

IPv6 and Out-of-Box Experience of Network Configuration

Hiroki Sato <hrs@FreeBSD.org>

FreeBSD Developer Summit, June 2021

Outline

- Using IPv6 is not rare these days, finally
 - Deployment scenarios are not well-known compared to IPv4. What is normal?
 - On FreeBSD, what is supported and what is not?
- Where can we improve "how to configure"?
 - rc.conf(5) sometimes becomes cryptic
- Please share your good/bad experiences

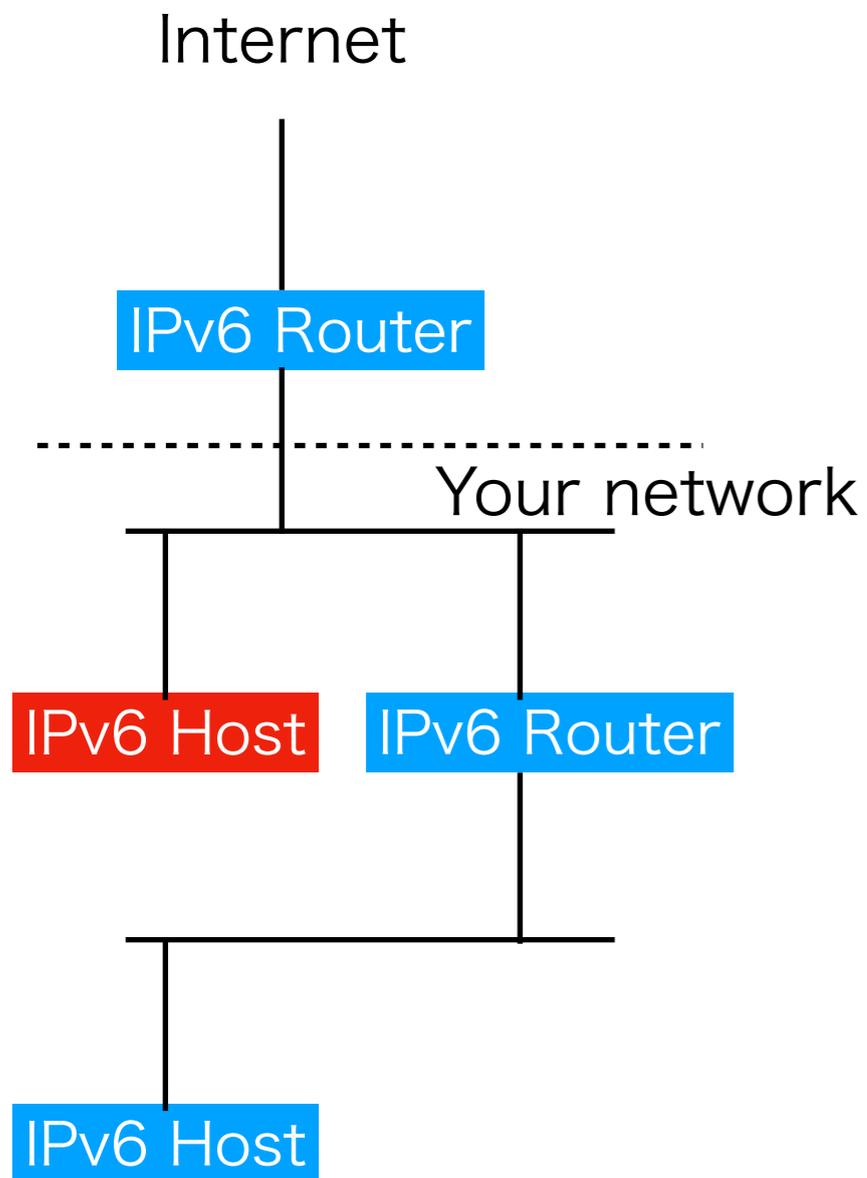
IPv6

- Over 20 years old (RFC1883 in 1995)
- Core protocol is mature, but the deployment is still changing:
 - Automatic configuration (SLAAC, DHCPv6)
 - Privacy/security enhancement
 - Multi-prefix environment

2001:0db8:0000:0000:0001:0000:0000:4444

- 128-bit long, in hexadecimal (16-bit field x 8)
RFC4291: "IP Version 6 Addressing Architecture"

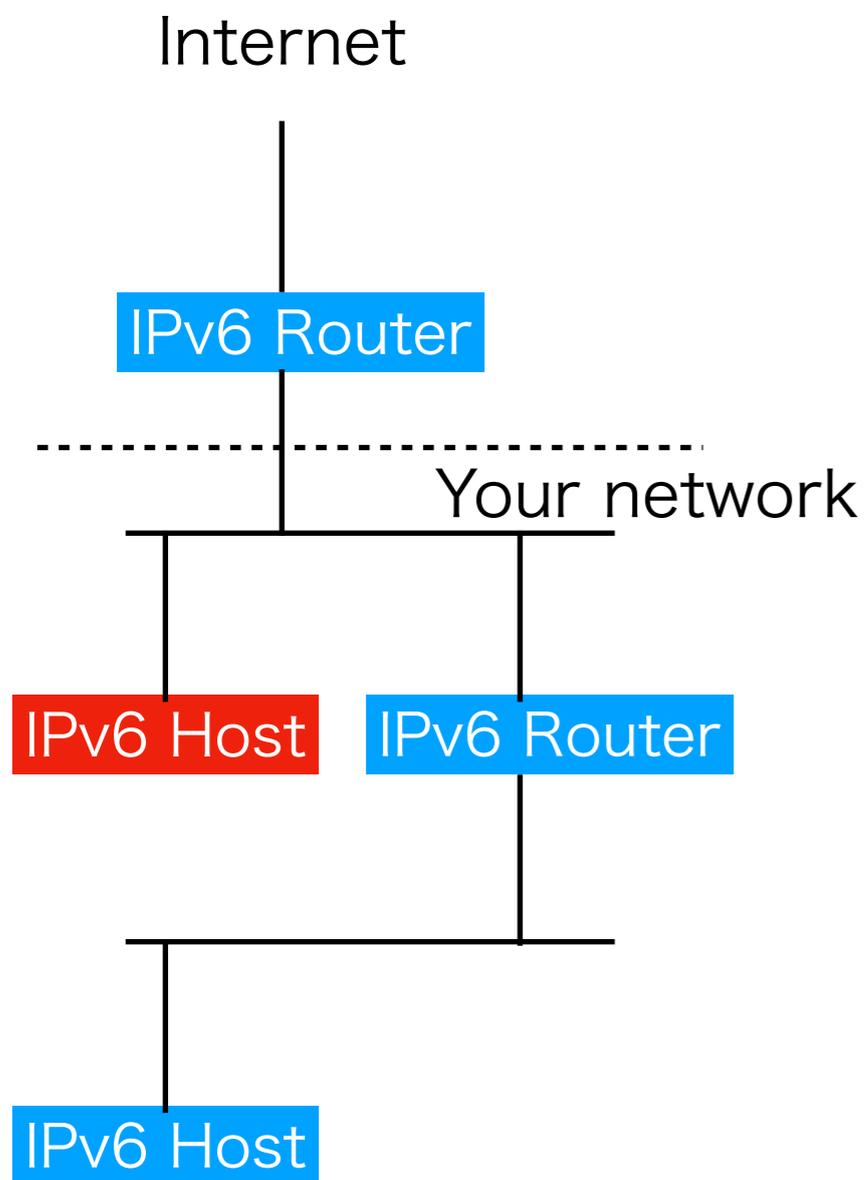
4 Config Scenarios



- A: Manual Configuration
- B: SLAAC (native)
- C: SLAAC + DHCPv6 (native)
- D: PPPoE (IPv6CP) + DHCPv6

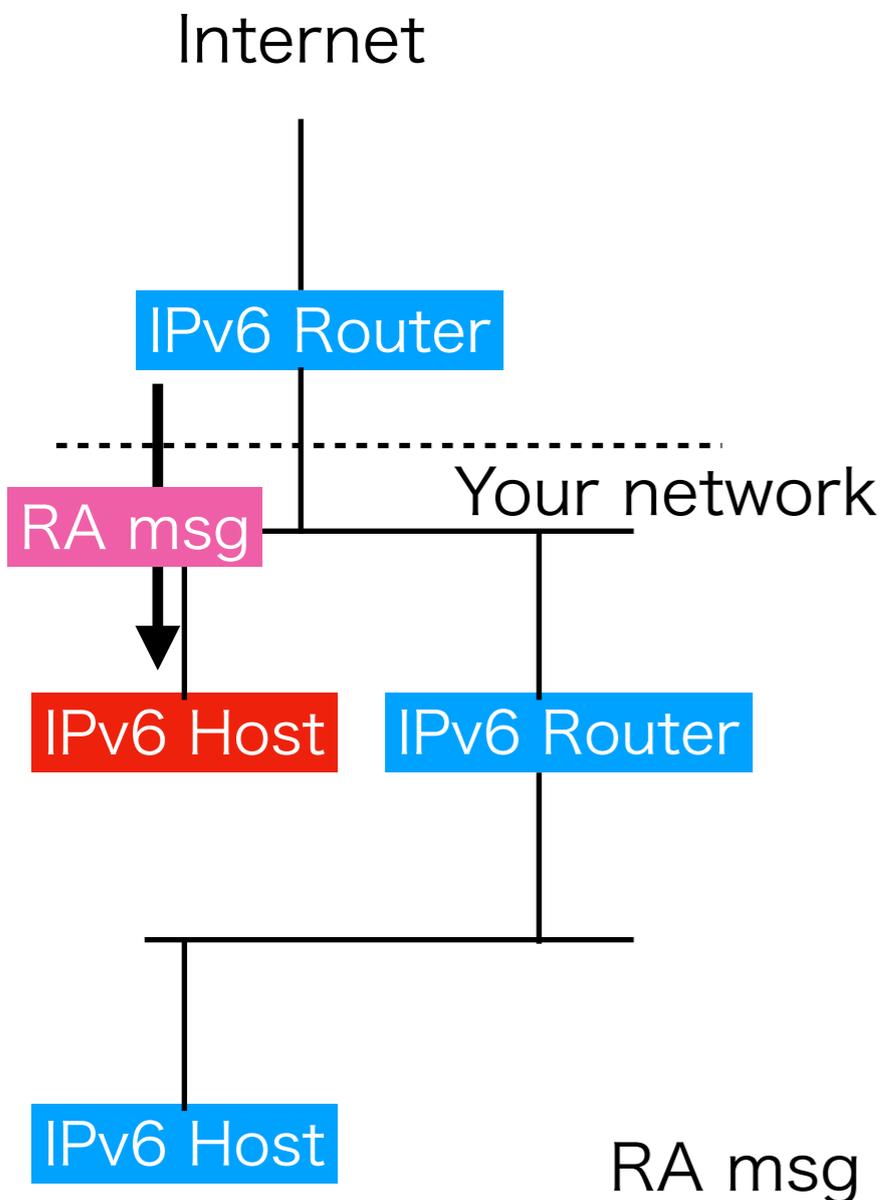
SLAAC: StateLess Address AutoConfiguration

4 Config Scenarios



- A: Manual Configuration
 - `ifconfig_em0_ipv6="inet6 ..."`
 - `ifconfig_em0_alias0="inet6 ..."`
 - `"_ipv6"` is the per-if knob to declare to use IPv6. If missing, IPv6 capability of the interface is disabled by `rc.d(8)` script.
- `ipv6_defaultrouter, ipv6_gateway`

4 Config Scenarios



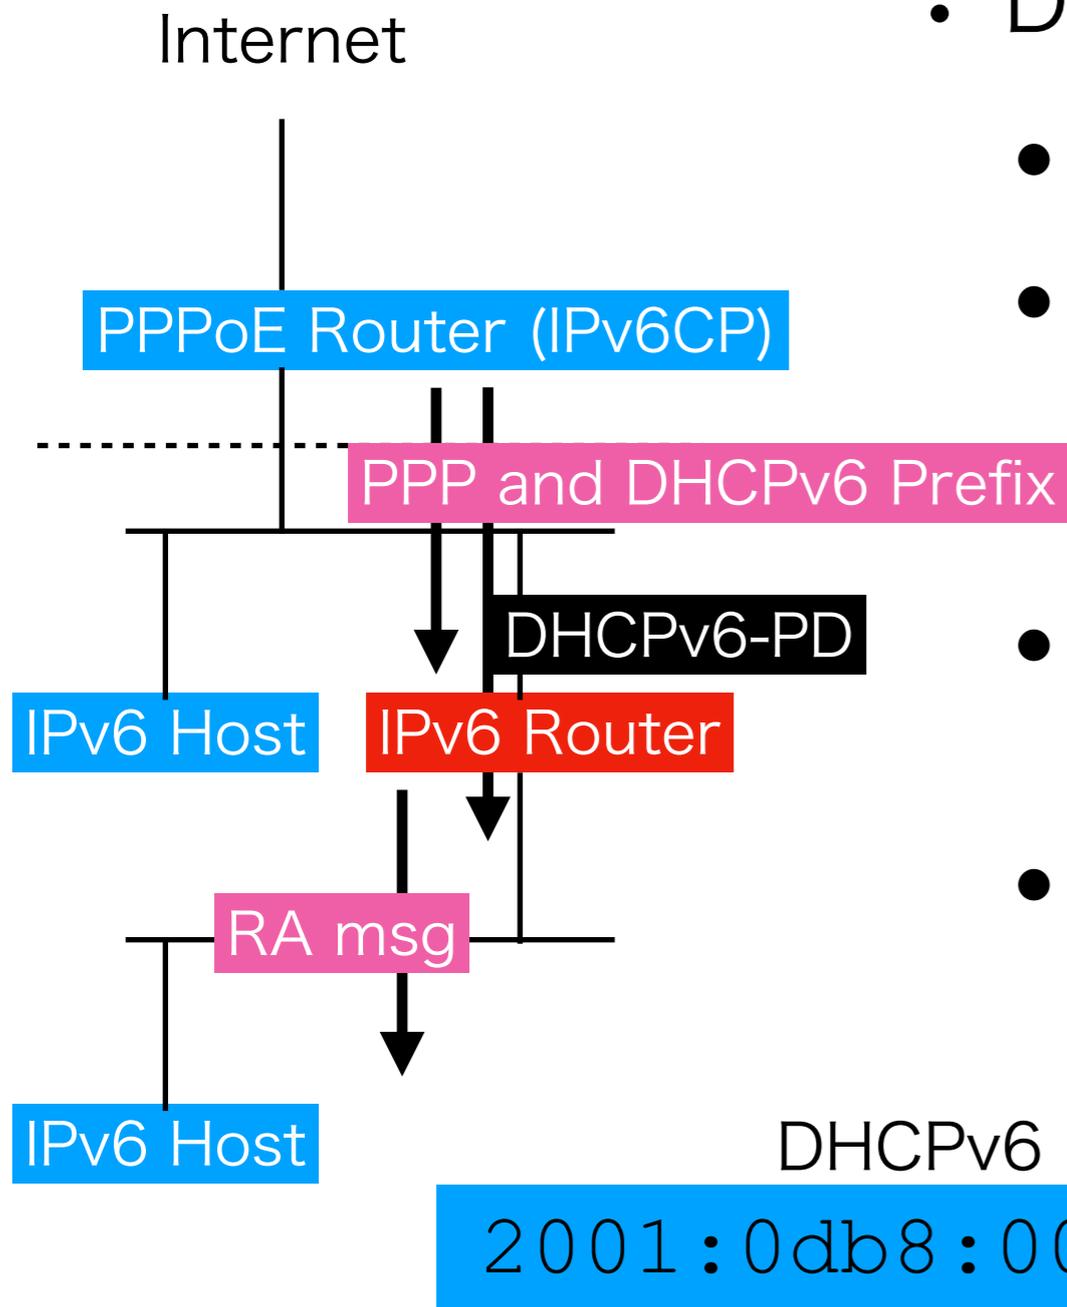
- B: SLAAC (native)
 - `ifconfig_em0_ipv6="inet6 accept_rtadv"`
 - "accept_rtadv" enables receiving RA messages from the router
 - Prefix (network address of IPv6) and default router are configured **by kernel**
 - DNS server information can be configured **by rtsold(8)** if RA messages have the option
 - `rtsold_enable="YES"`
 - `rtsold_flags="em0"`

RA msg (prefix opt.)

SLAAC

`2001:0db8:0000:0000:0001:0000:0000:4444`

4 Config Scenarios



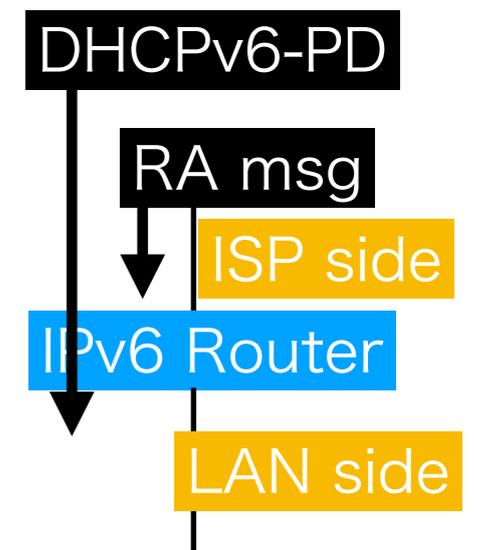
- D: PPPoE (IPv6CP) + DHCPv6
 - PPPoE works for IPv6: `net/mpd5`
 - IPv6CP (a part of PPPoE) provides ID part (lower 64-bit) of the WAN-facing address on the router
 - DHCPv6 option provides prefix information on the same link
 - DHCPv6-PD provides prefix information for your router

IPv6 on FreeBSD

- FreeBSD can support (and rc.d(8) was designed to support) the following:
 - A: Manual Configuration
 - B: SLAAC (native)
 - C: SLAAC + DHCPv6 (native)
 - D: PPPoE (IPv6CP) + DHCPv6
- You have to install a DHCPv6 client for C and D. DHCPv6-PD (for a router) is popular but DHCPv6 (for a host) is not, however.
- **Question:** do you have any other configuration scenario?

IPv6 on FreeBSD

- One pitfall is that you cannot have an RA-receiving interface (`accept_rtadv`) if packet forwarding is enabled. SLAAC is for host node only.
- A router facing your ISP may require if it uses SLAAC + DHCPv6-PD
- Setting `net.inet6.ip6.rfc6204w3=1` allows it.



DHCPv6 Client in Base?

- We want, but I **personally** want a very small implementation which covers only DHCPv6.
- DHCPv6 is not independent from the viewpoint of configuration; it should be invoked depending on RA messages. The client should work well with other utilities such as `rtsold(8)`.
- Candidates: `dhcpcd` or `wide-dhcp6`. A tight integration is preferred for good out-of-box experience

Issues and Missing Feature

- **Link-local address does not always work**
 - Utilities in the base system should be "address-family agnostic" whenever possible
- **Anycast address**
 - L3 communication is now allowed by a revised RFC.
- **Privacy extensions**
 - Stable address instead of EUI-64 IDs
- mDNS?
- **Question:** Please share your idea!

Do you know what happens, BTW?

```
% ping6 ff02::1%em0
% ping6 ff02::2%em0
% ping6 -w ff02::1%em0
```

/etc/resolv.conf - Does this work?

```
nameserver fe80::ffff:1:35%epair3b
```

/etc/ntp.conf - Does this work?

```
restrict -6 fe80::ffff:1:7b%vlan100 noquery nomodify
server -6 fe80::ffff:1:7b%vlan100 iburst
```

/etc/rc.conf - Does this work?

```
ipv6_defaultrouter="fe80::%em0"
syslogd_flags="-b [fe80::%bridge100]/64"
```

/etc/exports - Does this work?

```
/a -alldirs fe80::a6ba:dbff:fe11:2290%lagg0
/b -alldirs -network fe80::%lagg0/10
```

Issues and Missing Feature

- **Link-local address does not always work**
 - Utilities in the base system should be "address-family agnostic" whenever possible
- **Anycast address**
 - L3 communication is now allowed by a revised RFC.

You can safely add `fe80::/64` anycast to multiple routers for redundancy (like CARP or VRRP):

```
epair4b: flags=8843<UP,BROADCAST,RUNNING,SIMPLEX,MULTICAST> metric 0 mtu
1500
options=8<VLAN_MTU>
inet6 fe80::4f:6cff:feb5:d30b%epair4b prefixlen 64 scopeid 0x4
inet6 fe80::%epair4b prefixlen 64 anycast scopeid 0x4
```

Issues and Missing Feature

- **Link-local address does not always work**
 - Utilities in the base system should be "address-family agnostic" whenever possible
- **Anycast address**
 - L3 communication is now allowed by a revised RFC.
- **Privacy extensions**
 - Stable address instead of EUI-64 IDs

Privacy issue: the MAC address is used to generate the IPv6 address:

```
ether 02:4f:6c:b5:d3:0b  
inet6 fe80::4f:6cff:feb5:d30b%epair4b prefixlen 64 scopeid 0x4
```

Issues and Missing Feature

- **Link-local address does not always work**
 - Utilities in the base system should be "address-family agnostic" whenever possible
- **Anycast address**
 - L3 communication is now allowed by a revised RFC.
- **Privacy extensions**
 - Stable address instead of EUI-64 IDs
- mDNS?
- **Question:** Please share your idea!

- Your bad experience is still valuable to improve the future releases of FreeBSD. Please share!

Questions?