Allocation Schemes

CIDR representation and IPv6 allocations

IPv4 Subnet Masking

- Originally the network size was based on the first few bits (classful addressing)
- Getting rid of address classes was *painful!*
  - routing protocols, stacks, applications
- Modern IPv4 allows subnet boundaries anywhere within the address (classless addressing)
- But decimal addresses still make figuring out subnets unnecessarily difficult. . .

CIDR

In IPv4 you would see representations like:

129.93.0.0/16

129.93.0.0 255.255.0.0

At the bit level this is:

10000001.01011101.00000000.00000000

Reasons for CIDR

- To try to preserve the address space.
- To control the growth of the routing table.

IPv6 Notation

- In IPv6 every address is written:
  - IPv6 address / prefix length
- For example:
  - 2001:0468::/35
  - 2001:0468::/32
- At the bit level:
  - 0010 0000 0000 0001: 0000 0100 0110 1000: 000 0::/35
  - 0010 0000 0000 0001: 0000 0100 0110 1000 0::/32

Allocation Strategies Example

- We wish to allocate /48s out of the /35.
- Which are available:
  - 2001:0468:0000 through
  - 2001:0468:1fff
- Recall that the bit structure is:
  - 0010 0000 0000 0001: 0000 0100 0110 1000: 000 | 0:0000:0000:0000
  - 0010 0000 0000 0001: 0000 0100 0110 1000: 000 | 1:1111:1111:1111
- So there are 8192 /48s in a /35

Why Allocation?

- To try to control the growth of the routing table in the default-free zone.
- It is a necessary consequence of using a provider-based aggregatable address scheme.
- It makes the address space more manageable.

How would allocations work?

- Suppose you wish to give out /40s in the /35.
  - 2001:0468:000 | 0 0000 | or 2001:0468::/40
  - 2001:0468:000 | 1 1111 | or 2001:0468:1f00::/40
- Thus there are 32 /40s in the /35, each of which has 256 /48s.
  - 5 bits
  - 8 bits

How would allocations work?

- The same idea holds for /41s or /42s.
  - 2001:0468:000 | 0:0000:0 | or 2001:0468::/41
  - 2001:0468:000 | 1:1111:1 | or 2001:0468:1f80::/41
  - 2001:0468:000 | 0:0000:00 - :000 | 1:1111:11
  - 2001:0468::/42 – 2001:0468:1fc0::/42

Mixed Allocations

- The interesting case is how to handle mixed allocations.
- Some sites need a /40, others a /42. How can you handle this case?
- See
  - RFC 3531 (Marc Blanchet)
  - A flexible method for managing the assignment of bits of an IPv6 address block
  - A perl script is included.

Mixed Allocations

- Take 2001:468::/32. Out of that allocate:
  - 2 subnets of /34
  - 3 subnets of /37
  - 5 subnets of /38
- Review address allocations (separate slide)
- Assign addresses:
  - Assign /34s for the two top-tier routers.
  - Assign /35s for their downstream routers.
  - Assign /37s for the third-tier routers.
  - Remember at each level to retain some /64s for "local" use, and allocate them for point-to-point links in the network diagram.
  - When you're done, your network diagram will have loopbacks, point-to-points, and appropriately-sized network blocks allocated at each level.

Neighbor Solicitation

Neighbor Solicitation

- This protocol solves a set of problems related to the interaction between nodes attached to the same link. It defines mechanisms for solving each of the following problems...

Problems Solved by Neighbor Solicitation

- Router Discovery: How hosts locate routers that reside on an attached link.
- Prefix Discovery: How hosts discover the set of address prefixes that define which destinations are on-link for an attached link. (Nodes use prefixes to distinguish destinations that reside on-link from those only reachable through a router.)
- Parameter Discovery: How a node learns such link parameters as the link MTU or such Internet parameters as the hop limit value to place in outgoing packets.

Problems Solved by Neighbor Solicitation

- Address Autoconfiguration: How nodes automatically configure an address for an interface.
- Address resolution: How nodes determine the link-layer address of an on-link destination (e.g., a neighbor) given only the destination's IP address.
- Next-hop determination: The algorithm for mapping an IP destination address into the IP address of the neighbor to which traffic for the destination should be sent. The next hop can be a router or the destination itself.

Problems Solved by Neighbor Solicitation

- Neighbor Unreachability Detection: How nodes determine that a neighbor is no longer reachable. For neighbors used as routers, alternate default routers can be tried. For both routers and hosts, address resolution can be performed again.
- Duplicate Address Detection: How a node determines that an address it wishes to use is not already in use by another node.
- Redirect: How a router informs a host of a better first-hop node to reach a particular destination.

ICMP Packet Types

- Neighbor Discovery defines five different ICMP packet types: a pair of Router Solicitation and Router Advertisement messages, a pair of Neighbor Solicitation and Neighbor Advertisement messages, and a Redirect message. The messages serve the following purposes…

ICMP Packet Types

- Router Solicitation: When an interface becomes enabled, hosts may send out Router Solicitations that request routers to generate Router Advertisements immediately rather than at their next scheduled time.
- Router Advertisement: Routers advertise their presence together with various link and Internet parameters either periodically, or in response to a Router Solicitation message. Router Advertisements contain prefixes that are used for on-link determination and/or address configuration, a suggested hop limit value, etc.

Valid & Preferred Prefixes

- Valid & Preferred lifetime values in Router-Advertisements can be used to re-number.
- During a prefix's preferred life, new connections can be opened at will.
- During a prefix's valid life, existing connections can be used, but new connection may not be opened.
- These values are continually refreshed by default.

ICMP Packet Types

- Neighbor Solicitation: Sent by a node to determine the link-layer address of a neighbor, or to verify that a neighbor is still reachable via a cached link-layer address. Neighbor Solicitations are also used for Duplicate Address Detection.
- Neighbor Advertisement: A response to a Neighbor Solicitation message. A node may also send unsolicited Neighbor Advertisements to announce a link-layer address change.
- Redirect: Used by routers to inform hosts of a better first hop for a destination.

Stateless Autoconfiguration

Why does this matter?

- Manual configuration of individual machines before connecting them to the network should not be required.
  - Address autoconfiguration assumes that each interface can provide a unique identifier for that interface (i.e., an "interface token")
- Plug-and-play communication is achieved through the use of link-local addresses
  - Small sites should not need stateful servers
- A large site with multiple networks and routers should not require the presence of a stateful address configuration server.
- Address configuration should facilitate the graceful renumbering of a site's machines

Stateless Autoconfiguration

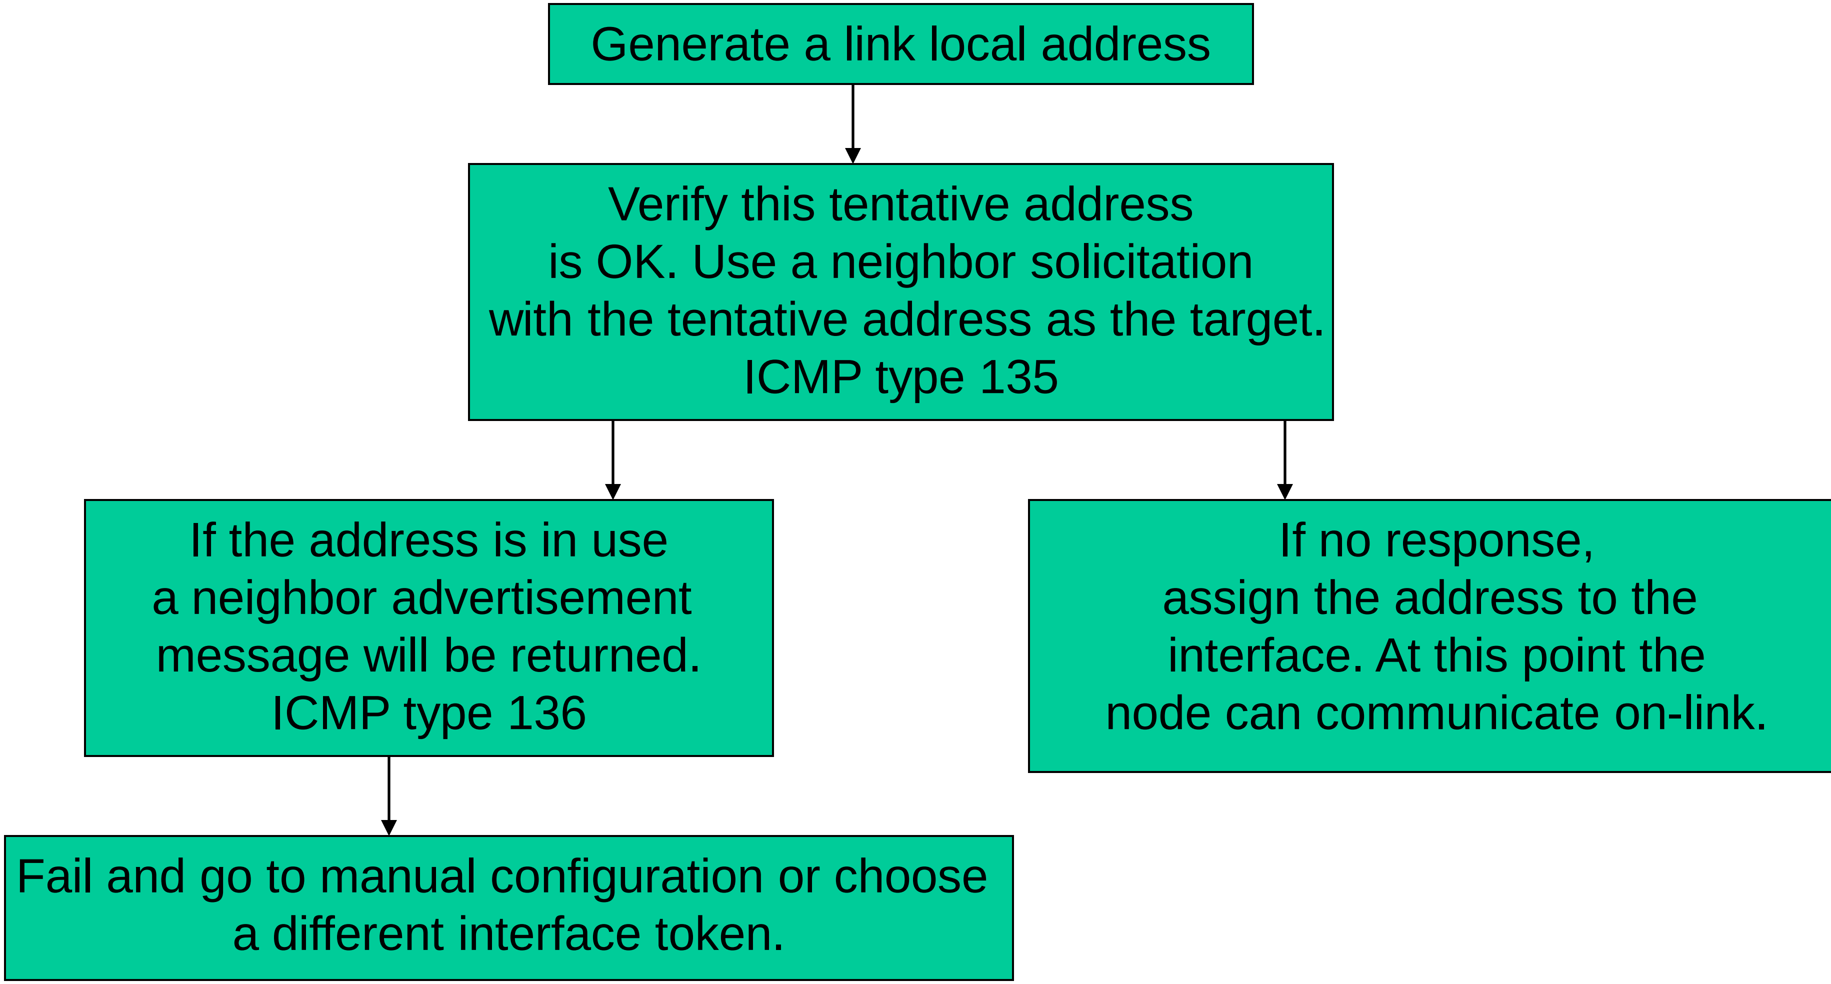

Stateless Autoconfiguration

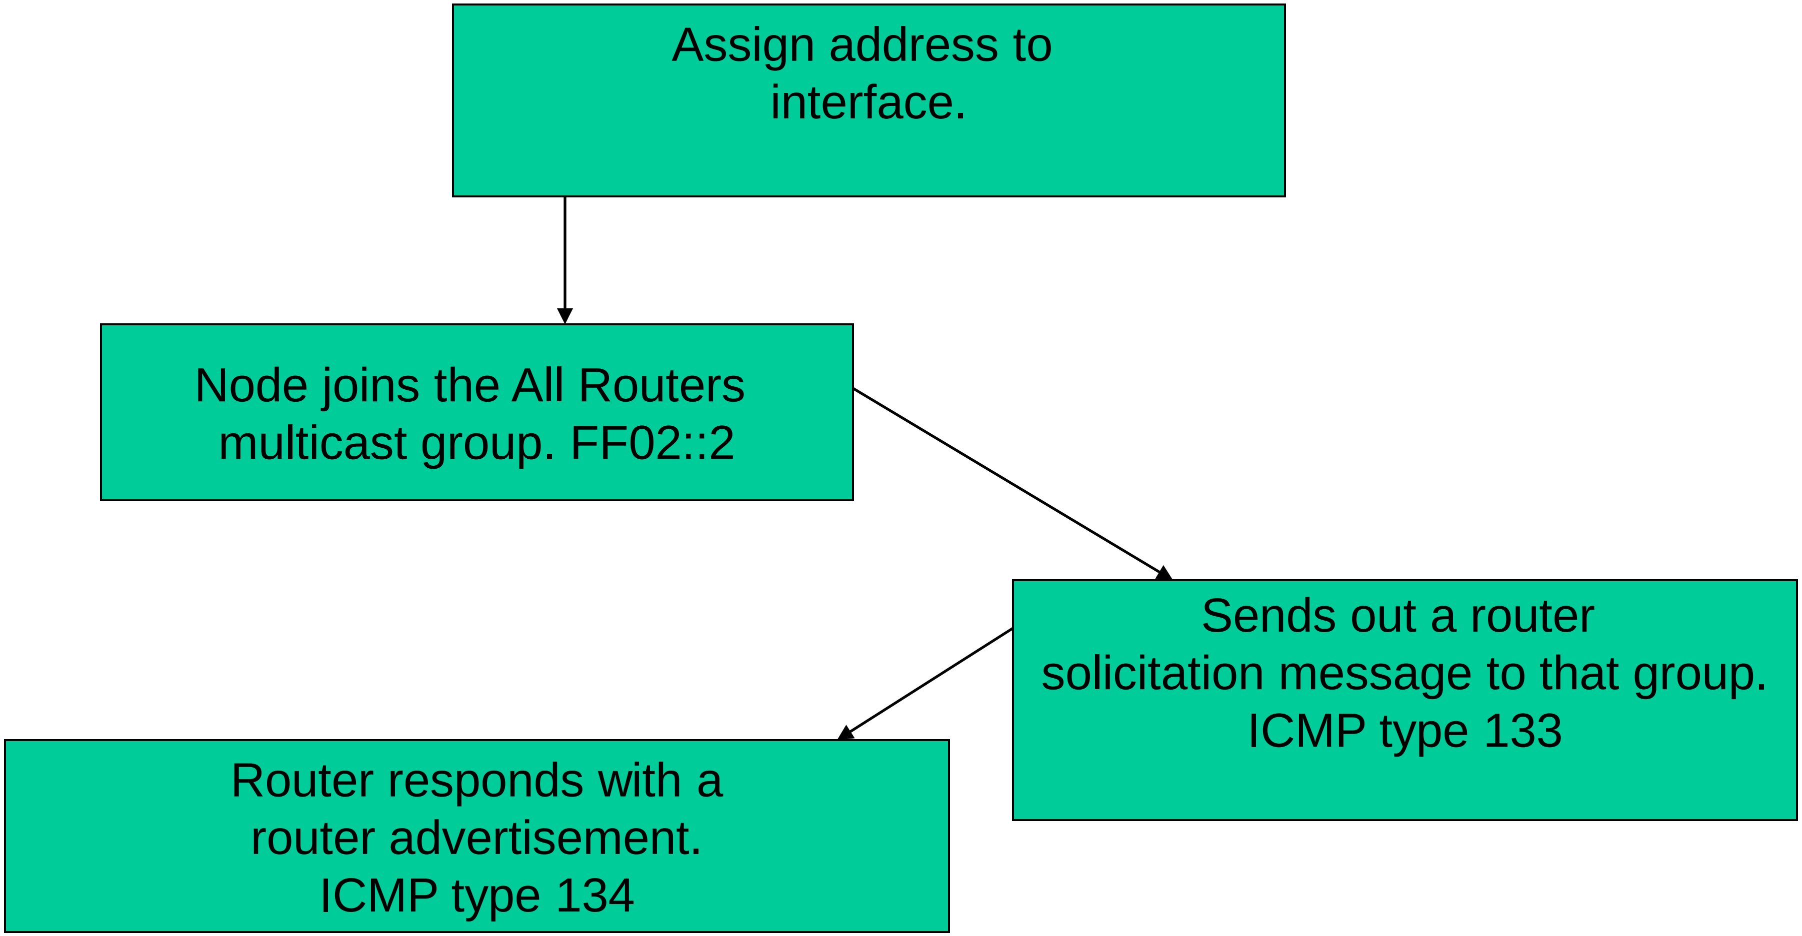

Stateless Autoconfiguration

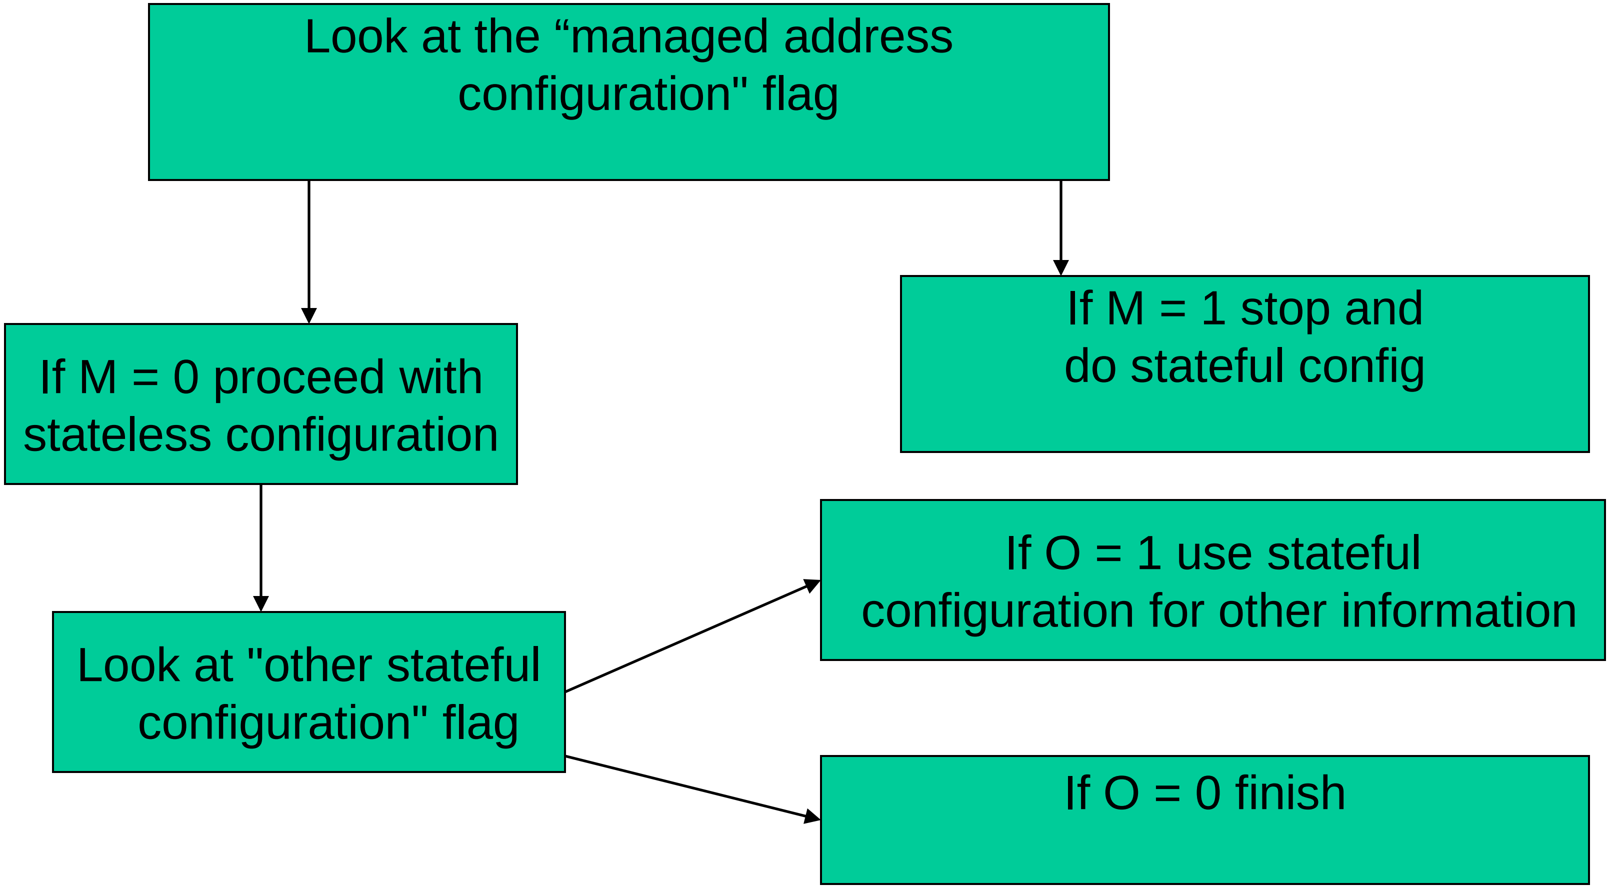

Router Solicitation

Type = 133	Code = 0	Checksum	
Reserved			
Possible option: Source Link Layer Address			

Router Advertisement

Type = 134	Code = 0			Checksum
Cur. Hop Limit	M	O	**Reserved**	Router Lifetime
Reachable Time				
Retransmission Timer				
Possible options: -Source Link Layer Address -MTU -Prefix Information				

Neighbor Solicitation

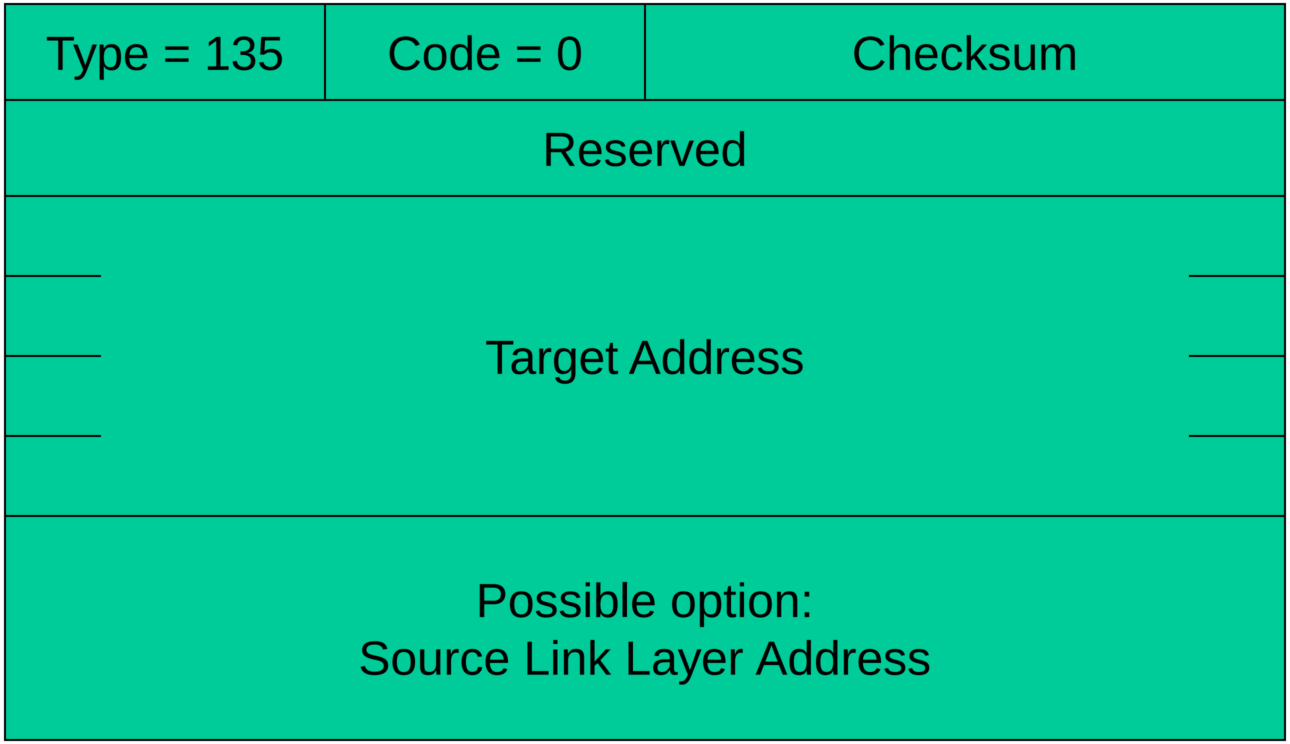

Neighbor Advertisment

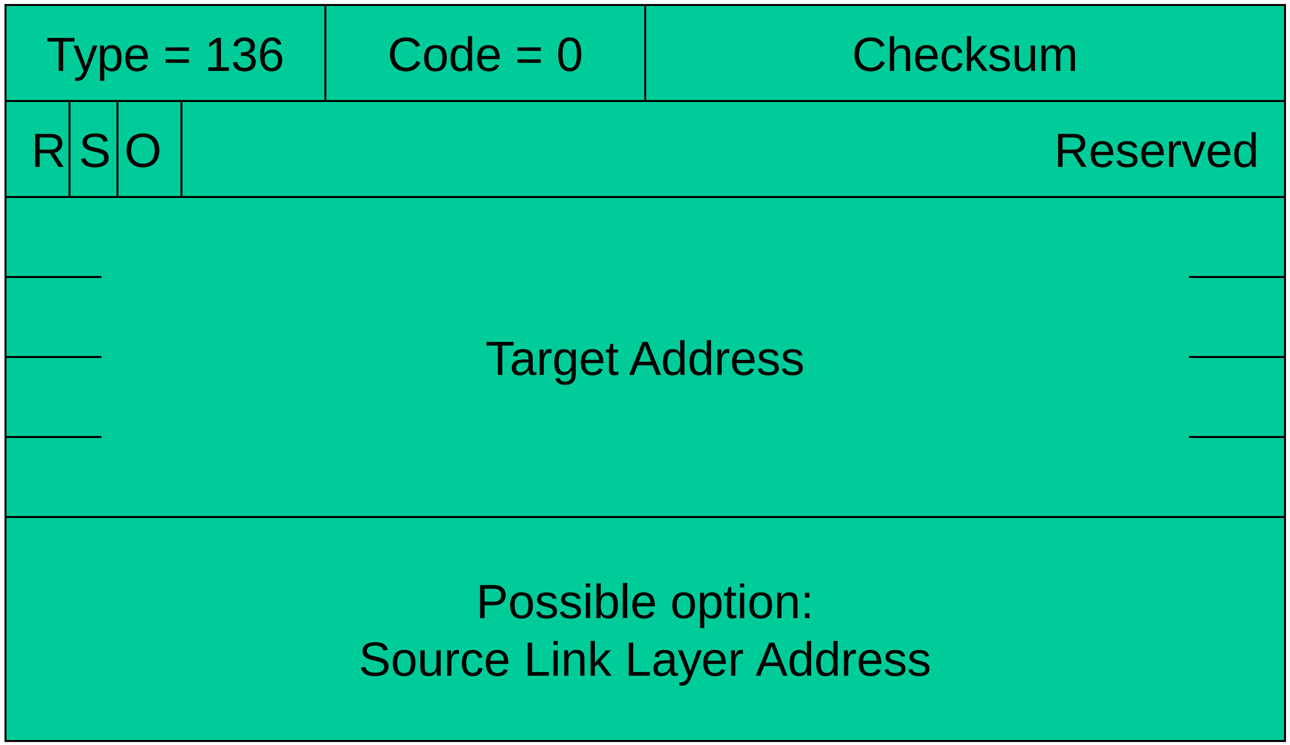

Prefix Option

Type	Length	Prefix Length	L	A	Reserved
Valid Lifetime					
Preferred Lifetime					
Reserved					
Prefix List					

Router Solicitation Options
Prefix Information

- This should include all prefixes the router is aware of
- Flag bits:
  - On-link = 1
    - Prefix is specific to the local site
  - Autonomous Configuration bit = 1
    - Use the prefix to create an autonomous address

Router Solicitation Options
Prefix Information

- Valid Lifetime
  - 32-bit unsigned integer. The length of time in seconds before an address is invalidated.
- Preferred Lifetime
  - 32-bit unsigned integer. The length of time in seconds before an address is deprecated.

Stateless Autoconfig

- Routers are to send out router advertisements at regular intervals to the all-hosts address.
  - This should update lifetimes.
- Note that stateless autoconfig will only configure addresses.
  - It will not do all the host configuration you may well want to do.
- RFC 2462 defines IPv6 Stateless Autoconfig

Stateful Configuration

- When you do not wish to have stateless configuration done you will need to provide a configuration server (DHCP most likely) to provide configuration information to the hosts as they come up.
  - RFC 3315 defines DHCP, updated by RFC 4361
  - Dibbler – DHCPv6 implementation
    - http://sourceforge.net/projects/dibbler

DHCP 'Lite'

- Used in combination with stateless address configuration, to provide other information:
  - DNS resolver
  - domain suffix

```
 ipv6 dhcp pool v6lite
 dns-server 2001:4::1
 domain-name example.com
!
interface FastEthernet0/1
 ipv6 address 2001:4:1::1/64
 ipv6 nd other-config-flag
 ipv6 dhcp server v6lite
```

Address Configuration Lab

- Start Ethereal running on host R
- Disconnect and reconnect the Ethernet cable
- Observe the neighbor discovery and attempted address configuration packets

- Log in to router E
- Enable IPv6 on the interface:

```
interface f0/1
ipv6 address 2001:468:110z:xxxx::/64
```

- Disconnect and reconnect the Ethernet, and observe the address autoconfiguration process
- Verify the address with ifconfig