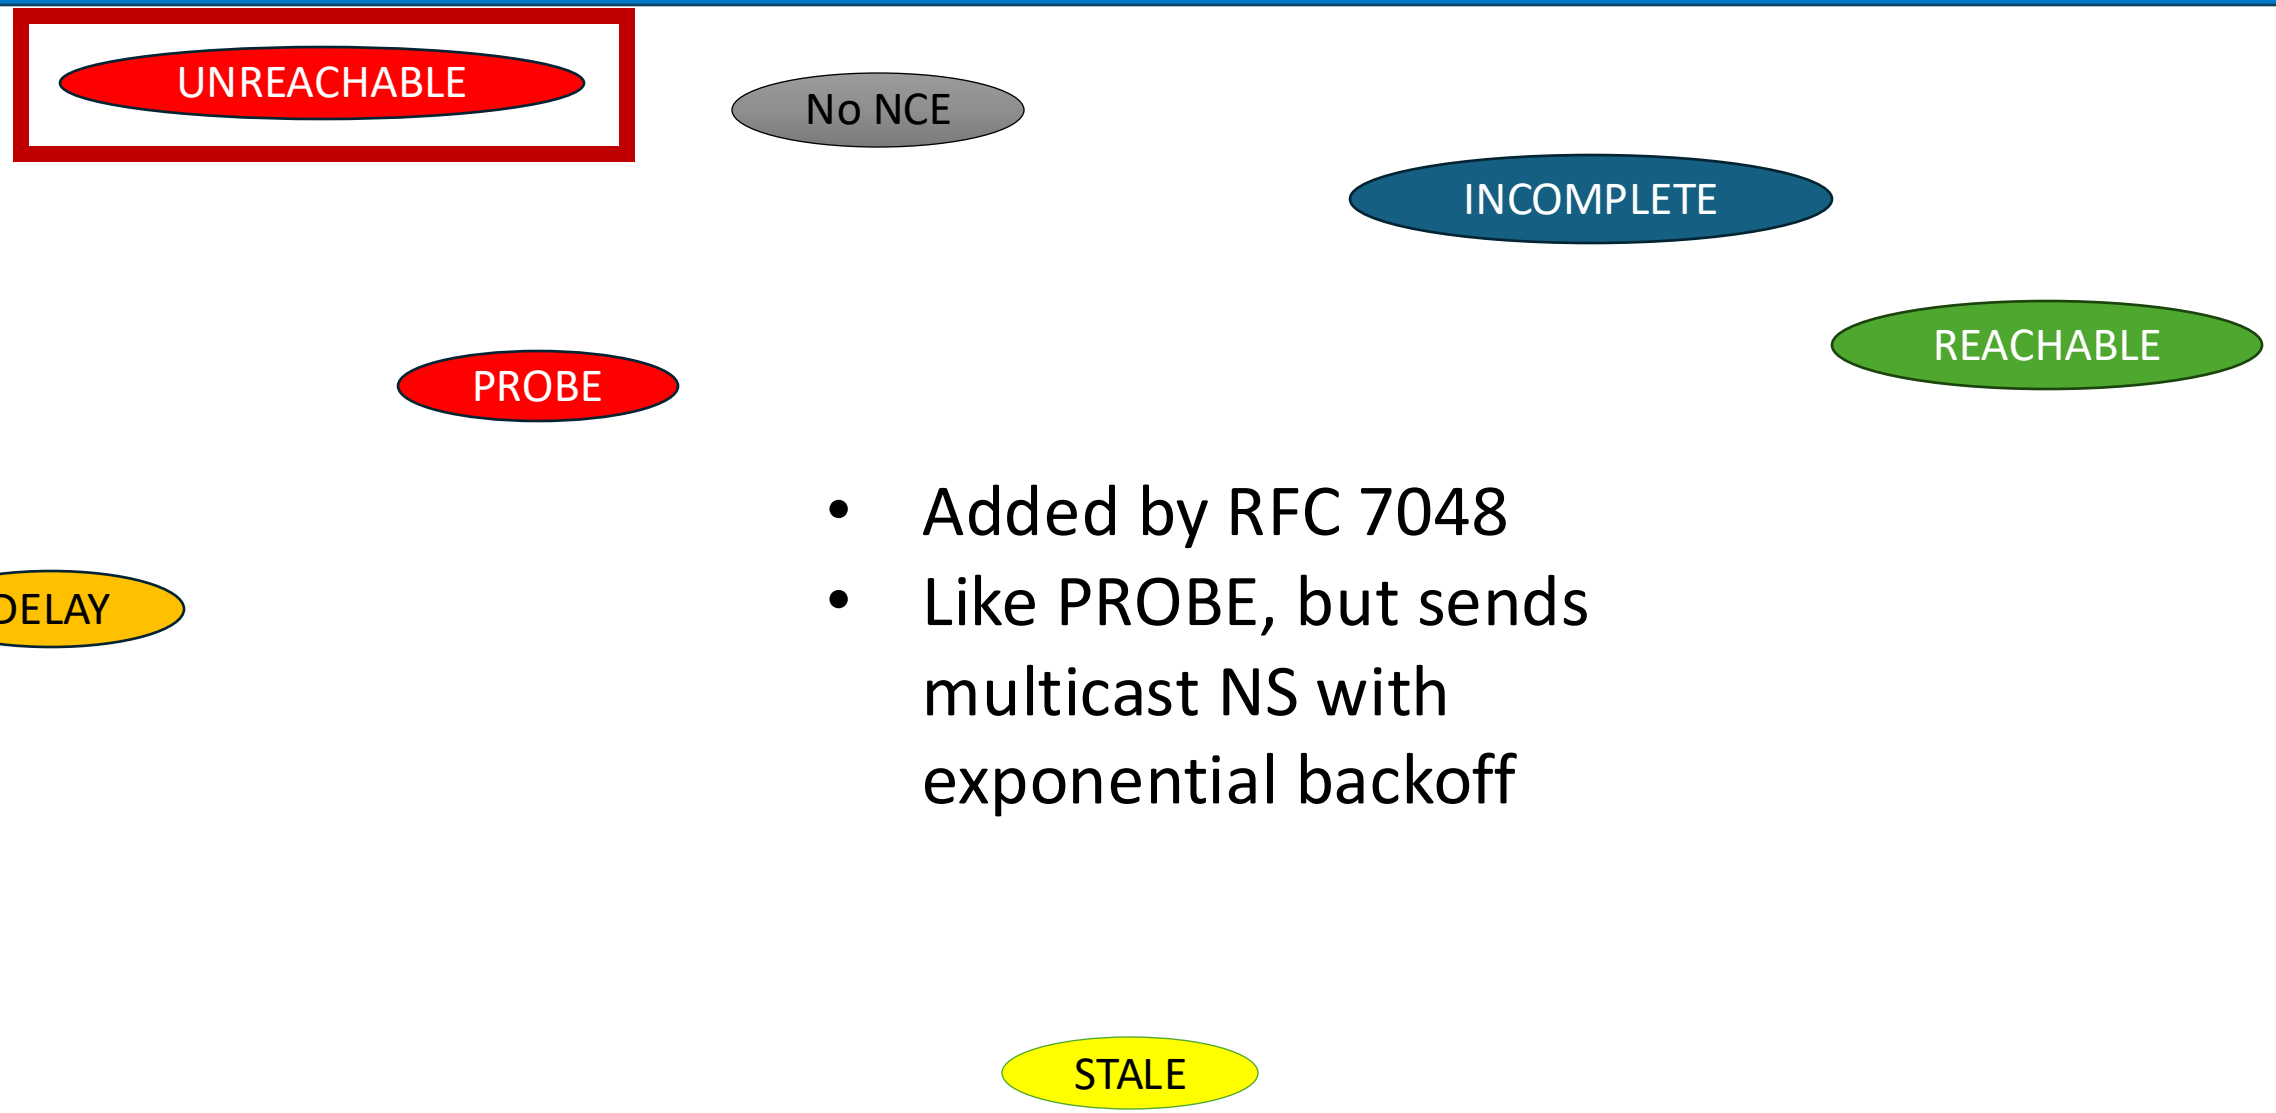

- Added by RFC 7048
- Like PROBE, but sends multicast NS with exponential backoff

NCE Typical Lifecycle

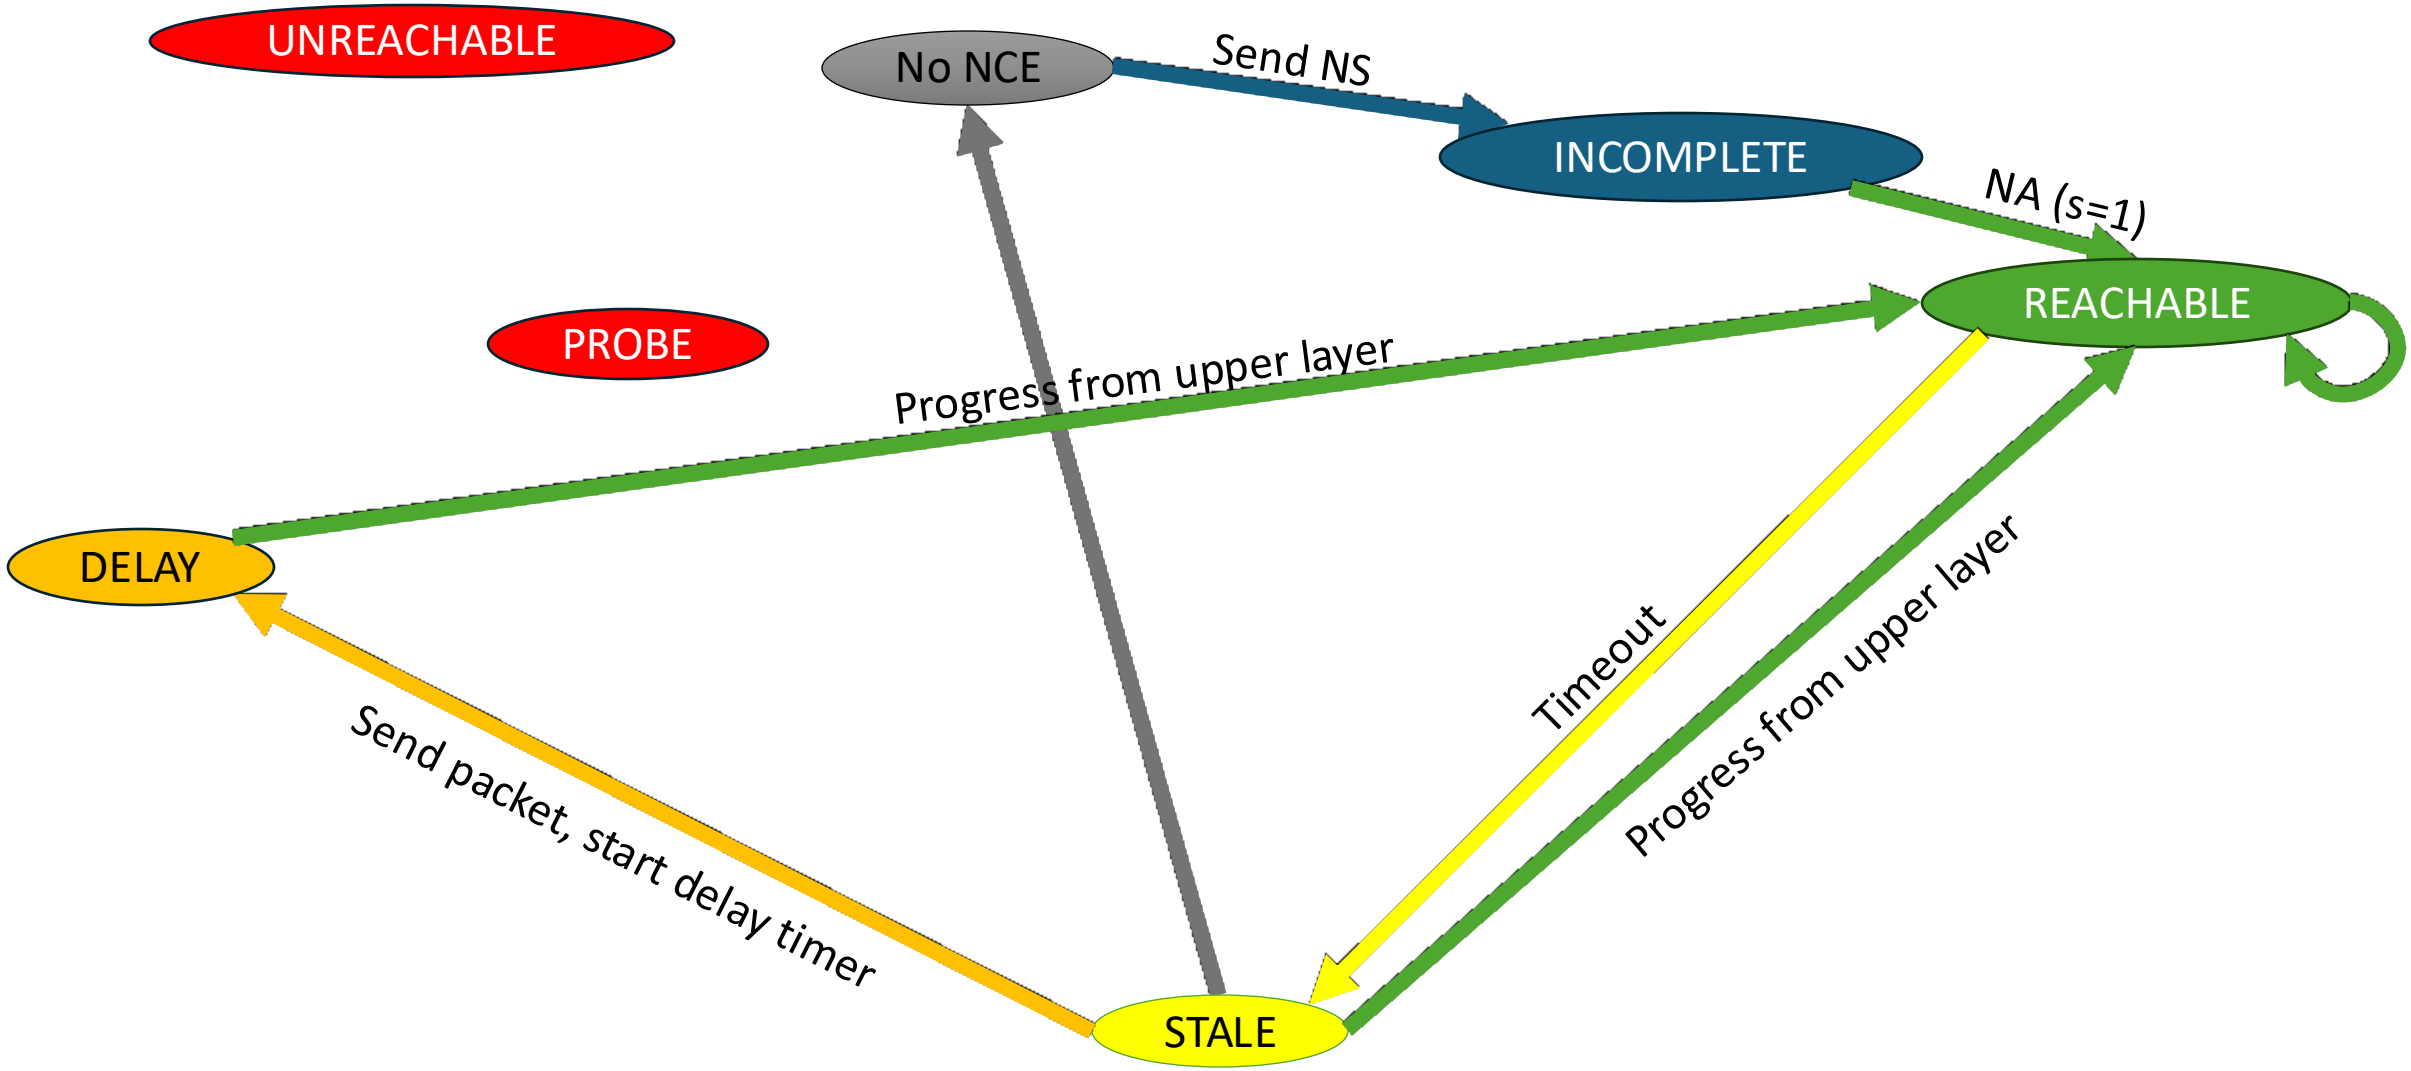

NCE Node Connectivity Lost

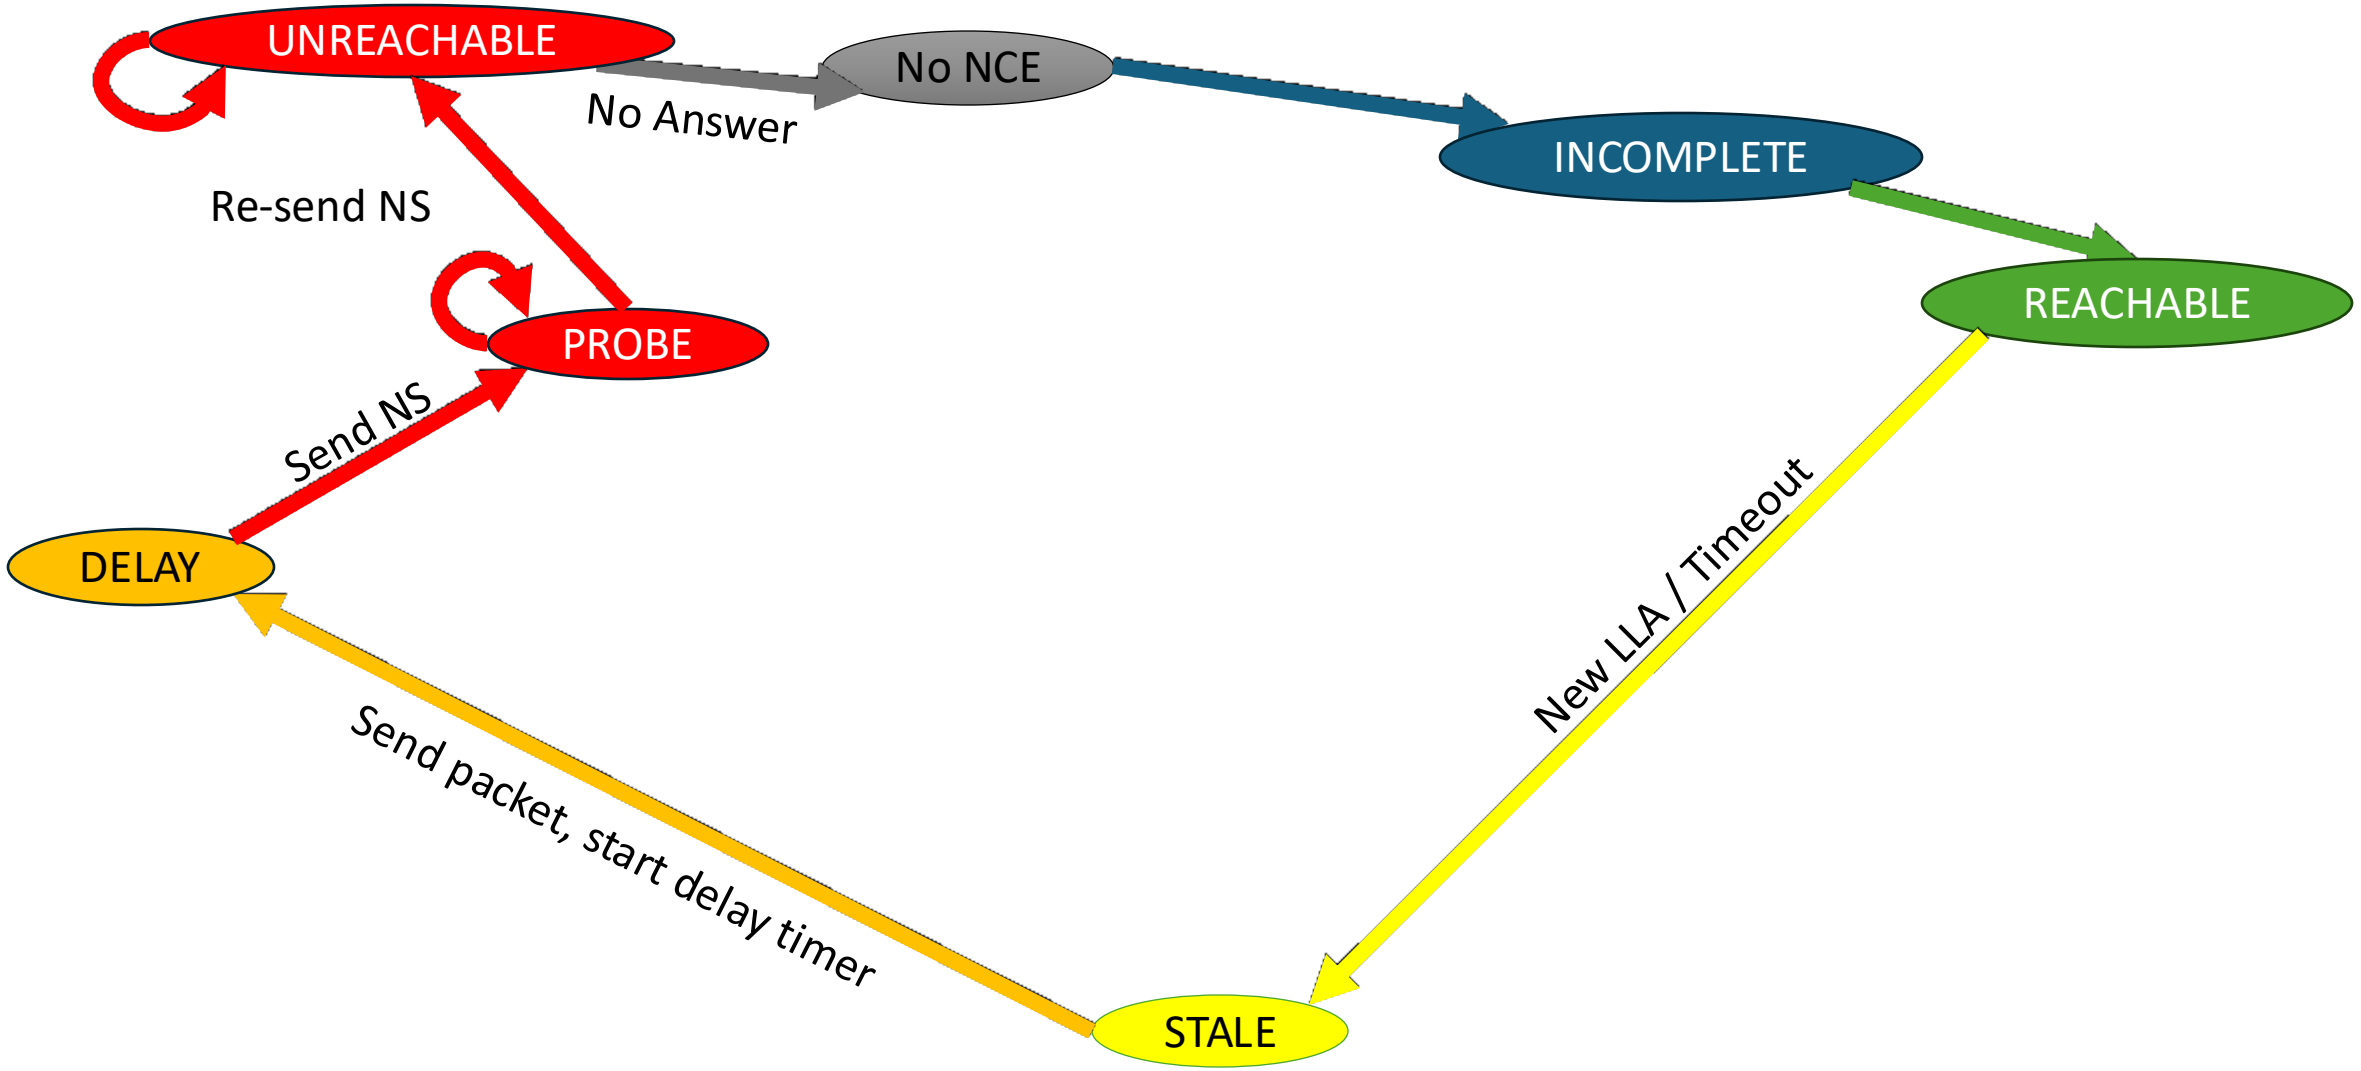

NCE Node LLA Changed

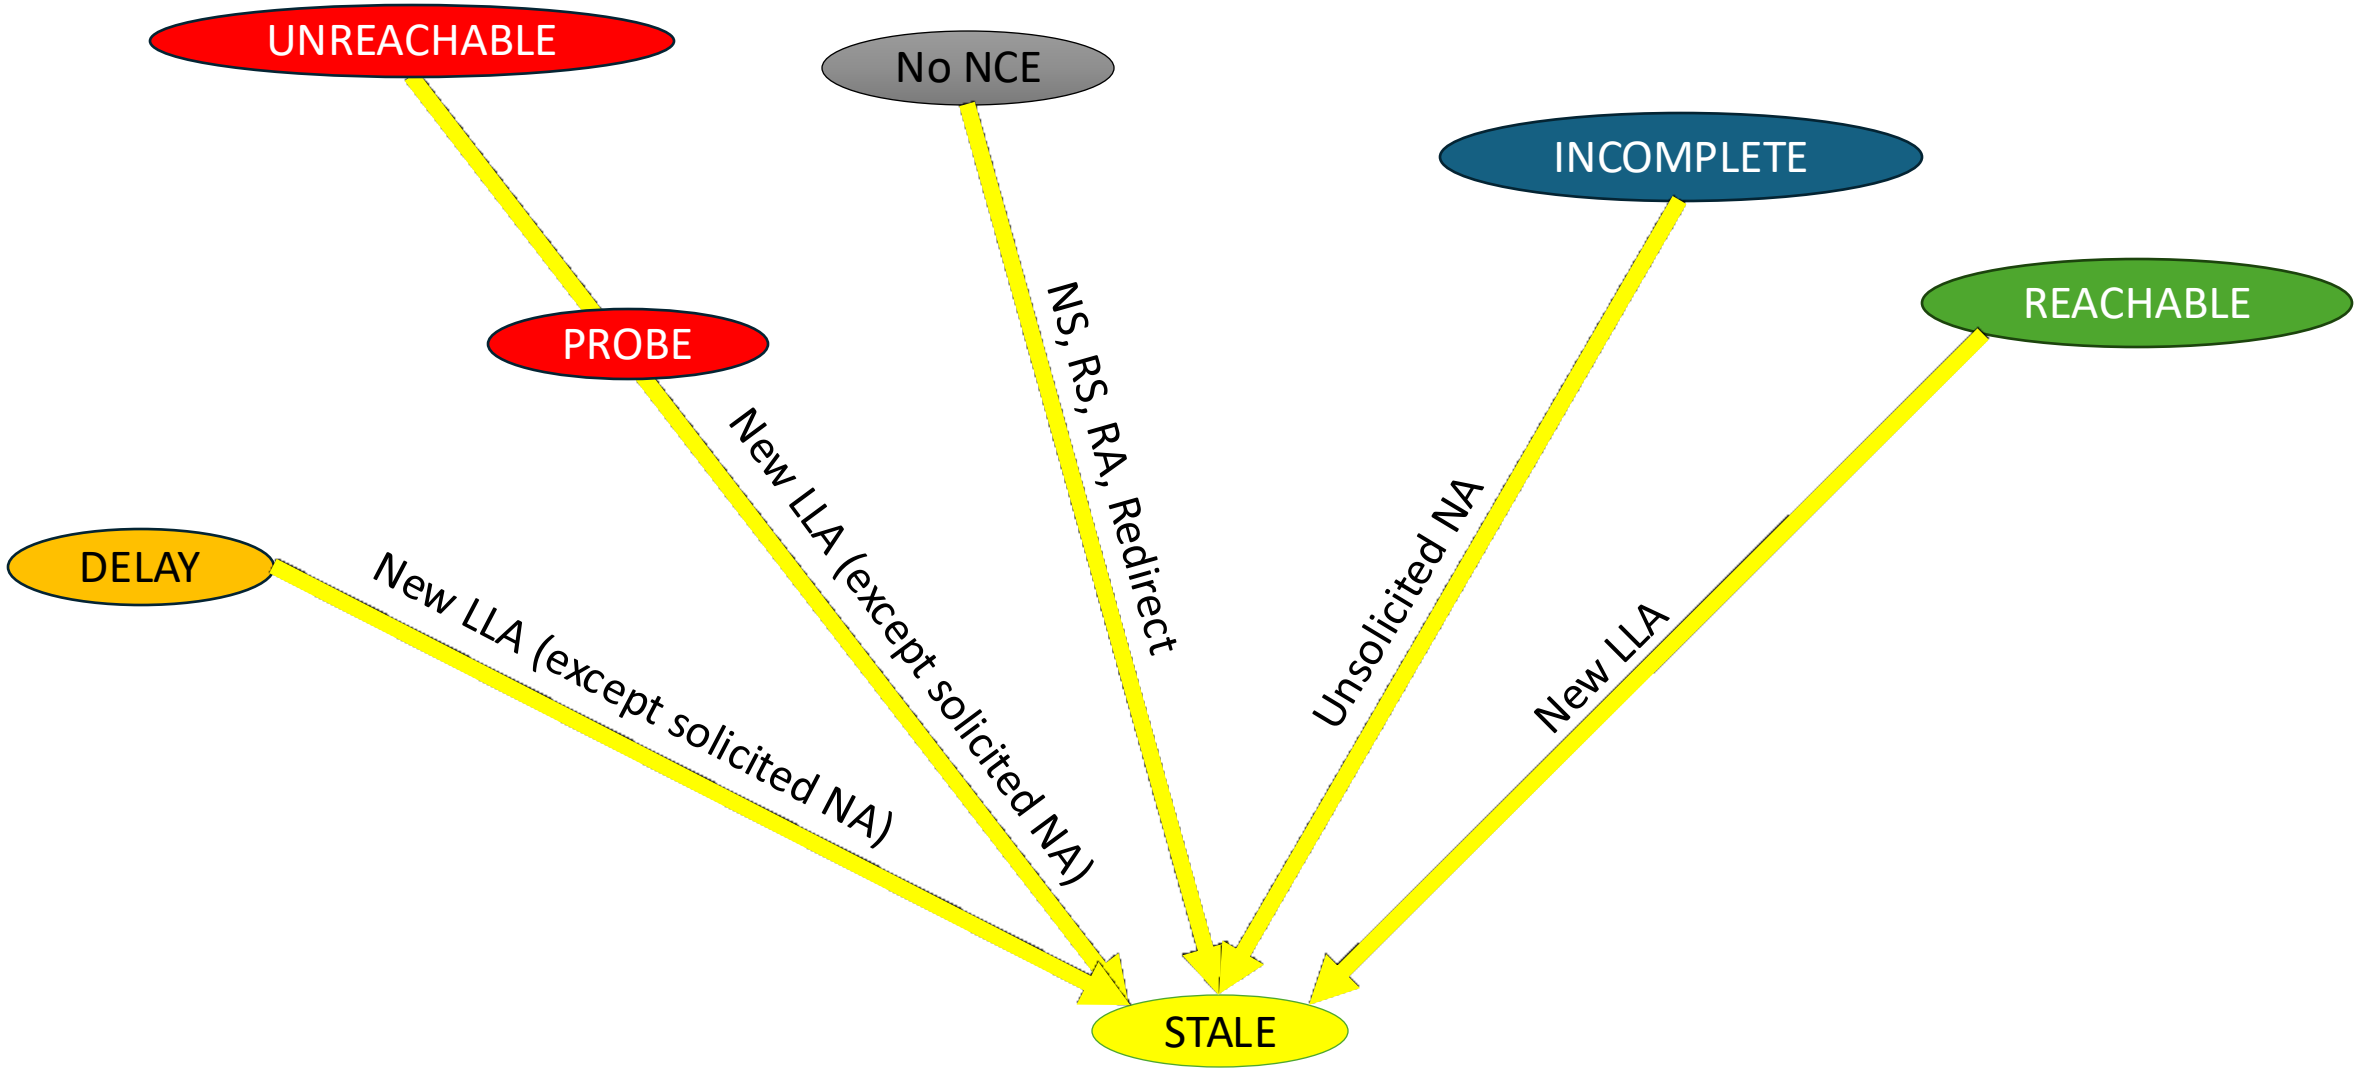

Trust, but verify…

NCE Total State Diagram

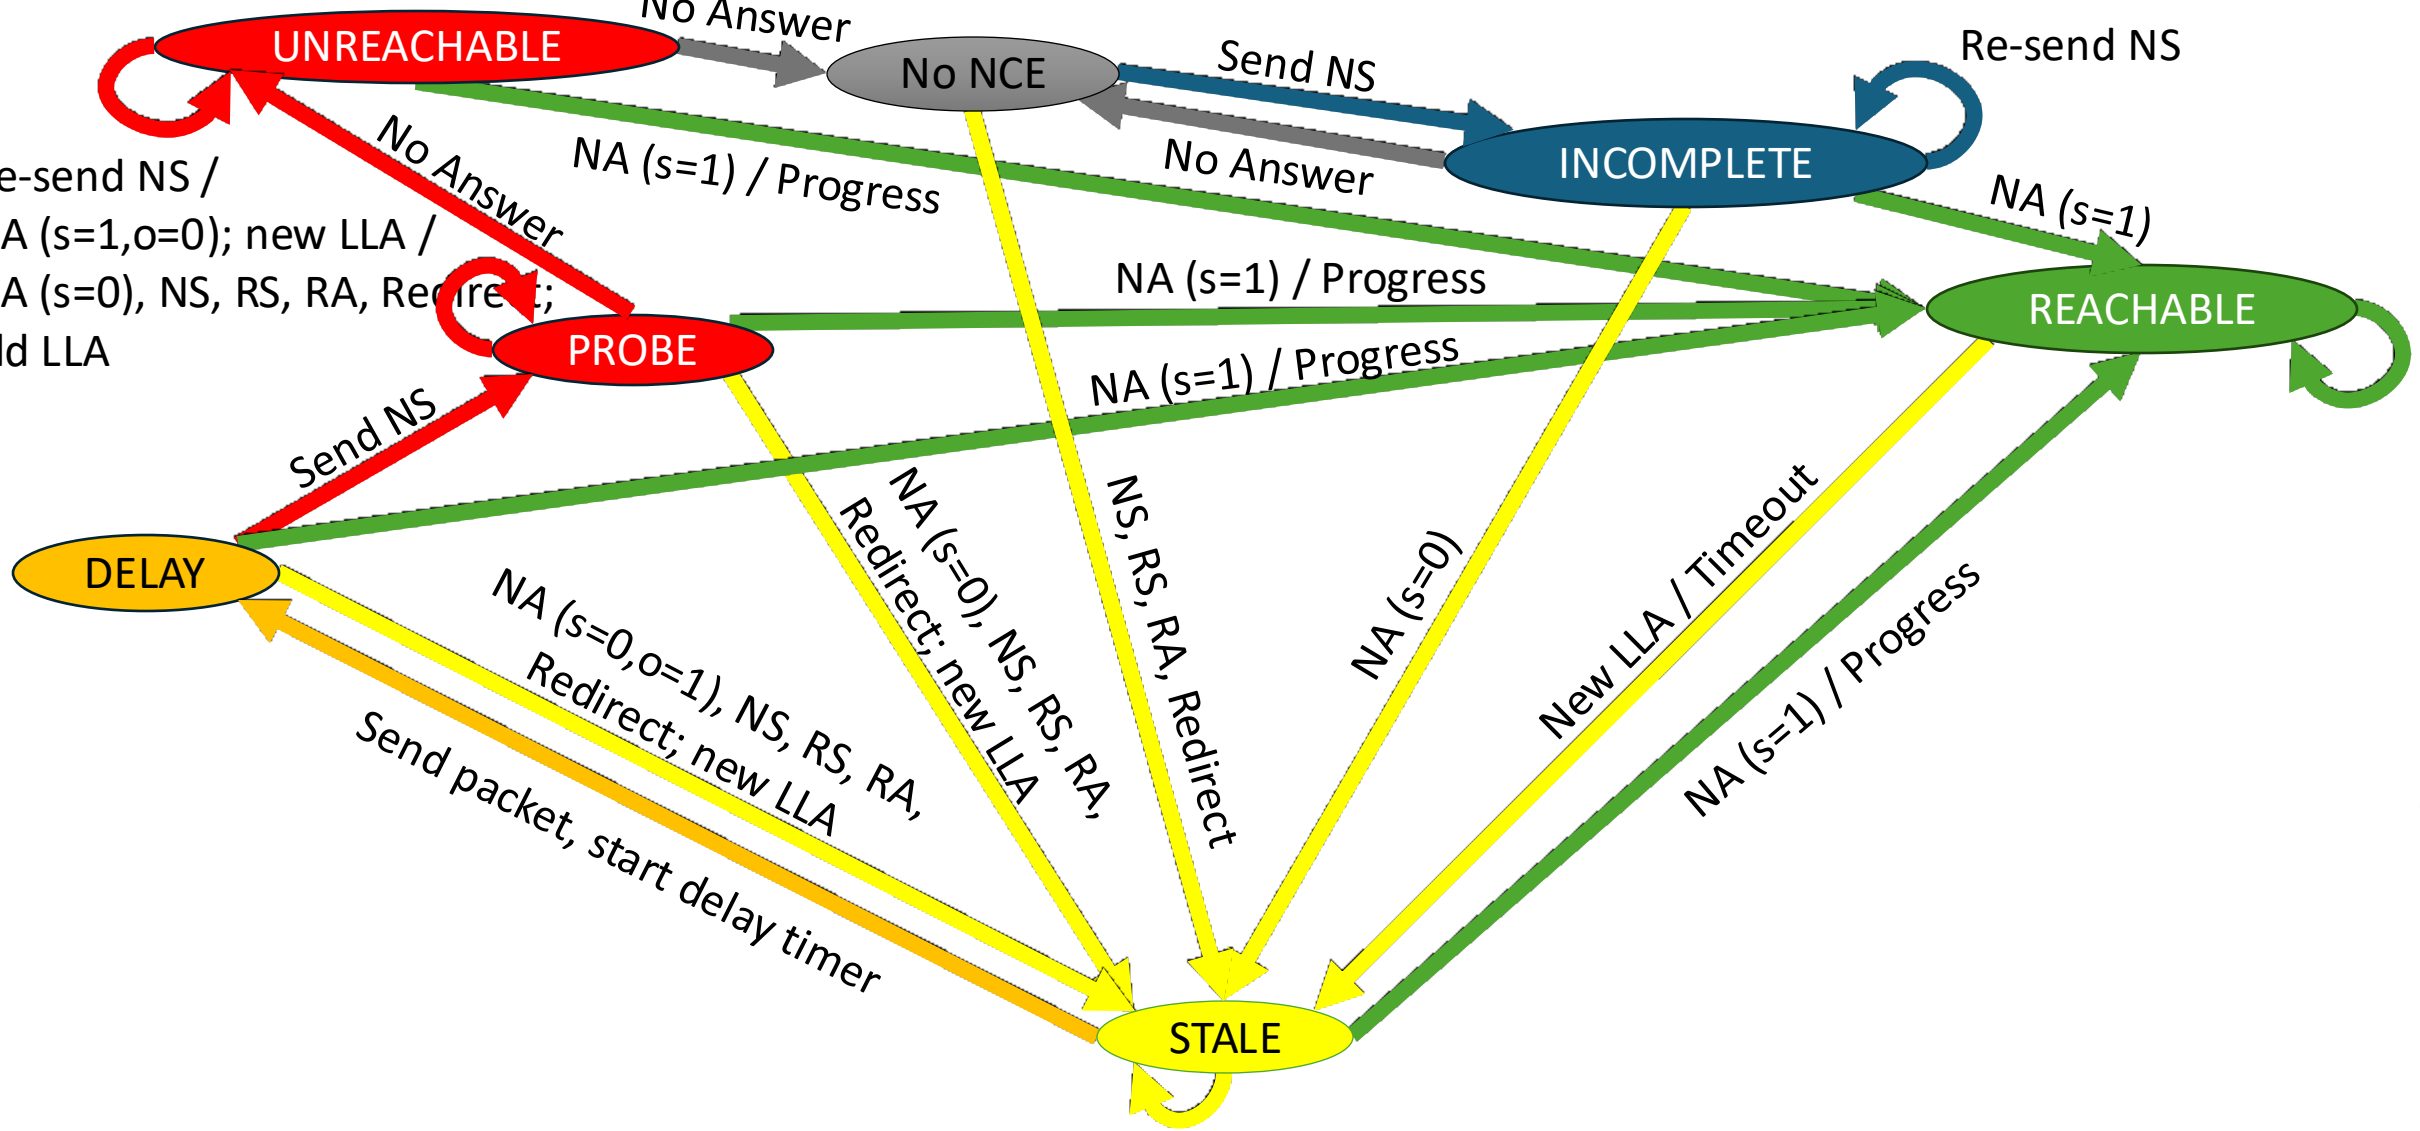

Copies of slides

THANK YOU

We greatly appreciate your time today and continued support!

Please take a moment to rate this session in the *CEN Events* mobile app.