

Network Fundamentals

Matsuzaki 'maz' Yoshinobu

<maz@ij.ad.jp>

IPv4 and IPv6

- almost the same
 - forwarding architecture
 - a bit different on IP header
- address length
 - IPv4 32-bit long (4 octets)
 - IPv6 128-bit long (16 octets)

address notation

- IPv4
 - 255.255.255.255
 - 4 fields of 1octet in decimal
 - separated by a dot [.]
- IPv6
 - ffff:ffff:ffff:ffff:ffff:ffff:ffff:ffff
 - 8 fields of 2octets in hexadecimal
 - separated by a colon [:]


it's just a number

- you can use different notations if your application accepts it
- IPv4
 - ping 2130706433
 - ping 017700000001
 - ping 127.0.0.1
- IPv6
 - ping ::1
 - ping ::0.0.0.1

IPv4 packet forwarding

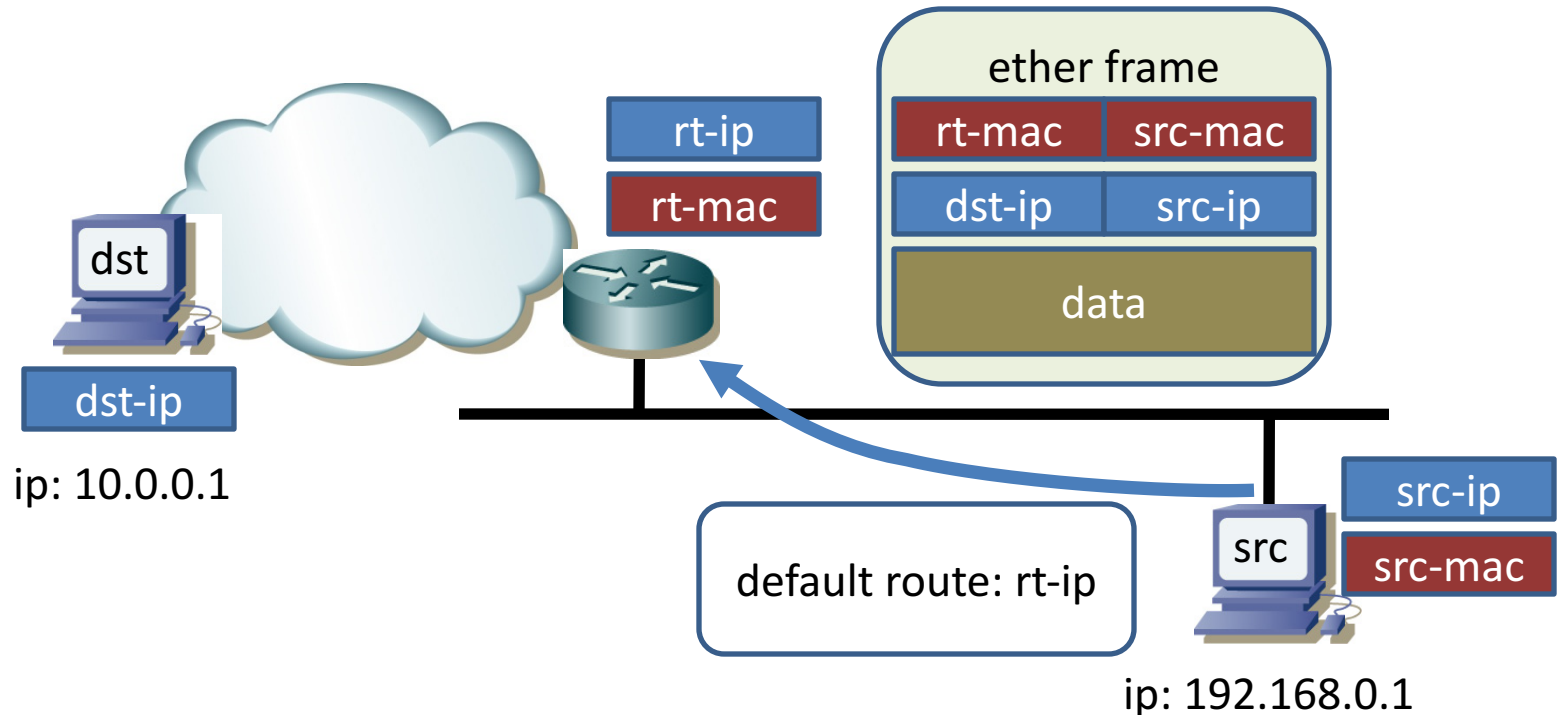
- send directly if the destination is on the same network

```
inet 192.168.0.1 netmask 255.255.255.0
      :
192.168.0.0 - 192.168.0.255 are on the link
```



IPv4 packet forwarding

- send to router based on routing information if destination is not on the same network



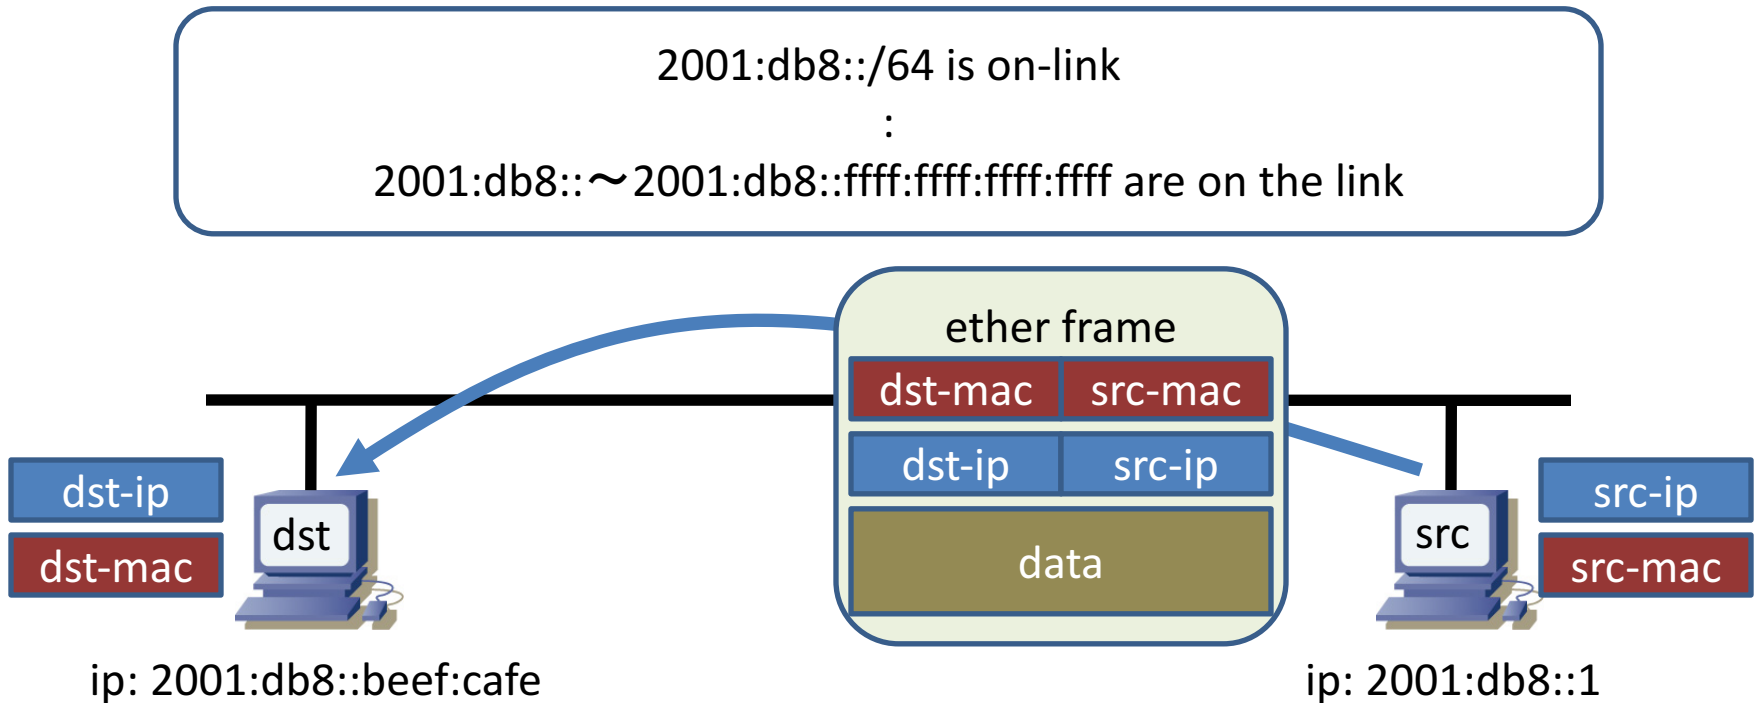
arp (Address Resolution Protocol)

- MAC address is needed to forward a packet over Ethernet
 - IPv4 address to MAC address resolution
- arp
 - RFC826

```
arp who-has 192.168.0.2 tell 192.168.0.1
0x0000:  ffff ffff ffff 0019 bb27 37e0 0806 0001
0x0010:  0800 0604 0001 0019 bb27 37e0 c0a8 0001
0x0020:  0000 0000 0000 c0a8 0002
arp reply 192.168.0.2 is-at 00:16:17:61:64:86
0x0000:  0019 bb27 37e0 0016 1761 6486 0806 0001
0x0010:  0800 0604 0002 0016 1761 6486 c0a8 0002
0x0020:  0019 bb27 37e0 c0a8 0001 0000 0000 0000
0x0030:  0000 0000 0000 0000 0000 0000
```


IPv6 packet forwarding

- send directly if the destination is on-link



IPv6 packet forwarding

- send to router based on routing information if destination is not on the same link



ndp (Neighbor Discovery Protocol)

- MAC address is needed to forward a packet over Ethernet
- ndp resolution [RFC4861]
 - query MAC address using ICMP6
 - destination is Solicited-Node multicast address
 - IP: ff02::1:ff00:0000 ~ ff02::1:ffff:ffff
 - generated by lower 24-bit of the destination address
 - MAC: 33:33:00:00:00:00 ~ 33:33:ff:ff:ff:ff
 - generated by lower 32-bit of the destination address

MAC address resolution by ndp

```
IP6 2001:db8::1 > ff02::1:ffef:cafe
```

```
ICMP6, neighbor solicitation, who has 2001:db8::beef:cafe  
source link-address option: 00:19:bb:27:37:e0
```

```
0x0000: 3333 ffe0 cafe 0019 bb27 37e0 86dd 6000  
0x0010: 0000 0020 3aff 2001 0db8 0000 0000 0000  
0x0020: 0000 0000 0001 ff02 0000 0000 0000 0000  
0x0030: 0001 ffe0 cafe 8700 9a90 0000 0000 2001  
0x0040: 0db8 0000 0000 0000 0000 beef cafe 0101  
0x0050: 0019 bb27 37e0
```

```
IP6 2001:db8::beef:cafe > 2001:db8::1
```

```
ICMP6, neighbor advertisement, tgt is 2001:db8::beef:cafe  
destination link-address option: 00:16:17:61:64:86
```

```
0x0000: 0019 bb27 37e0 0016 1761 6486 86dd 6000  
0x0010: 0000 0020 3aff 2001 0db8 0000 0000 0000  
0x0020: 0000 beef cafe 2001 0db8 0000 0000 0000  
0x0030: 0000 0000 0001 8800 c1fd 6000 0000 2001  
0x0040: 0db8 0000 0000 0000 0000 beef cafe 0201  
0x0050: 0016 1761 6486
```

IPv6 history

- 1990 concerns about Routing and Class B exhaustion
- 1991 IETF Routing and Addressing group
 - to solve IP address exhaustion and routing explosion
 - 1992 RFC1380 introducing CIDR and recommendation of new internet protocol
- 1992-1994 proposals for IPng
 - 1992 IAB proposed IPv7 based on CLNP, but rejected
 - 1992 IPng proposals: TUBA, CATNIP, SIPP, etc
 - 1994 based on SIPP, 128bit long IPng was adopted
- 1994-1996 IPng basic spec
 - 1995 RFC1883 IPv6 Spec

IPv6 spec

- 128-bit long address space
- multiple IP addresses on a interface
 - node has a policy of source address selection
- scope
 - linklocal, global

IPv6 address notation

- hexadecimal digit of the eight 16-bit pieces of the address separated by “:”
 - 2001:0db8:0000:0000:0000:0000:0000:0001
 - suppress leading zeros
 - use “::” to shorten 2 or more groups of zeros
 - 2001:db8::1
- **RFC5952**
 - A Recommendation for IPv6 Address Text Representation

IPv6 prefix notation

- the same as IPv4
 - the first IP address of the ip block, and
 - trailing “/” and decimal digit of prefix length
- 2001:db8::/64
 - 2001:db8:: - 2001:db8::ffff:ffff:ffff:ffff

IPv6 addressing

- Unicast
 - an identifier for a single interface. A packet sent to an unicast address is delivered to the interface identified by that address
- Multicast
 - an identifier for a group of interfaces. A packet sent to a multicast address is delivered to all interfaces identified by that address

unicast address

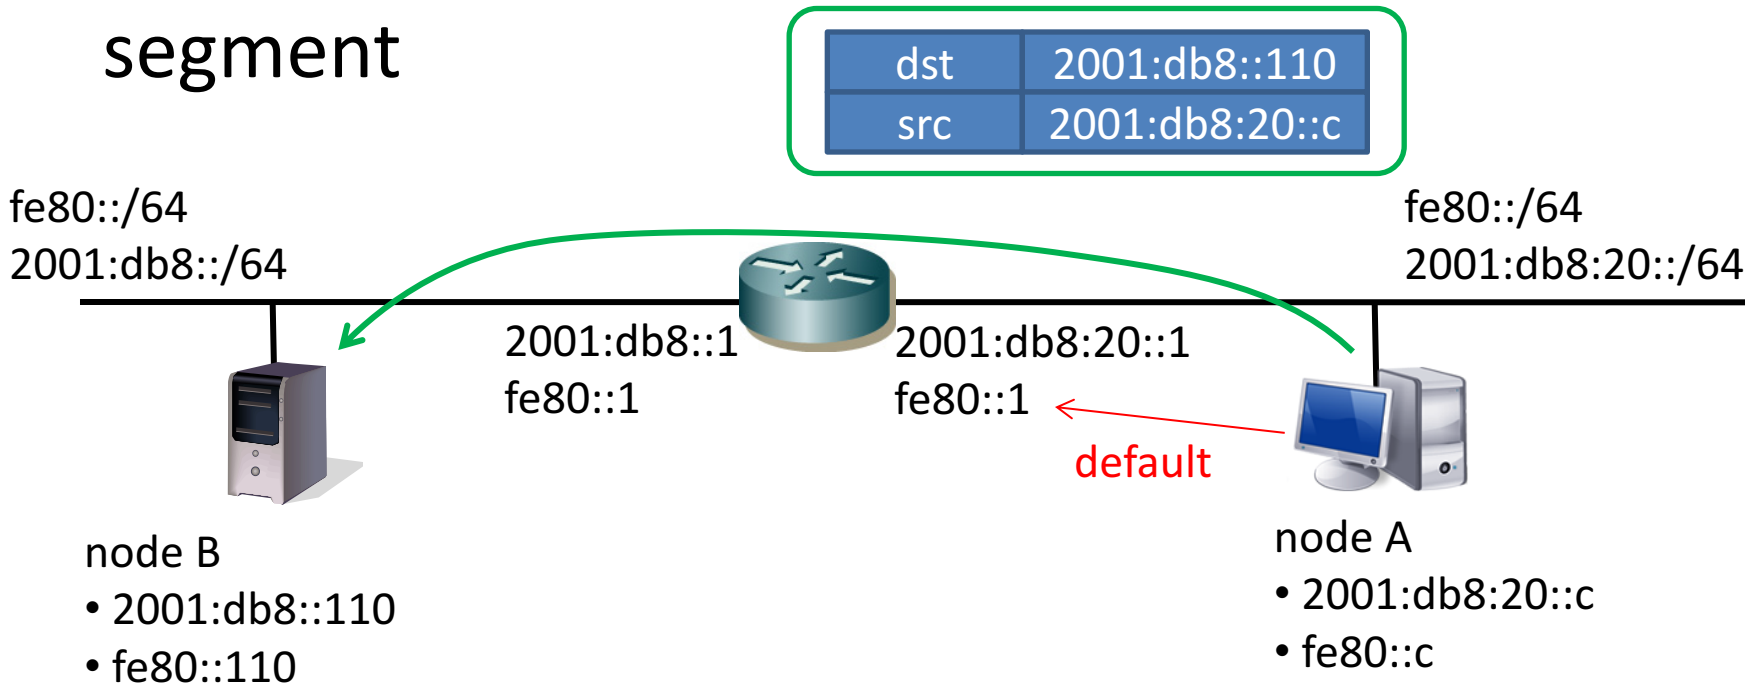
- link-local
 - for use on a single link
 - fe80::/64 is used
- global unicast
 - other than special addresses

link-local address

- for use on a single link
 - an address from fe80::/64 is assigned for an interface
 - routers **MUST** not forward any packet with link-local source or destination addresses to other links
- for automatic address configuration, routing
 - communication between nodes on the same link
 - as an immediate nexthop of ip routing

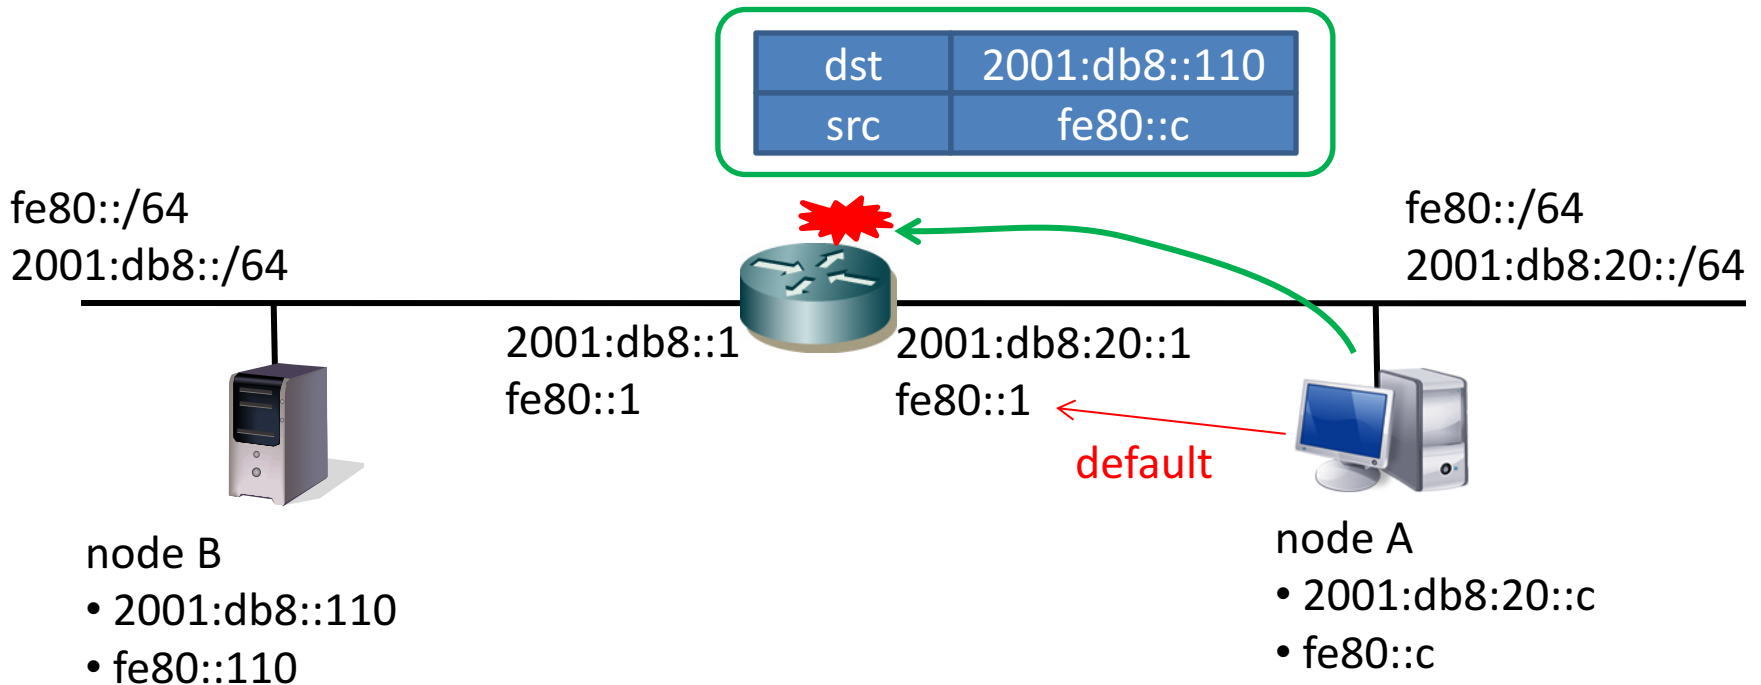
inter-segment communication

- a packet has global unicast source addresses and destination addresses in the other segment



routers must not forward link-locals

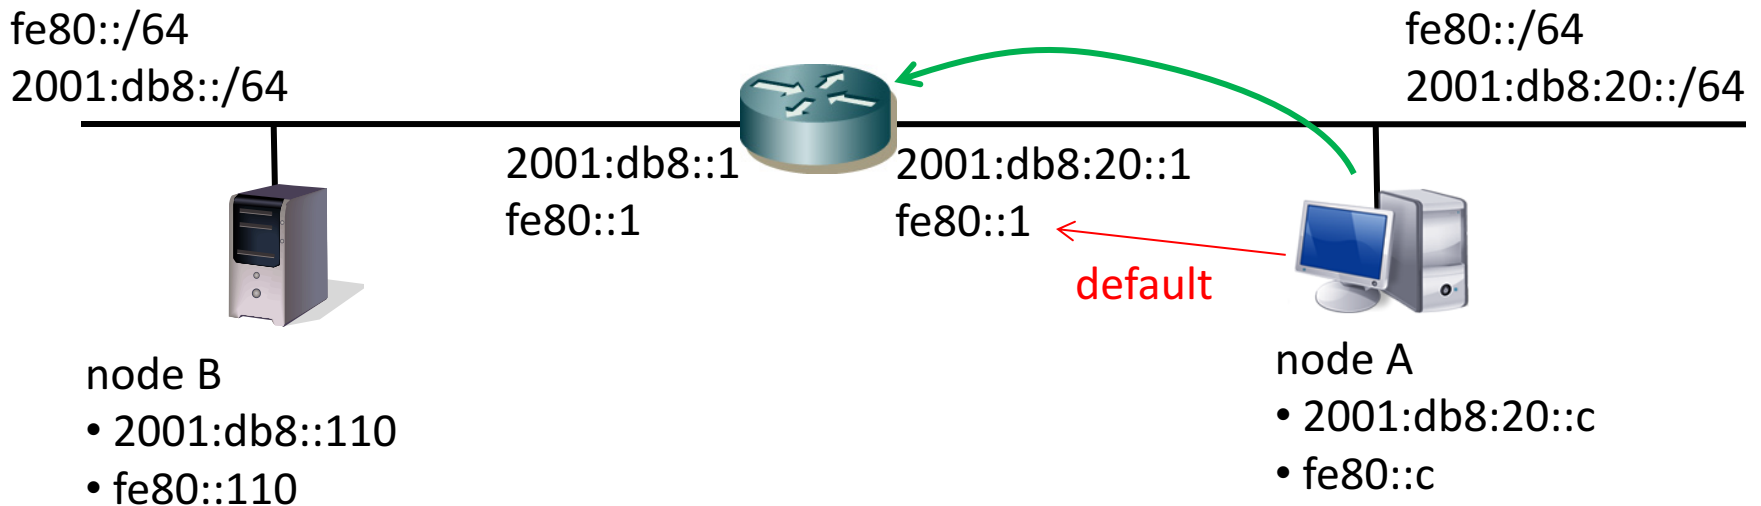
- Routers must not forward a packet with a link-local source address



link-local for a single link

- Routers receive a packet with a link-local source address if the destination is themselves

dst	fe80::1
src	fe80::c



multicast address

- could be a group of nodes
- All nodes on a link
 - ff02::1
- Solicited-Node multicast address
 - ff02::1:ff00:0/104
 - generate using lower 24-bit of an unicast address
 - 2001:db8::1234:5678 -> ff02::1:ff34:5678

allocated space

Global Unicast

- 2000::binary (001x xxxx xxxx xxxx xxxx)

Unique Local Unicast

- fc00::binary (1111 110x xxxx xxxx xxxx)

Link Local Unicast

- fe80::binary (1111 1110 10xx xxxx xxxx)

Multicast

- ff00::binary (1111 1111 xxxx xxxx xxxx)

<http://www.iana.org/assignments/ipv6-address-space>

special addresses

2001:db8::/32

- document address

::

- unspecified address

::1

- loopback address

ff02::1

- All IPv6 Nodes on a link


ff02::2

- All IPv6 Routers on a link

IPv6 address structure

- 128-bit long address space
- basically /64 anywhere
 - upper 64-bit is prefix
 - lower 64-bit is interface identifier

2001:0db8:0000:0000:0000:0000:0000:0001



interface identifier

- to identify an interface on a link
 - unique within a subnet prefix
- several ways to auto-generate the identifier
 - modified EUI-64 format

MAC address - 00:19:bb:27:37:e0 (example)

1. separate to two 24-bit pieces
2. insert "ff:fe"
3. invert universal/local bit(modified from EUI-64)

modified EUI-64 format - 0219:bbff:fe27:37e0

- randomly generated 64-bit long address

for “plug and play”

- IPv6 address = **prefix** + **interface identifier**
 - interface identifier can be generated automatically
 - prefix is needed
 - link-local (fe80::/64) is available
- generating a link-local address
 - MAC address - 00:19:bb:27:37:e0 (example)
 - modified EUI-64 - 0219:bbff:fe27:37e0
 - tentative address - fe80::0219:bbff:fe27:37e0

Duplicate Address Detection (DAD)

- to test for uniqueness of the unicast address
 - source address “::”
 - destination address “Solicited-Node multicast”

```
IPv6 :: > ff02::1:ff27:37e0
ICMP6, neighbor solicitation,
who has fe80::219:bbff:fe27:37e0
  0x0000:  3333 ff27 37e0 0019 bb27 37e0 86dd 6000
  0x0010:  0000 0018 3aff 0000 0000 0000 0000 0000
  0x0020:  0000 0000 0000 ff02 0000 0000 0000 0000
  0x0030:  0001 ff00 0001 8700 5764 0000 0000 fe80
  0x0040:  0000 0000 0000 0219 bbff fe27 37e0
```

Router Advertisement (RA)

- Routers advertise RA by ICMP6
 - a host can ask routers to advertise if necessary

```
IP6 fe80::20b:fdff:fed5:e17f > ff02::1
```

```
ICMP6, router advertisement,
```

```
  prefix info option: 2001:db8::/64,
```

```
  valid time 2592000s
```

```
0x0000:  3333 0000 0001 000b fdd5 e17f 86dd 6e00
```

```
0x0010:  0000 0040 3aff fe80 0000 0000 0000 020b
```

```
0x0020:  fdff fed5 e17f ff02 0000 0000 0000 0000
```

```
0x0030:  0000 0000 0001 8600 3aa1 4000 0708 0000
```

```
0x0040:  0000 0000 0000 0101 000b fdd5 e17f 0501
```

```
0x0050:  0000 0000 05dc 0304 40c0 0027 8d00 0009
```

```
0x0060:  3a80 0000 0000 2001 0db8 0000 0000 0000
```

```
0x0070:  0000 0000 0000
```

RA informs

- a default router
- prefixes on the link and its lifetime
- a host can communicate with other nodes
 - how to configure a DNS resolver on the host
 - DHCPv6
 - RA option for DNS configuration
 - use IPv4 DNS resolver if the host already has

summary of IPv6 plug & play

- RA – stateless autoconfiguration
 - RFC4862 IPv6 Stateless Address Autoconfiguration
 - RFC6106 IPv6 RA Option for DNS Configuration
- DHCPv6 – stateless/full autoconfiguration
 - RFC3315 DHCPv6
 - RFC3646 DNS Configuration options for DHCPv6

steps of address auto-configuration

1. auto-generation of link-local address

example) fe80::219:bbff:fe27:37e0

2. Duplicate Address Detection (DAD)

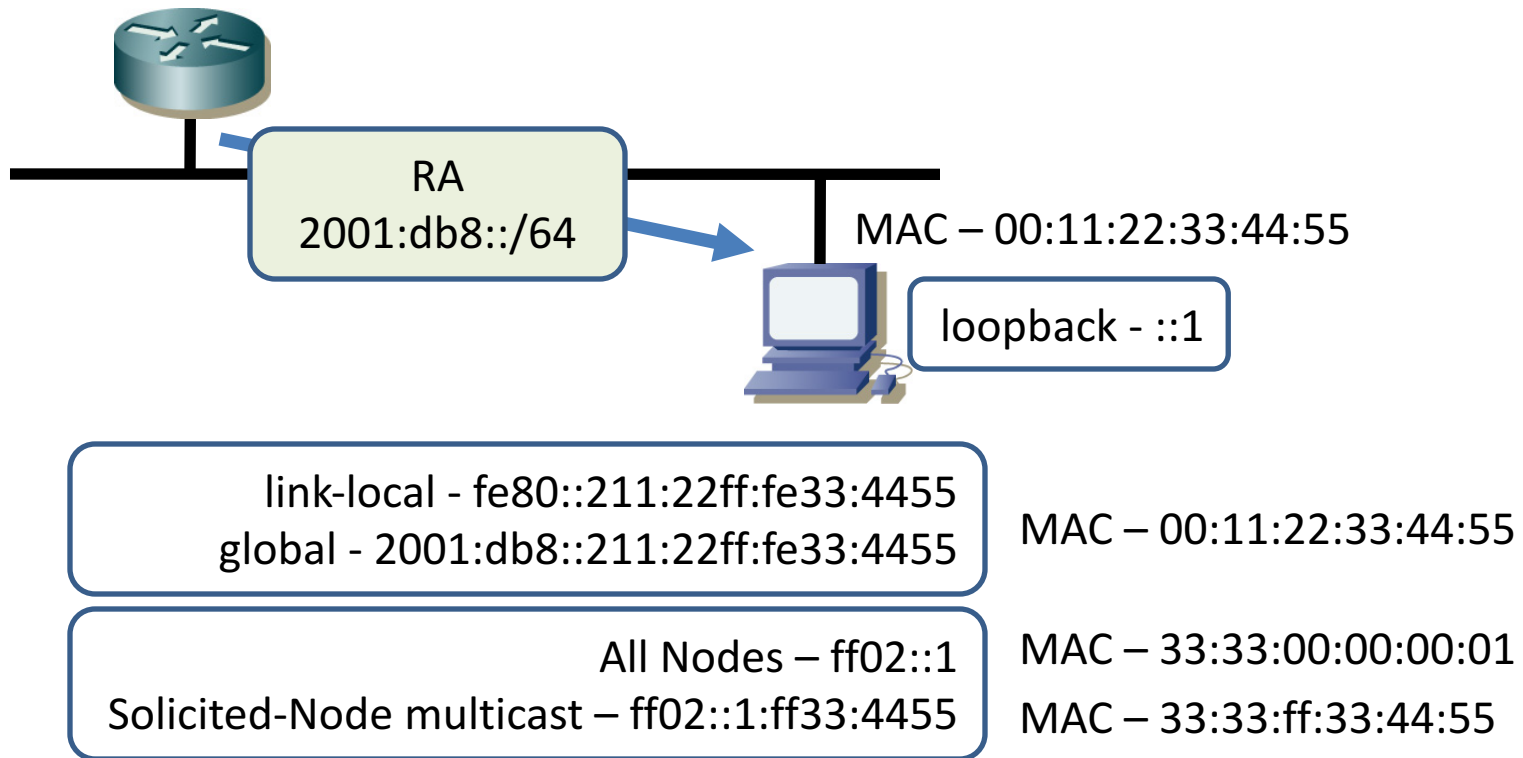
3. accept Router Advertisement (RA)

example) 2001:db8::/64

4. auto-generation of global unicast address

example) 2001:db8::219:bbff:fe27:37e0

IPv6 addresses of a host



address and interface

- local scope and multiple interfaces
 - “ping6 ff02::1” is not deterministic if a host has 2 or more interfaces
 - several OSes assume a ‘default’ interface
- address notation with outbound interface
 - <ipv6 address>%<outgoing interface>
 - example) ping ff02::2%8
 - example) ping6 ff02::1%em0

Exercise 1

- Try the followings
- IPv4
 - ping 2130706433
 - ping 017700000001
 - ping 127.0.0.1
- IPv6
 - ping ::1
 - ping ::0.0.0.1

Exercise 2

- What's your IP address(es)
- tips:
 - Windows: `ipconfig /all`
 - Unix: `ifconfig -a`

Exercise 3

- ping your IPv6 linklocal address
- tips:
 - Windows: ping <ipv6>%<zone id>
 - Unix: ping6 <ipv6>%<outgoing if>