

Do you remember legacy IP^(*)?

What happened to FreeBSD's IPv6?

Bjoern Zeeb, BSDCan 2019

Today's Menu

Hungry for more bits? Then join for the bigger addresses and crunchy colons and learn what that 10 year old badge said.

- ❖ **Starters**
Evolution of the IPv6 stack
- ❖ **Main**
*Advocacy,
knobs and features,
and for the developers*
- ❖ **Dessert**
the menu for next time

<picture removed, URL below>

Timeline

- 1999-11 KAME, first svn log entry for sys/netinet6

r52904 | shin | 1999-11-05 14:41:39

KAME related header files additions and merges.
(only those which don't affect c source files so much)

Reviewed by: cvs-committers
Obtained from: KAME project

Timeline

- 1999 KAME

r53541 | shin | 1999-11-22 02:45:11

KAME netinet6 basic part(no IPsec,no V6 Multicast Forwarding, no UDP/TCP for IPv6 yet)

With this patch, you can assigne IPv6 addr automatically, and can reply to IPv6 ping.

Reviewed by: freebsd-arch, cvs-committers

Obtained from: KAME project

Timeline

- 1999 KAME

r54263 | shin | 1999-12-07 17:39:16

udp IPv6 support, IPv6/IPv4 tunneling support in kernel,
packet divert at kernel for IPv6/IPv4 translator daemon

This includes queue related patch submitted by
jburkhol@home.com.

Submitted by: queue related patch from jburkhol@home.com

Reviewed by: freebsd-arch, cvs-committers

Obtained from: KAME project

Timeline

- 1999 KAME

r55009 | shin | 1999-12-22 19:13:38

IPSEC support in the kernel.

pr_input() routines prototype is also changed to support IPSEC and IPV6 chained protocol headers.

Reviewed by: freebsd-arch, cvs-committers

Obtained from: KAME project

Timeline

- 1999 KAME

r55163 | shin | 1999-12-28 02:37:14

Getaddrinfo(), getnameinfo(), and etc support in libc/
net.

Several udp and raw apps IPv6 support.

Reviewed by: freebsd-arch, cvs-committers

Obtained from: KAME project

Timeline

- 2000 KAME

r55679 | shin | 2000-01-09 19:17:30

tcp updates to support IPv6.

also a small patch to `sys/nfs/nfs_socket.c`, as `max_hdr` size change.

Reviewed by: `freebsd-arch`, `cvsc-committers`

Obtained from: KAME project

Timeline

- 2000 KAME

r56722 | shin | 2000-01-28 05:10:56

IPv6 multicast routing.

kernel IPv6 multicast routing support.

pim6 dense mode daemon

pim6 sparse mode daemon

netstat support of IPv6 multicast routing statistics

Merging to the current and testing with other existing multicast routers is done by Tatsuya Jinmei <jinmei@kame.net>, who writes and maintains the base code in KAME distribution.

Make world check and kernel build check was also successful.

Timeline

- 2000 KAME

r56815 | shin | 2000-01-29 13:54:44

Add ip6fw.

Yes it is almost code freeze, but as the result of many thought, now I think this should be added before 4.0...

make world check, kernel build check is done.

Reviewed by: green

Obtained from: KAME project

Timeline

It's been 3 months,
split up into some junks,
and IPv6 and IPsec
was in FreeBSD.

FreeBSD 4 (March 2000) to ship this!

Timeline

- 2000 later ..

We start seeing commits in `sys/netinet6` for all kinds of things and from various committers doing code cleanup and fixes.

- 2001-06 sync to kame-20010528-freebsd43-snap.tgz
- 2002-04 cleanup for freebsd4-snap-20020128
- 2002-07 default of IPv4-mapped IPv6 address to off
- ~~2002-10 "Fast IPsec" (not with v6 yet)~~
- 2003-03 mention of TAHL in the commit log (testing!)
- 2003-09 Enable IPv6 for Token Ring

Timeline

- 2003-10 Advanced Sockets API for IPv6 to RFC3542
- 2003-11 _____
- 2005-04 Add IPv6 support to IPFW and Dummmynet.
- 2005-07 SA (IPsec)
- 2005-10 sync with KAME
- 2005-11 Conclusion of KAME announced for 2006-03
<https://lists.freebsd.org/pipermail/freebsd-net/2005-November/008968.html>
- 2006-03 (final) KAME merges
- 2006-04 More work on harmonising v4 and v6 bits
- 2006-05 _____
- 2007-04 Routing Header [RH0] processing off
- 2007-07 Fast IPsec IPv6 support, KAME IPsec removed
- ~~2008-05 Multi FIB IPv4~~
- 2008-05 VIMAGE gets mentioned

Timeline

- 2008-07 Giant starts to go away
- 2008-08 1st VIMAGE commit
- 2008-09 More SAs
- 2008-11 Multi-IPv[46] jails committed
- 2008-12 L2/L3 tables split (ARP, NDP) vs routing
- 2009-03 Start removing RH0 code (RFC5095)
- 2009-04 IPv6 SSM and MLDv2 mega-commit
- 2009-05 _____
- 2009-06 SA

SeND

- 2010-08 Secure Neighbor Discovery (SeND), RFC 3971
- NDP open to attack (RFC3756)
- SeND provides:
 - NDP option for public-key based signature
 - Cryptographically Signed Addresses to prove address ownership of nodes
 - X.509 certs authorising nodes as routers to delegate prefixes
- GSoC project
<https://wiki.freebsd.org/SOC2009AnaKukec>
- Kernel hooks, user space handling
- AsiaBSDCon paper & talk

Up-to World IPv6 Day

- 2011-04 compile without INET
- 2011-06-06 RDNSS & DNSSL support
- 2011-06-06 accept_rtadv, no_rdr, ifdisabled, ..
- Created FreeBSD and PC-BSD “IPv6-only” snapshots
- Created web pages
- Had an IRC channel #freebsd-w6d
- Did not sleep enough ;-)
- Foundation & iX press release
<http://www.prweb.com/releases/2011/6/prweb8529718.htm>

- » [Developers](#)
- » [Developer's Handbook](#)
- » [Porter's Handbook](#)
- » [Source code repositories](#)
- » [Release Engineering](#)
- » [Platforms](#)
- » [Project Ideas](#)
- » [Contributing](#)

IPv6 in FreeBSD

Category Navigation

- [IPv6 in FreeBSD](#)
- [IPv6-only](#)
- [World IPv6 Day](#)
 - [www stats](#)

Introduction

FreeBSD has shipped tightly integrated IPv6 support in its first release to include "out-of-the-box" IPv6 support. development in the FreeBSD community, including pa

Latest news

- **August 3, 2011:** 9.0-BETA1 IPv6-only snapshot see [IPv6Only wiki page](#).
- **June 9, 2011:** Thanks to everyone having joined (unspectacular) [statistics of www.freebsd.org](#) after the event for you.
- **June 7, 2011:** New set of IPv6-only snapshot nameserver support in [rtsol\(8\)](#) and [rtsold\(8\)](#), see [here](#) and find download links and netinstall doc
- **June 6, 2011:** The FreeBSD Foundation and I support the efforts of World IPv6 Day to accelerate [release](#) and find more information on FreeBSD
- **June 6, 2011:** New set of IPv6-only snapshot and find download links [here](#).
- **June 6, 2011:** Hiroki Sato [commits RFC 6106](#) Configuration" (RDNS and DNSSEC) support.

IPv6 in FreeBSD

FreeBSD is a widely used, open source operating system for decades of research, as well as a reference implementation. FreeBSD first shipped IPv6 support in March 2000 as

IPv6 and the FreeBSD Project

The FreeBSD Project has been an early adopter and with the help of the community, we have been serving releases. FreeBSD's website, mailing lists, and developer infras

FreeBSD is used by critical Internet infrastructure such as the world's busiest and most reliable web sites as well as for the best IPv6 support. To read more about some of their products, see the [FreeBSD Foundation Testimon](#)

- » [Developers](#)
- » [Developer's Handbook](#)
- » [Porter's Handbook](#)
- » [Source code repositories](#)
- » [Release Engineering](#)
- » [Platforms](#)
- » [Project Ideas](#)
- » [Contributing](#)

FreeBSD and World IPv6 Day

Category Navigation

- [IPv6 in FreeBSD](#)
- [IPv6-only](#)
- [World IPv6 Day](#)
 - [www stats](#)

About World IPv6 Day

While Internet Protocol version 6 is in fact more than 15 years old, World IPv6 Day, a 24 hour test flight day on **8 June 2011** has motivated a lot of different organizations to get ready for IPv6, or improve their already existent IPv6 support. Major content providers will enable IPv6 for their websites that day, network operators and hosting companies have been working to provide IPv6 to their customers and operating system vendors like FreeBSD have been improving IPv6 support. You can find more information about World IPv6 Day on The Internet Society's web page [www.worldipv6day.org](#).

You and FreeBSD on World IPv6 Day

What is special about that day for FreeBSD you might wonder? Indeed, with the help of the community, FreeBSD has been serving releases on IPv6 since 2003. We have our major infrastructure like [www.freebsd.org](#) IPv6 enabled since [2007](#). FreeBSD itself has been supporting IPv6 since the 4.0 Release for over a decade now, using the [KAME](#) based reference implementation.

Nonetheless there are things we can do during that day:

- help promoting IPv6. During the day there might be a lot more IPv6 users suddenly, and some people are expecting significantly more traffic. A lot of FreeBSD servers are out there actively using IPv6, some people have put FreeBSD into their networking products, but only few tell us. We would love to hear about your FreeBSD IPv6 experiences on World IPv6 Day.
- help you in case of problems. Obviously if you will run into problems you might want help immediately for the one day test flight. We will more closely monitor the [FreeBSD networking mailing list](#), as well as being around on IRC in #freebsd-w6d on EFNet to answer your questions. In case you will file an IPv6 related [bug report](#), you may want to tag it with [IPv6].
- working on improving your IPv6 experience. As we get your questions and feedback we will take notes to further improve IPv6 support in FreeBSD. There might not be an immediate change but rest assured that your feedback will not be lost.

Join us for World IPv6 Day, spread the word, to help improving IPv6 support in FreeBSD even further!

FAQ

What is this IPv6 thing?

Please see the [chapter on IPv6](#) in the FreeBSD [Handbook](#) for an explanation.


Do you provide IPv6 connectivity?

FreeBSD is an operating system, not an Internet Service Provider. There are multiple ways to connect to an IPv6 network with FreeBSD however:



Last updated: 2011-09-12 00:00 UTC

World IPv6 Day



No-INET kernel&world

Why?

Maybe should have called it
the NO-NO-INET6 option?

Timeline

- 2011-08 ipfw fwd IPv6 support
- 2011-09 ipv6_cpe_wanif
- 2012-02 Multi-FIB IPv6 support
- 2012-05 IPv6 offloading support (checksum, TSO, LRO)
- 2012-12 more locking work, hash tables, ..
- 2013-07 pcpu statistics framework coming
- 2014-04 UDPLite
- 2014-05 v4-in-v6 and v6-in-v4 IPsec tunnels
- 2014 RSS work
- 2014 a lot of code cleanup, address handling, scope
- 2014-11 gre over v6
- 2014-11 remove faith(4)/faithd(8)
- 2014 IPsec code improvements

Timeline

- 2014/15 A lot of nd6/dad/... code improvements
- 2015-04 RA DoS mitigation SA
- 2015-08 merges from projects/routing
- 2015-08 RFC 7527 (enhanced dad)
- 2015-08 ongoing RSS work
- 2015 Lots of nd6_* work
- 2015/16 Lots more routing KPI work
- 2016-02 ND default router list locking
- 2016 route cache comes back
- 2016 A lot of code quality time
- 2016 VIMAGE teardown cleanup
- 2016-07 ipfw_nptv6
- 2016-08 ipfw_nat64 (stateful, stateless)

Timeline

- 2016-10 Lock the ND prefix list + refcounting
- 2017-02 projects/ipsec merge
- 2017-05 Fix connect(2) in some cases
- 2017-07 RFC8200
- 2017-12 RFC6980 ignore IPv6 NDP with frag header
- 2018 UNH test framework fixes
- 2018 epoch(9) happened
- 2018-06 IP encapsulation code rework
- 2018-07 Allow implicit TCP connection setup for TCP/IPv6
- 2018-08 IPv6 reassembly SAs & related improvements
- 2018-08 UDPv6 locking parity to v4 (from 2012)
- 2019-01 Fix lo traffic on non-lo0 link-local IPv6 addresses
- 2019-03 NAT64 CLAT (v4 to v6) support/LSN update

There was a saying:

“The Internet runs on Proposed Standards.”

(Bob Hinden)

RFC8200

Timeline

- 2016-10 Lock the ND prefix list + refcounting
- 2017-02 projects/ipsec merge
- 2017-05 Fix connect(2) in some cases
- 2017-07 RFC8200
- 2017-12 RFC6980 ignore IPv6 NDP with frag header
- 2018 UNH test framework fixes
- 2018 epoch(9) happened
- 2018-06 IP encapsulation code rework
- 2018-07 Allow implicit TCP connection setup for TCP/IPv6
- 2018-08 IPv6 reassembly SAs & related improvements
- 2018-08 UDPv6 locking parity to v4 (from 2012)
- 2019-01 Fix lo traffic on non-lo0 link-local IPv6 addresses
- 2019-03 NAT64 CLAT (v4 to v6) support/LSN update

Summary

- About 1850 commits in `sys/netinet6`.
- Ignored most outside `sys/netinet6` and userspace.
- KAME
- IPv4-IPv6 feature parity
- It's become business as usual (mostly)



Timeline

- 1999-11 KAME, first svn log entry for sys/netinet6

r52904 | shin | 1999-11-05 14:41:39

KAME related header file additions and merges.

(only those which don't affect cscope files so much)

Reviewed by: cvs-committers

Obtained from: AM project

**APPROACHING
20 YEARS
THIS YEAR**

Quote from IRC

2019-02-13 phk:

battlez, you realize that ipv6 is old enough to be a "vintage-protocol" now ? :-)

IPv6 and you?



How To

rc.conf (static IPv6 addresses):

```
ifconfig_ed0_ipv6="inet6 2001:db8:1::1  
prefixlen 64"  
ifconfig_ed0_alias0="inet6 2001:db8:2::1  
prefixlen 64"
```


How To

rc.conf (static IPv6 addresses):

```
ifconfig_ed0_ipv6="inet6 2001:db8:1::1/64"
```

```
ipv6_defaultrouter="2001:db8:1::fffe"
```

```
ipv6_defaultrouter="fe80::1%ed0"
```

```
ipv6_default_interface="ed0"
```

```
ipv6_defaultrouter="fe80::1%${ipv6_default_interface}"
```

The simplest thing

rc.conf (for IPv6 auto-configuration):

```
ifconfig_<IF>_ipv6="up"
```

```
ipv6_activate_all_interfaces="YES"
```

rtsol

rc.conf (for IPv6 auto-configuration):

```
rtsol_flags=""
```

```
rtsold_enable="NO"
```

```
rtsold_flags="-a"
```

We will run rtsol once automatically for you on “up”.
For continuous operation, enable rtsold.

Address selection

rc.conf:

```
ip6addrctl_enable="YES"  
ip6addrctl_policy="AUTO"  
# (ipv4_prefer, ipv6_prefer, or AUTO)
```

If you configure IPv6, we'll prefer IPv6
so defaults should be fine.

[If you have a special setup, you can tune the table.]

Privacy

```
ipv6_privacy="YES"
```

Will use RFC 4941 privacy addresses on interfaces where we receive RAs with prefix.

v4-mapped v6

```
ipv6_ipv4mapping="NO"
```

```
"::ffff:a.b.c.d"
```

Forwarding

very much like legacy IP:

```
ipv6_gateway_enable="YES"
```

```
ipv6_static_routes="ula"
```

```
ipv6_route_ula="fc00::/7 ::1"
```

Tuning your interface

ifconfig(8) section:

The following parameters are for ICMPv6 Neighbor Discovery Protocol.

Note that the address family keyword "inet6" is needed for them:

```
accept_rtadv no_radr auto_linklocal  
defaultif ifdisabled ipv6_only  
nud no_prefer_iface no_dad
```

Note: normally you should not need to!

CPE mode

```
ipv6_cpe_wanif="ed0"
```

Will set: global sysctl defaults

```
net.inet6.ip6.no_radr=1
```

```
net.inet6.ip6.rfc6204w3=1
```

for you and

```
-no_radr_accept_rtadv
```

on the CPE interface.

Advertising yourself

```
rtadvd_enable="YES"  
rtadvd_interfaces=""
```

Configuration file in termcap(5) style
/etc/rtadvd.conf

Advertising yourself

```
default:\  
:raflags#8:
```

```
ed0:\  
:addr="2001:db8:1::":prefixlen#64:\  
:rdnss="2001:db8:1::10,2001:db8:f::2:43":\  
:dnssl="example.com":
```

Tools

```
netstat [-r] [-f inet6] [-W]
```

```
sockstat -6
```

```
route -6 get 2001:db8::1
```

```
...
```

Tools

netstat [-r] [-f inet6] [-W]

sockstat -6

route -6 get 2001:db8::1

...

+ ndp (kind-of like arp)

ndp -a

ndp -r

ndp -p

Firewall



Firewall

- We have 3 firewalls.
- All support IPv6.
- I have used 3 of these 4 over time.
- End-to-end: central vs. host-FW or both
- Feature parity-ish

Firewall

- Filtering. We provide samples.
- NAT64: ipfw can do stateless and stateful
- CLAT: ipfw
- Transition technologies which help to remove IPv4 (from the transport path).

Some golden rules

1. Do not think IPv4
2. Do not think IPv4
3. Go to 1, back again? fall-through
4. Start somewhere to roll it out
5. Monitor!
6. Dual-Stack is a transition technology
7. Go IPv6-only where you can

It is dying

<BSD is dying, Slashdot screenshot removed, URL below>

<https://bsd.slashdot.org/comments.pl?sid=189013&cid=15569908>

It is dying

~~<BSD is dying, Slashdot screenshot removed, URL below>~~

<https://bsd.slashdot.org/comments.pl?sid=189013&cid=15569900>

***They say (said): BSD is dying.
I say: IPv4 will be dying faster!***

Developer

1. Do not think IPv4
2. Do not think IPv4
3. Go to 1, back again? fall-through
4. Prepare your code for IPv6-only
5. Make dual-stack work
6. Do not rely on one protocol family or the other
7. `gethostby*()` is a no-go

On the menu for next time

- IETF: IPv6-only flag?
- IETF: Jumbogram removal?
- DHCPv6 in base
- IPv6-only FreeBSD tree
- Working backwards from NFS mount to UEFI boot
- Renumbering

2009-05

BSDCan 2009
The Technical BSD Conference



Bjoern A. Zeeb
bz@FreeBSD.org

Do you remember legacy IP?



Thank you

Questions?

Bjoern Zeeb

bz@FreeBSD.org / info@zabbadoz.com