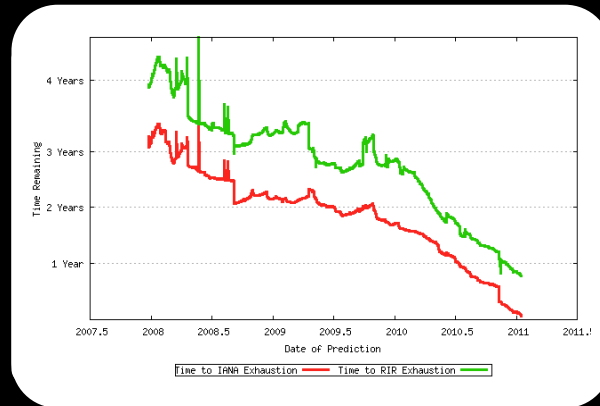


Looking Towards Tomorrow: Implementing IPv6 Now

Gerard Hickey
January 27, 2011
Macworld 2011 Conference
IT813

How long will IPv4 last?

- IPv4 exhaustion in 21 days on Feb. 18th



Why move to IPv6?

- It is not just because of IP exhaustion
- IPv6 can be easily extended to support emerging technologies
- Enhanced security features
- Flow labeling supports resource allocations

IPv6 Features

- Much larger address space
- Link-local addressing
- Protocol extensions
- End to end security
 - Packet encryption
 - Endpoint authentication

IPv6 Overview

- 128 bit addresses (340 trillion³)
- Headers are a series of link lists
 - Header does not include a checksum
 - Most headers are only touched by destination

IPv6 Overview

- No more broadcast address
 - More use of multicast addressing
 - ARP replaced with Neighbor Solicitation
- Only source host can fragment packet
- IPSec is a requirement

IPv6 Overview

- Host finds router through Router Solicitation and Router Advertisements
- New DNS record for IPv6 addressing

Increase in Addresses

- 340,282,366,920,938,463,463,374,607,431,768,211,456
- 340 Undecillion
- 282 Decillion
- 366 Nonillion
- 920 Octillion
- 938 Septillion
- 436 Sextillion
- 436 Quintillion
- 374 Quadrillion
- 607 Trillion
- 431 Billion
- 768 Million
- 211 Thousand
- 456 Hundreds

Approximately 50
octillion addresses for
each of the 6.5 billion
inhabitants on the earth

6.65×10^{23} addresses / m²

IPv6 Addresses

- Addresses consist of eight 16 bit hexadecimal values separated by colons
- Leading zeros can be omitted
- Consecutive zeros can be replaced with a double colon (::), but only once

IPv6 Addresses

- All of these addresses are equivalent

2001:0470:000B:01F7:0000:0000:0000:0013

2001:470:B:1F7:0:0:0:13

2001:470:B:1F7::13

Special Addresses

Purpose	Long Address	Short Address
Loopback	0:0:0:0:0:0:0:1	::1
Unspecified	0:0:0:0:0:0:0:0	::
IPv4 Mapped	0:0:0:0:FFFF:###.###.###	::FFFF:###.###.###

- IPv4 mapped addresses for support of IPv4 only devices

What is radvd?

- Router Advertisement Daemon
- Listens for Router Solicitations (RS) and responds with Router Advertisements (RA)
- Provides auto configuration data to hosts
- Not installed on OS X Server

Privacy Addressing

- Auto configured IPv6 addresses use the host's MAC address
- Security and tracking considerations
- Privacy Addressing will generate random host address periodically

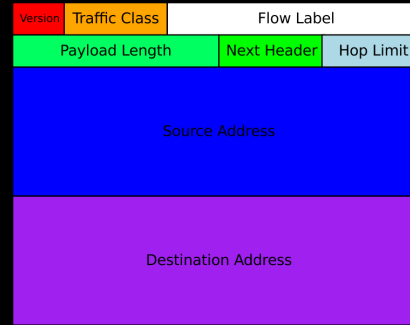
```
sudo sysctl -w net.ipv6.ip6.use_tempaddr=1
```

Dual Stacks

- A dual stack is when an IPv4 and IPv6 address are on the same IP stack (i.e. same network interface)
- Nearly every operating system now supports dual stacks

IPv6 Header

- Reduced amount of fields need to be examined
- Only header use to route packet around the network
- Other header types are provided as extensions



How to Prepare

- Need to develop plans now!
- Start with getting an IPv6 tunnel
- Insure your routers support IPv6
 - 6to4, 6over4, 6in4 or Teredo tunnels
 - Routing protocol 41 for tunnels
 - radvd

IPv6 Tunnels

- Hurricane Electric
- SixXS
- Freenet6
- Tunnelbroker.ru

OS X IPv6 Commands

- ip6
- ping6
- traceroute6
- ifconfig
- dig
- netstat

ping6

```
hertz% ping6 www.kame.net
PING6(56=40+8+8 bytes) 2001:470:b:1f7::33 --> 2001:200:dff:fff1:216:3eff:feb1:44d7
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=0 hlim=53 time=174.150 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=1 hlim=53 time=168.068 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=2 hlim=53 time=168.991 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=3 hlim=53 time=167.674 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=4 hlim=53 time=166.538 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=5 hlim=53 time=166.780 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=6 hlim=53 time=166.085 ms
16 bytes from 2001:200:dff:fff1:216:3eff:feb1:44d7, icmp_seq=7 hlim=53 time=167.555 ms
```

traceroute

```
hertz% traceroute www.kame.net
traceroute to orange.kame.net (203.178.141.194), 64 hops max, 52 byte packets
 1 fw (172.99.1.1) 1.475 ms 0.685 ms 0.539 ms
 2 73.105.120.1 (73.105.120.1) 9.111 ms 8.609 ms 9.615 ms
 3 ge-3-18-ur02.bellevue.wa.seattle.comcast.net (68.86.113.25) 10.786 ms 9.057 ms 9.039 ms
 4 be-50-ar01.seattle.wa.seattle.comcast.net (68.85.240.65) 21.644 ms 12.677 ms 13.126 ms
 5 pos-1-13-0-0-cr01.seattle.wa.ibone.comcast.net (68.86.95.21) 23.597 ms
   pos-1-10-0-0-cr01.seattle.wa.ibone.comcast.net (68.86.93.97) 10.190 ms 11.070 ms
 6 pos-0-0-0-0-pe01.seattle.wa.ibone.comcast.net (68.86.86.138) 12.827 ms 11.922 ms 18.370 ms
 7 if-4-0-0.core2.s00-seattle.as6453.net (66.110.64.1) 11.476 ms 9.719 ms 10.532 ms
 8 if-15-0-0.mcore4.pdi-paloalto.as6453.net (207.45.196.65) 37.907 ms 39.969 ms 39.563 ms
 9 if-0-0-0.core3.sqn-sanjose.as6453.net (216.6.33.5) 1136.041 ms 60.521 ms 75.933 ms
10 vlan56.icore1.sqn-sanjose.as6453.net (209.58.116.49) 32.915 ms 33.372 ms 32.019 ms
11 vlan504.icore1.sqn-sanjose.as6453.net (209.58.116.22) 38.561 ms 38.864 ms 37.629 ms
12 ae-7.r21.snjsca04.us.bb.gin.ntt.net (129.250.5.54) 39.393 ms 40.164 ms 39.866 ms
13 as-0.r21.tokyjp01.jp.bb.gin.ntt.net (129.250.5.82) 156.501 ms
   as-2.r20.tokyjp01.jp.bb.gin.ntt.net (129.250.2.35) 156.146 ms 149.750 ms
14 * * *
15 203.105.72.18 (203.105.72.18) 149.532 ms 144.915 ms 134.299 ms
16 ve-51.foundry6.otemachi.wide.ad.jp (203.178.141.141) 136.826 ms 141.759 ms 136.474 ms
17 ve-42.foundry4.nezu.wide.ad.jp (203.178.136.66) 142.190 ms 145.992 ms 144.905 ms
18 cloud-net1.wide.ad.jp (203.178.138.9) 154.134 ms 136.395 ms 138.605 ms
19 * * *
```

traceroute6

```
hertz% traceroute6 www.kame.net
traceroute6 to orange.kame.net (2001:200:dff:fff1:216:3eff:feb1:44d7) from 2001:470:b:
1f7::33, 64 hops max, 12 byte packets
 1 2001:470:b:1f7::1 0.734 ms 0.675 ms 0.688 ms
 2 kineticcompute-1.tunnel.tserv14.seal.ipv6.he.net 33.795 ms 34.882 ms 34.679ms
 3 gige-g2-6.core1.seal.he.net 33.781 ms 35.035 ms 35.442 ms
 4 10gigabitethernet1-2.core1.pdx1.he.net 37.892 ms 41.836 ms 39.800 ms
 5 10gigabitethernet7-1.core1.sjc2.he.net 57.951 ms 59.843 ms 59.156 ms
 6 xe-0.equinix.snjsca04.us.bb.gin.ntt.net 53.456 ms 53.574 ms 56.159 ms
 7 as-1.r21.osakjp01.jp.bb.gin.ntt.net 227.605 ms 227.568 ms 227.347 ms
 8 ae-4.r21.tokyjp01.jp.bb.gin.ntt.net 225.314 ms 226.047 ms 278.376 ms
 9 po-2.a15.tokyjp01.jp.ra.gin.ntt.net 227.630 ms 228.041 ms 227.832 ms
10 ge-8-2.a15.tokyjp01.jp.ra.gin.ntt.net 168.110 ms 169.261 ms 167.753 ms
11 ve44.foundry6.otemachi.wide.ad.jp 167.676 ms 167.854 ms 167.689 ms
12 ve42.foundry4.nezu.wide.ad.jp 168.634 ms 169.705 ms 168.122 ms
13 cloud-net1.wide.ad.jp 168.968 ms 168.834 ms 173.126 ms
14 * * *
```

ifconfig

- OS X provides dual stack network interfaces

```
hertz% ifconfig
lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
    inet6 ::1 prefixlen 128
    inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
    inet 127.0.0.1 netmask 0xff000000
    inet6 fd02:8d48:1afb:4a39:3615:9eff:fe2e:b13a prefixlen 128
en0: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
    ether 34:15:9e:2e:b1:3a
    inet6 fe80::3615:9eff:fe2e:b13a%en0 prefixlen 64 scopeid 0x4
    inet6 2001:470:b:1f7::33 prefixlen 48
    inet 172.99.1.33 netmask 0xfffff00 broadcast 172.99.1.255
    media: autoselect (100baseTX <full-duplex,flow-control>)
    status: active
```

dig

- dig by default retrieves IPv4 A records

```
hertz% dig www.kame.net

; <<>> DiG 9.6.0-APPLE-P2 <<>> www.kame.net
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 61981
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 13,
ADDITIONAL: 0

;; QUESTION SECTION:
;www.kame.net.          IN A

;; ANSWER SECTION:
www.kame.net.         86292 IN CNAME orange.kame.net.
orange.kame.net.     86292 IN A   203.178.141.194
```

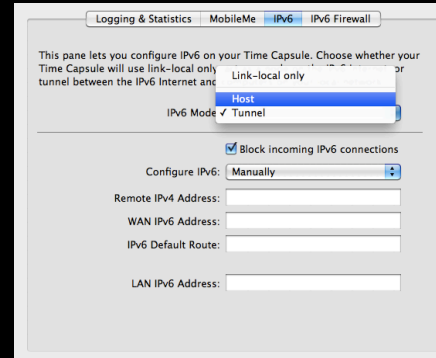

netstat

```
hertz% netstat -rn -f inet6
Routing tables
```

```
Internet6:
Destination          Gateway              Flags      Netif  Expire
default              2001:470:b:1f7::1   UGSc      en0
::1                  ::1                 UH        lo0
2001:470:b:1f7::1    8:0:20:bc:6d:1     UHL       en0
2001:470:b:1f7::33  34:15:9e:2e:b1:3a  UHL       lo0
2002::/16            2002:0:0:1::1      Uc        stf0
2002:0:0:1::1       link#3              UHL       lo0
fd02:8d48:1afb:4a39:3615:9eff:fe2e:b13a link#1              UHL       lo0
fe80::%lo0/64        fe80::1%lo0         Uc        lo0
fe80::1%lo0          link#1              UHL       lo0
fe80::%en0/64        link#4              Uc        en0
fe80::224:36ff:feal:8113%en0 0:24:36:a1:81:13   UHLW      en0
fe80::a00:20ff:febc:6d01%en0 8:0:20:bc:6d:1     UHLW      en0
fe80::3615:9eff:fe2e:b13a%en0 34:15:9e:2e:b1:3a  UHL       lo0
ff01::/32            ::1                 Um        lo0
ff02::/32            ::1                 UmC       lo0
ff02::/32            link#4              UmC       en0
ff02::fb             link#4              UHmLW    en0
```

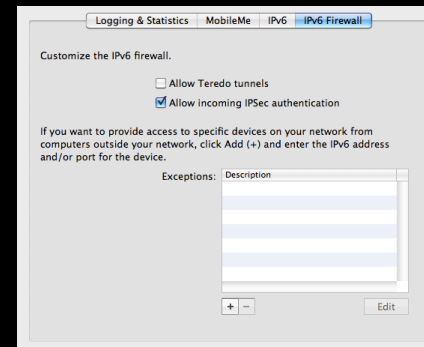
Airport Extreme

- v7.5 or better
- 6to4 tunnel
- radvd enabled
- No DHCPv6
- Cheap IPv6 gateway device



Airport IPv6 Firewall

- Only enabled if blocking incoming IPv6 connections
- Simple filter on inbound connections



IPv6 DNS Considerations

- Flush DNS cache - otherwise may be stuck with IPv4 addresses
- Use IPv6 DNS server
 - Google has 8.8.8.8 available

Reference Sites

- <http://www.sabi.co.uk/Notes/swIPv6Prefixes.html>
- <http://test-ipv6.com/>