

Introduction to IPv6

(Chapter 4 in Huitema)

IPv6 addresses

- 128 bits long
- Written as eight 16-bit integers separated with colons
 - E.g. 1080:0000:0000:0000:0000:0008:200C:417A
= 1080::8:800:200C:417A
- Types
 - Unicast
 - Defines one interface within their scope of validity
 - Multicast
 - Delivers packets to all members of a group
 - Anycast
 - Delivers packets to the *nearest* member of a group


Special IPv6 addresses

- Unspecified = $0:0:0:0:0:0:0:0 = ::$
 - Only as source address
- Loopback = $0:0:0:0:0:0:0:1 = ::1$
 - For sending datagrams to itself
- IPv4 addresses prepended with zeroes
 - $0:0:0:0:0:0:AABB:CCDD = ::a.b.c.d$
- Site-local addresses
 - $FEC0:0000:0000:subnet:station$ (subnet 16 bits, station 64 bits)
- Link-local addresses
 - $FEB0:0000:0000:0000:station$

IPv6 header

Version=6 (4)	Traffic class (8)	Flow label (24 bits)	
Payload length (16 bits)		Next header type (8)	Hop limit (8)
Source address (128 bits)			
Destination address (128 bits)			

- Differences between v4 and v6
 - No checksum (performed by lower layers)
 - No fragmentation (path MTU discovery instead, min. 1280)
 - No options (fixed length header, options in linked extension headers instead)
- Extension headers replace options



Source routing is implemented with the routing header

- Routing header:

Next header	Header ext. length	Routing type = 0	Segments left
Reserved			
IPv6 address 1			
IPv6 address 2			
...			
IPv6 address N			

- Only the router whose address is destination address in IPv6 header examines this extension \Rightarrow better performance
- Forwarder
 - Moves the next address to the IPv6 header
 - Decrements the number of segments left

Can be replaced by IP-in-IP

Only the sender can fragment packets

- Packets larger than the next hop's MTU are rejected
- Large packets must be fragmented by the sender
- Fragment header:

Next header	Reserved	Fragment offset	Reserved	M
Identification				


- Offset: Most significant 13 bits of 16-bit word
- M: More fragments
 - M=1 in all packets but the last

Other extensions

- Authentication Header (AH)
- Encrypted Security Payload (ESP)
- Destination options header
 - Only examined by the destination
 - Contains one or several options
 - Also defines handling for unrecognized parameters
- Hop-by-hop options header
 - Examined by each router
 - Similar format and coding as destination options header
 - E.g. jumbo payload
- Processing order is important
 - IPv6 → Hop-by-hop → Destination options (for tunneling) → Routing → Fragment → Authentication → Destination options → Upper layers (TCP/UDP)

Internet Control Message Protocol Version 6

- ICMPv6 header



- Also includes the functionalities of IGMP and ARP

- ICMP message types:

- 1. Destination unreachable
 - 2. Packet too big
 - 3. Time exceeded
 - 4. Parameter problem
 - 128. Echo request
 - 129. Echo reply
 - 133. Router solicitation
 - 134. Router advertisement
 - 137. Redirect
- } errors
- } for "ping"
- } router discovery

Router discovery

- For building a local list of routers on the same network

Type = 134	Code = 0			Checksum
Cur. hop limit	M	O	Res.	Router lifetime
Reachable time				
Retransmission timer				
Options (e.g. source link layer option, MTU option)				

- Curr.hop limit: Suggestion for initial hop limit value
 - Router lifetime: Seconds for holding in router list
 - Reachable time: Expected time neighbors remain reachable after advertising their media address (in milliseconds)
 - Reachable retransmission timer: Interval between successive solicitations of a neighbor that is not returning solicited neighbor advertisements (ms).
- + Source Link Layer option: contains media address of router

Neighbor discovery in IPv6 replaces ARP

- If there is no MAC address entry for the next hop, a *neighbor solicitation* message (comp. ARP-request) is sent:

Type = 135			Code = 0			Checksum		
R	S	O	Reserved					
Solicited address								
Options...								

- Hop count=1, own MAC address in *source link-level address* option
- The message is sent to a *solicited node multicast address* derived from the address of the next-hop. The MAC address is derived from this address.
 - FF02:0:0:0:0:1 + last 32 bits of address, 0x3333 + last 32 bits of multicast address
- The host recognizing its address, replies with a *neighbor advertisement* message (comp. ARP-reply)
 - Format similar, but Type=136
 - MAC address in *link layer address* option
 - R=1 if address is router, S=1 reply to solicitation, O=overrides previous cache entry

Redirect works like in IPv4 but may include the media address of the next hop

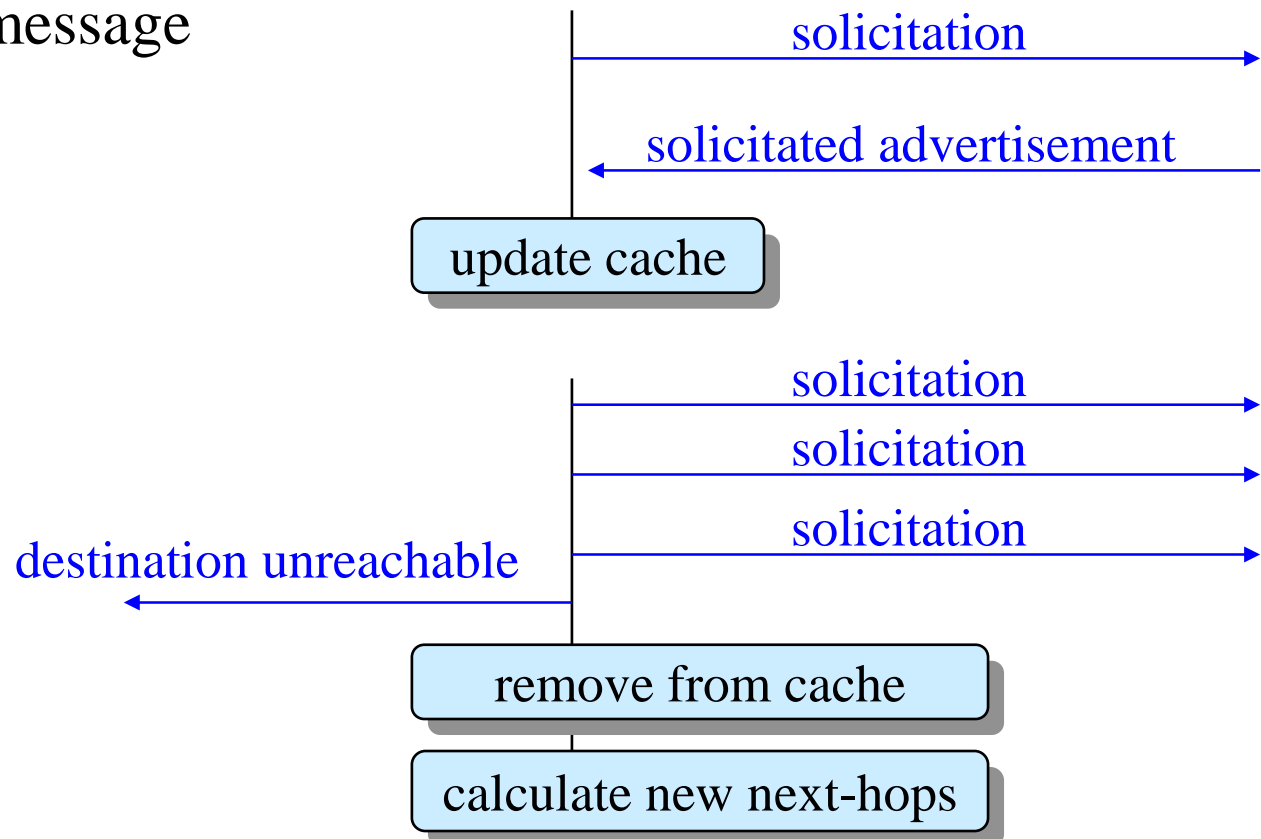
- Redirect message:

Type = 137	Code = 0	Checksum
Reserved		
Target address		
Destination address		
Options (e.g. target link layer address, redirected header option)		

- Target address contains the better next hop for the destination
- The media address of the next hop may be included in a *target link layer address* option.

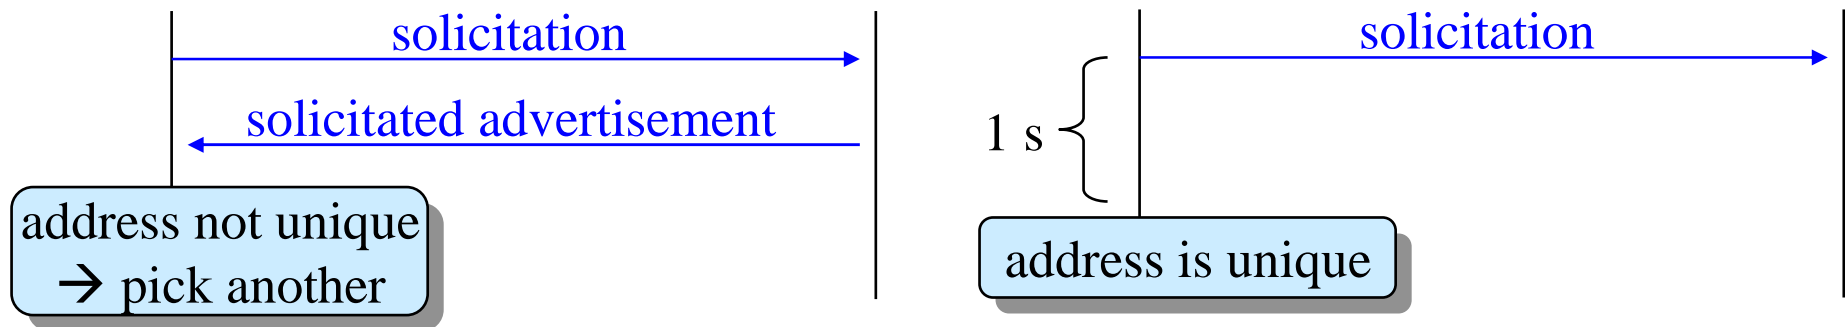
The sender needs feedback from the destination so that it does not send to a "black hole"

- If the sender does not get feedback (e.g. TCP acks) within 30 seconds, it checks the existence of the receiver with a solicitation message



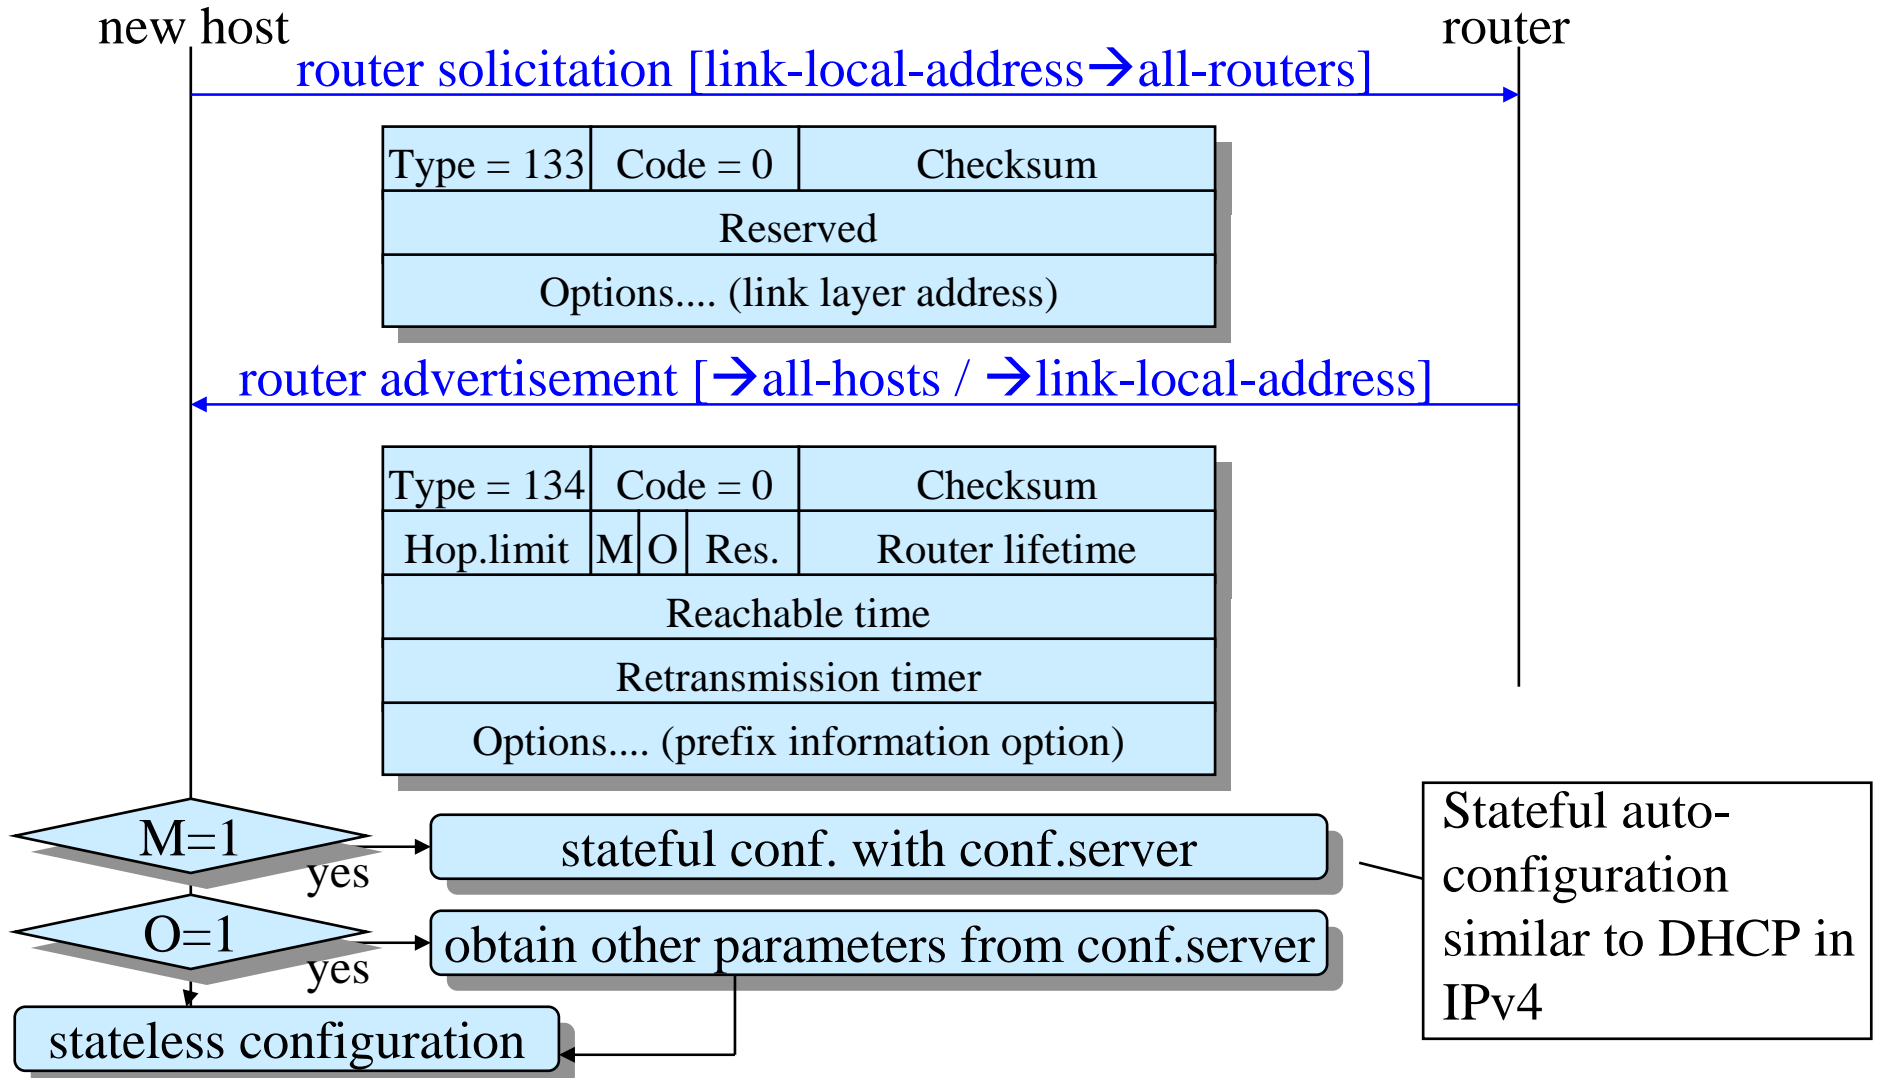
Before obtaining an address with autoconfiguration, the host uses a link local address

- FEB0:0000:0000:0000 + EUI-64 identifier
- The 64-bit EUI-64 identifier is generated from the 48-bit Ethernet address
- The host must check that the link local address is unique
 - In principle, addresses generated with the EUI-64 identifier should be unique, but...



- Lost messages \Rightarrow retry several times

Autoconfiguration can be stateful or stateless



Stateless autoconfiguration

Type = 134	Code = 0		Checksum	
Hop.limit	M	O	Res.	Router lifetime
Reachable time				
Retransmission timer				
Options.... (prefix information option)				

- Prefix information option contains list of prefixes with parameters
 - on-link bit → the prefix is specific to the local link
 - autonomous-bit → host can construct address by replacing the last bits of the prefix with EUI-64 identifier
- Stateless autoconfiguration properties
 - simple, no servers required
 - inefficient: 64 bits used for one local network
 - no access control

Mobile IP

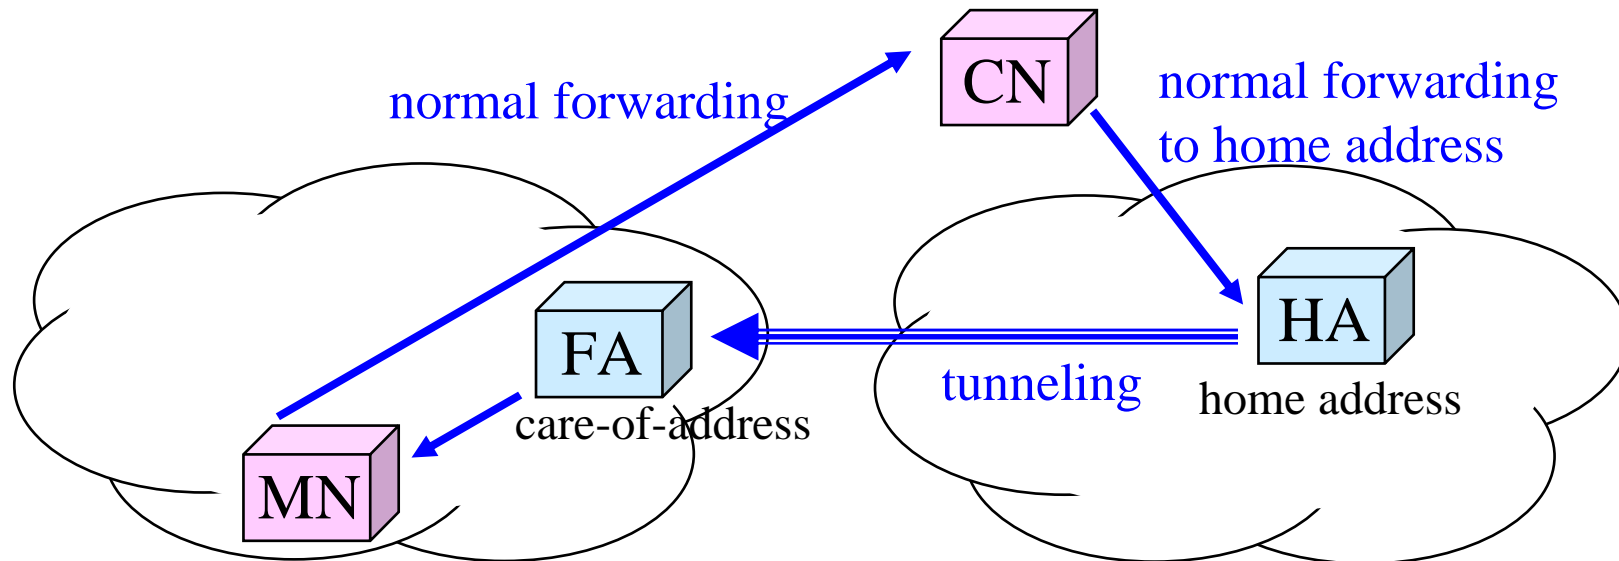
(Chapter 13 in Huitema)

Different types of mobility

- Computers transported and connected from different locations
 - Dynamic configuration \Rightarrow new IP address
 - Access through modem/ISDN
 - \Rightarrow new IP address
 - \Rightarrow TCP connection cut off
- Mobile computers, which stay connected during movements
 - Radio, infrared
 - \Rightarrow same IP address
- Mobile networks, e.g. in cars, planes, trains, ships
 - Recursive mobility (mobile host in mobile network)
- Mobility performed on lower layers, e.g. GPRS

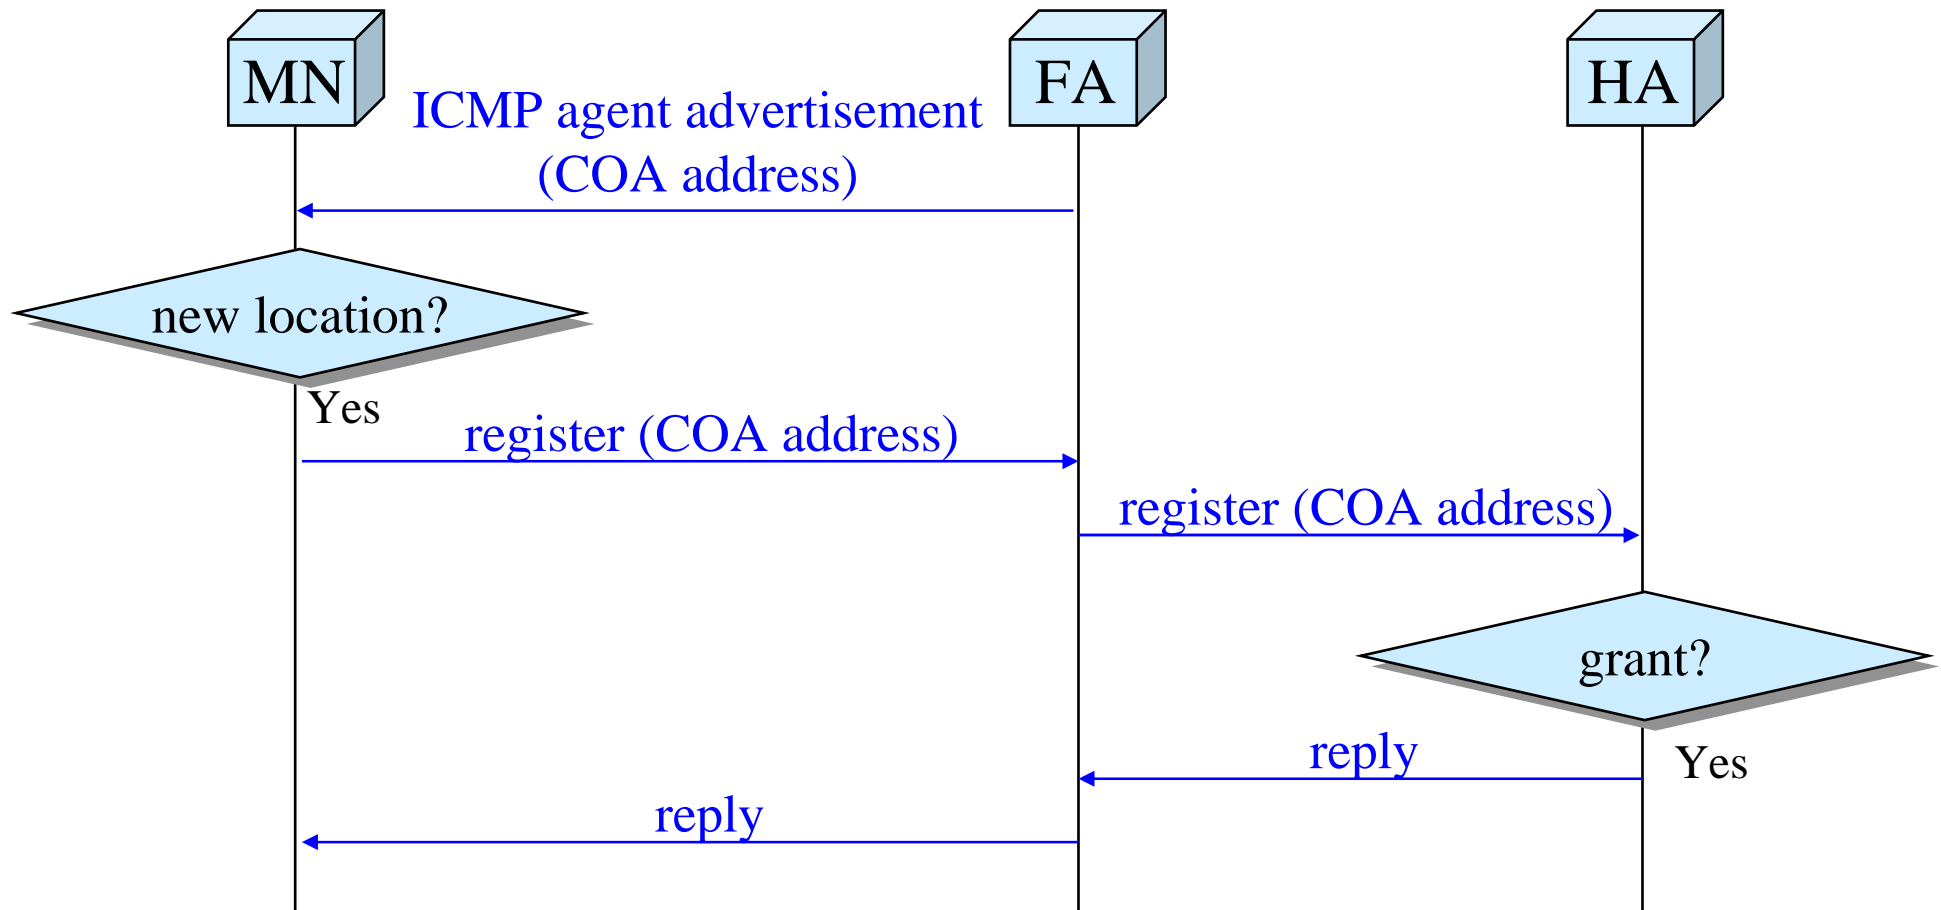
The traffic to a mobile node is tunneled from the home agent to the foreign agent

- **Mobile Node (MN)** – Node, who has a *home address* in the home network, and obtains a *care-of-address* (COA) in the visited foreign network
- **Home Agent (HA)** – Belongs to the home network and serves the home address
- **Foreign Agent (FA)** – Serves the visiting mobile node
- **Corresponding Node (CN)** – A node exchanging data with the mobile node



Home agents and foreign agents may be routers

When a host discovers that the location has changed, it must register the new COA with the HA




A lost request is resent by MN
FA never repeats the request.

Discovery of a Home Agent or Foreign Agent using periodical ICMP messages

- Agent advertisements are extensions to ICMP router advertisements
- The agent advertisements contain
 - Sequence number
 - Life-time of registration
 - Flags
 - Registration required
 - Foreign agent or home agent
 - Minimal encapsulation (RFC-2003)
 - Generic Routing Encapsulation (GRE) (RFC-1701)
 - Header compression used
 - List of care-of-addresses
 - Length of prefixes

The sequence numbers in the agent advertisement are similar to "lollipop" sequence numbers in OSPF



- If one of the number is < 256
 - The higher number is "higher"
- If both numbers are ≥ 256
 - If $(b-a) < (65635-256)/2$ then b is "higher"
- If the received is "lower" than the previous, then the server has been restarted
 - \Rightarrow Register again

Alternative discovery mechanisms

- Periodic broadcast of ICMP messages wastes transmission capacity, especially on wireless LANs
- The MN can detect changed location through media-level information
 - e.g. analyzing power of different basestations
- Instead of waiting, the MN can solicit the information
 - Similar to ICMP router solicitation
 - TTL = 1
 - Agent replies with agent advertisement


Registration request

- Registration request message contains
 - Message type = 1
 - Flags
 - FA co-located with MN
 - preferred encapsulation
 - Requested lifetime
 - 0 = cancellation of previous
 - Home address of MN
 - HA address
 - COA address
 - 64-bit request identification
 - Extensions
 - E.g. authentication


Registration reply

- Registration reply message contains
 - Message type = 3
 - Reply code (granted or denied)
 - Who denied (FA or HA)
 - Why denied
 - Accepted lifetime
 - Same or smaller than requested lifetime
 - Home address of MN
 - HA address
 - 64-bit request identification
 - Same as in request
 - Extensions
 - E.g. authentication

Security issues (1)

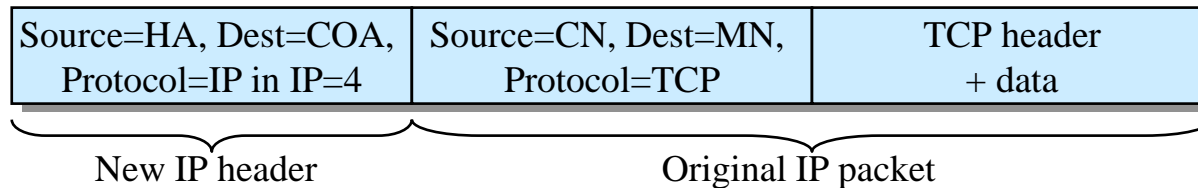
- Attack types
 - Attacker pretends to be a FA to capture traffic 
 - Attacker replays old registration messages
- Authentication extension proves the origin of the message and that the contents has not been changed
 - Security parameter index (SPI) together with HA, COA, or NM identifies security context
 - Shared secret, signature algorithm (e.g. keyed MD5) parameters of security context
 - Data and secret key → authentication field
 - MN to HA authentication mandatory
 - FA to HA and MN to FA authentications optional

Security issues (2)

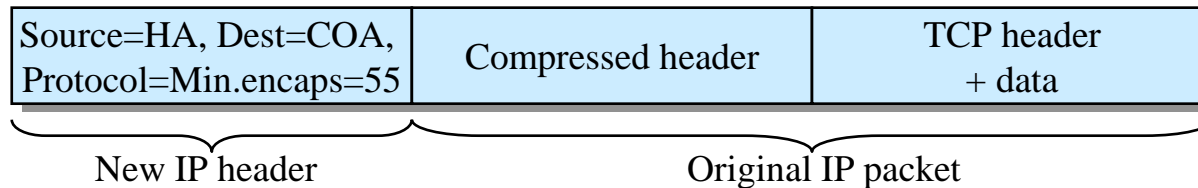
- Attack types
 - Attacker pretends to be a FA to capture traffic
 - Attacker replays old registration messages 
- Two requests must not contain the same identification
 - NTP timestamps (64-bit)
 - Only requests with higher timestamps are accepted
 - The timestamps must be close to the current time
 - Random numbers used only once (nonce)

Encapsulation

- Basic encapsulation, RFC-2003




- Minimal encapsulation, RFC-2004



Compressed header:
Protocol type of encaps. packet
(e.g. TCP), Destination address of
encaps. packet, Optional source
address of encaps. packet, Header
checksum

- Generic Routing Encapsulation (GRE), RFC-1701





Parameters: Protocol type (similar to the one in Ethernet packet), optional checksum, optional sequence number, optional authentication key, (source) routing field, flags (which options are present)

Broadcast and multicast should only be received by the MN, not the network of MN

- Easy if FA is co-located with MN



- Double encapsulation of broadcast/multicast traffic




- Joining multicast groups: ICMP messages are tunneled MN→HA
- More efficient: MN can subscribe to groups on the foreign network

Source address filtering is a problem in Mobile IP (1)

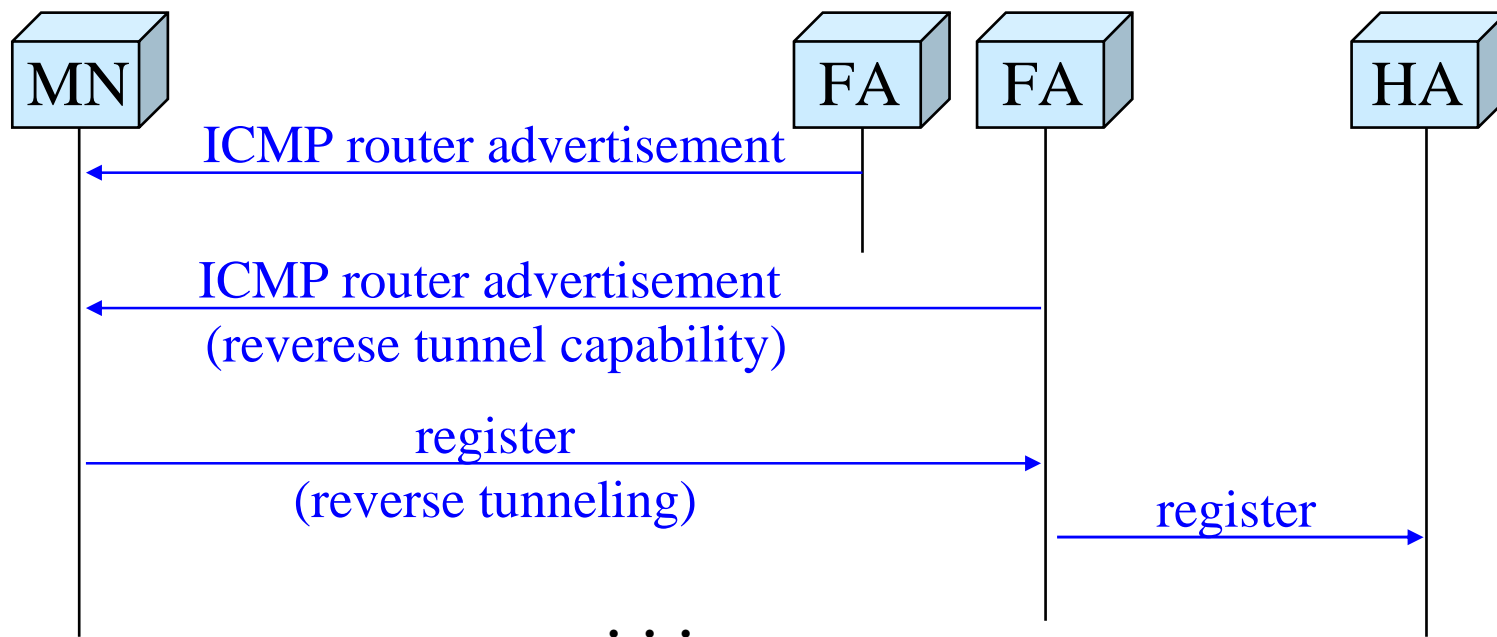
- Why source address filtering?
 - Address spoofing hides identity of attacker, helps targeting third parties' replies, helps gaining privileges
- Source address filtering is performed in firewalls, between ISP and customer, at peering points between providers, etc.

⇒ Packets sent by MN must be tunneled through the HA



Source address filtering is a problem in Mobile IP (2)

- FAs capable of tunneling packets back to HA, advertise it with a flag in agent advertisement message
- The MN requests reverse tunneling



Considerations

- Path $MN \rightarrow CN$ is shorter than the path $CN \rightarrow MN$
 - Asymmetry
- If the MN moves relatively fast, it must choose a new FA often
 - \Rightarrow Many registration messages to HA

Mobile IPv6

(Chapter 13 in Huitema)

Mobility in IPv6

- Discovery performed with IPv6 neighbor discovery and address configuration mechanisms
- Security \Rightarrow MN can notify their COA to the CN in addition to the HA
- Efficient encapsulation with the source routing header

Discovery

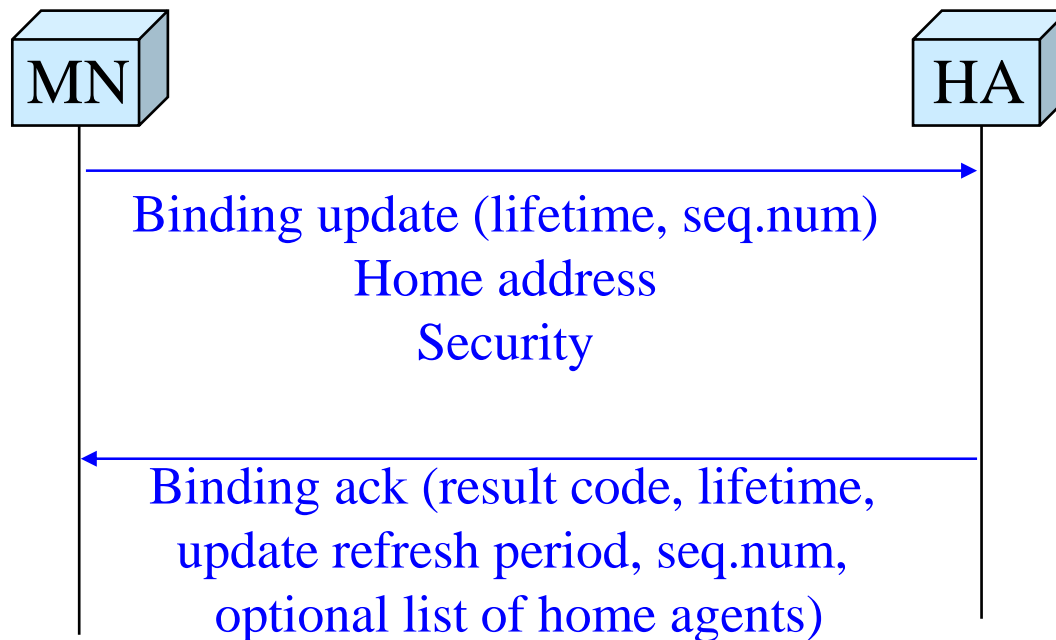
- The MN and FA are usually colocated \Rightarrow No separate FA
- Hosts listen to router advertisements to learn prefixes of the link
 - Hosts can detect that they are visiting a foreign network
- COA obtained with address configuration procedures
- Routers willing to act as home agents indicate it in the router advertisement

Binding updates (1)

- Binding performed using *destination options*
 - **Binding update** – informs about the new COA
 - **Binding ack** – acknowledges the COA
 - **Binding request** – To request information about the current COA
 - **Home address** – Identifies the home address of the MN
- Authentication with the *security option*

Binding updates (2)

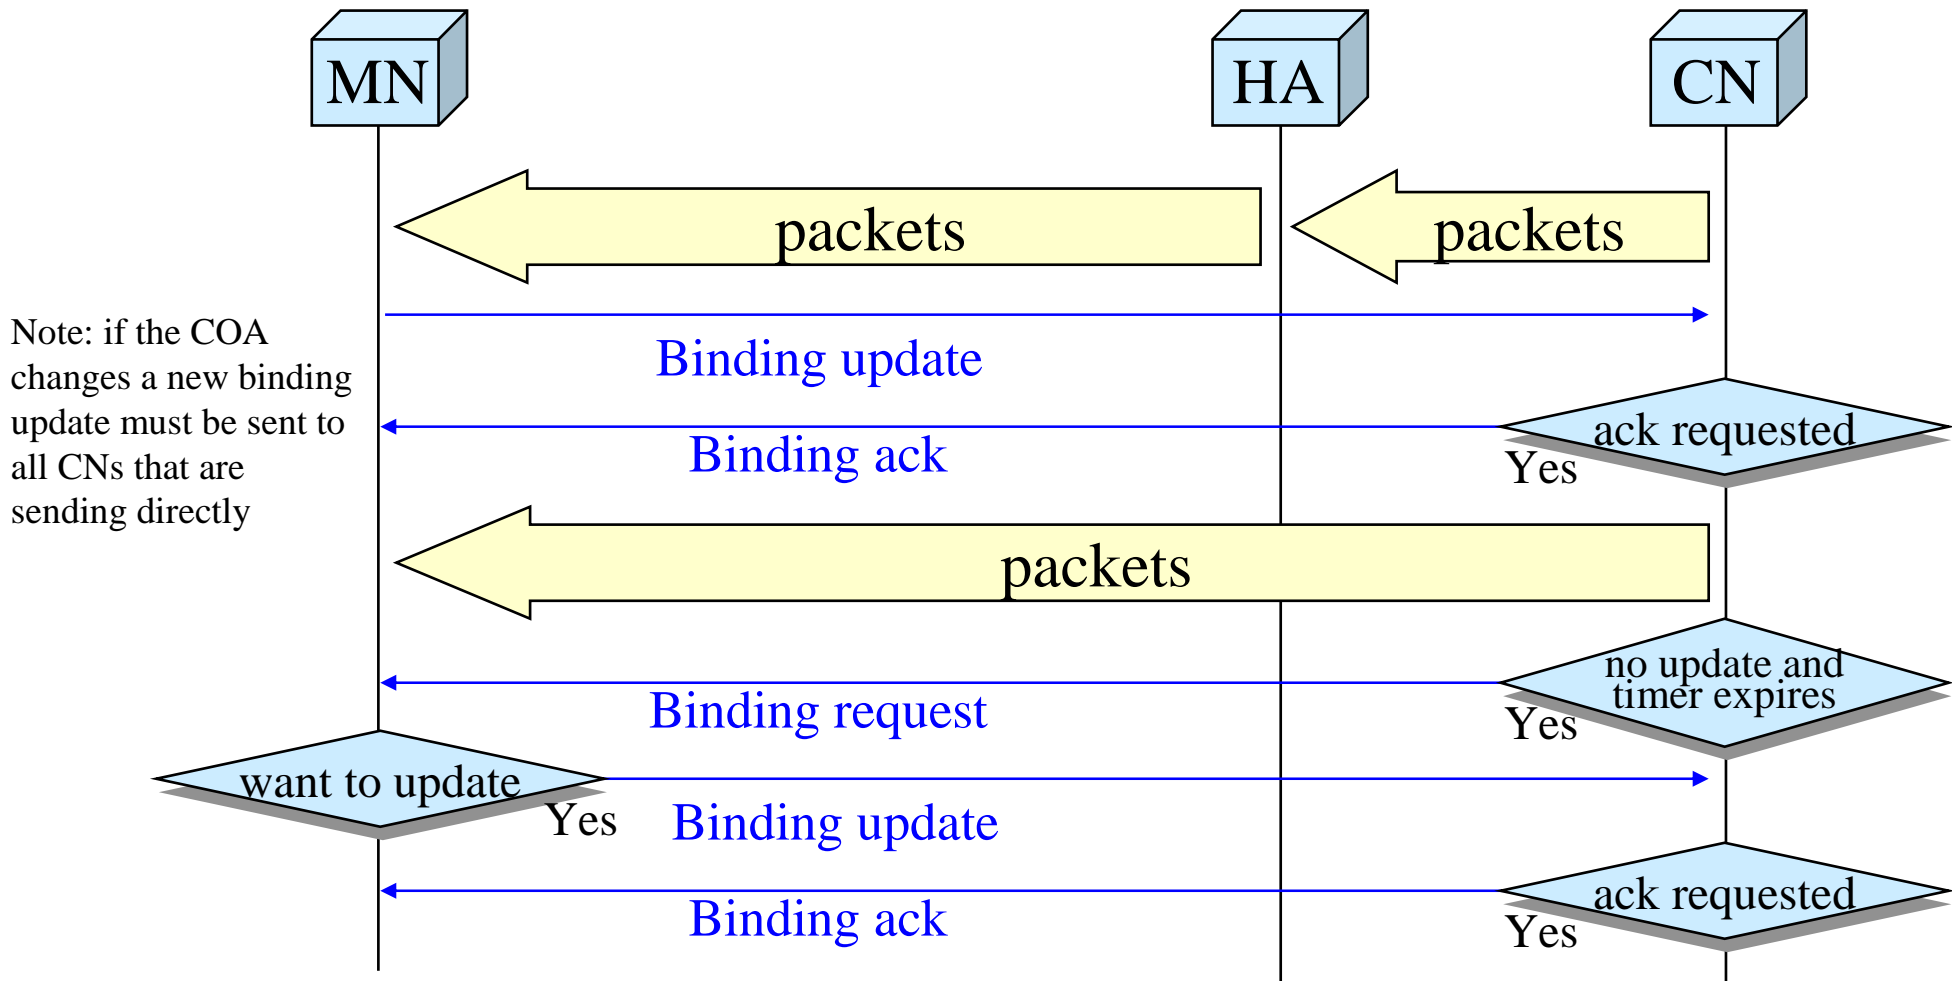
- COA transmitted in source address of IPv6 header
- Home address in the *Home Address option*



Source address filtering is not a problem in IPv6

- The mobile node does not put its home address in the IPv6 header. Instead, the home address is sent in the Home Address option. The IPv6 header contains the COA.
- Mandatory requirement.

The MN can send a binding update to the CN to optimize the route



IPv6 uses the routing header instead of encapsulation

