

Getting started with IPv6 on Linux

Jake Edge
LWN.net
jake@lwn.net

LinuxCon North America
19 August 2011

History and Motivation

- IPng project – July 1994
- IPv6 - RFC 2460 – December 1998
- “IPv5” - Internet Stream Protocol used version 5 in packet header – v4 uses 4, v6 uses 6
- Recognized IPv4 address exhaustion, but also wanted to solve other network issues

Advantages of IPv6

- 2^{128} addresses (plenty)
- Easier routing
- QoS support using Flow Label
- IPsec required
- Mobility support
- Minimum MTU increase (1280 vs. 576)

Addresses

- 128 bits (4x IPv4 32 bit addresses)
- Represented as 8 groups of 4 hex digits:
2001:0DB8:AC10:FE01:0000:0000:0000:0001
Normally written: 2001:DB8:AC10:FE01::1
(leading zeros, consecutive zeros eliminated)
- 64-bit network prefix, 64-bit host address
- Usually 48-bit routing prefix, 16-bit subnet
- 2001:470:bac3::/48 represents a routed network

Special Addresses

- Loopback ::1 (127.0.0.1 for IPv4)
- Unspecified :: (0.0.0.0 for IPv4 INADDR_ANY)
- Link local FE80:: - host address based on MAC

MAC: 00:1D:BA:06:37:64 becomes


FE80::021D:BAFF:FE06:3764

(FFFE inserted in middle and bit 1 in first byte turned on – Modified EUI-64)

- Privacy concerns

Ipv6 Packet Format

- Much simpler than Ipv4
- Fixed 40-byte length (IPv4 20-60 bytes)
- Moved options into additional headers



(Graphic shamelessly poached from Wikipedia: <http://en.wikipedia.org/wiki/IPv6>)

Ipv6 Headers

- Next Header is type of any following header
 - 6 for TCP, 17 for UDP, 59 for no next header
 - Other options (fragmentation, routing, ...)
- No header checksum (eliminates recalculation), uses lower (CRC on ethernet) and higher (checksum for UDP and TCP headers) layers
- Flags (SYN, ACK, etc. in TCP header)
- Network layer devices only need to see the mandatory header information

Stateless Autoconfiguration (SAC)

- One way to get an initial address
- Host sends router solicitation using link local address (FE80::modified-EUI) to FF02::2
- Routers reply with router advertisement to FF02::1
- Advertisements contain network prefix information and router lifetime
- Multiple routers may reply with different subnets
- Duplicate address detection is used

DHCPv6

- SAC doesn't provide DNS hostnames
- Either have to run IPv4 DNS, statically define, or use DHCPv6
- DHCPv6 can also assign addresses (or those can also be statically configured)
- “Managed” flag in router advertisement notes the presence of DHCPv6 server in subnet
- “Stateless” does DNS, while “stateful” does address assignment

Routing

- Many IPv6 hosts will have multiple addresses one for each router they can talk to (at least)
- IPv6 is supposed to easily enable renumbering networks by just changing network prefix
- Mobile routing is done by having a “home” address, and a “care-of” address, that is routed by the home agent via tunneling
 - When devices roam to a new network (e.g. WiFi to cell data) it informs the home agent of care-of addr
 - Keeps old IPv6 address to maintain connections

DNS

- Uses the same basic DNS structure as IPv4
- Instead of A records, uses AAAA records

```
$ dig aaaa ipv6.google.com
```

```
...
```

```
;; ANSWER SECTION:
```

```
ipv6.google.com.          604800  IN  CNAME
```

```
    ipv6.l.google.com.
```

```
ipv6.l.google.com.       300      IN  AAAA
```

```
    2001:4860:b006::68
```

- MX and CNAME records use hostnames

Firewall

- Many IPv4 hosts live behind NAT so they can't be connected to from the internet
- That is not true with IPv6, by default all hosts will be reachable from the internet
- Stateful firewalling will be required

Ipv6 Commands

- Some standard commands have a 6 added:
 - ping6:

```
$ ping6 ::1
```

```
$ ping6 -I eth0 ff02::1
```
 - traceroute6:

```
$ traceroute6 -i wlan0 fe80...
```
 - iptables6
 - ifconfig and ip used to configure IPv6

Applications

- Some still need changes to handle IPv6:
http://www.deepspace6.net/docs/ipv6_status_page_apps.html
- Some changes to user interfaces is required:
[https://\[2001:db8:85a3:8d3:1319:8a2e:370:7348\]:443/](https://[2001:db8:85a3:8d3:1319:8a2e:370:7348]:443/)
- Without routing set up, the interface is needed:
`ssh user@FE80::021D:BAFF:FE06:3764%eth0`
- Router uses `radvd` – configuration file:
`/etc/radvd.conf`
 - responds to solicitations with router advertisements

IPv6 and IPv4 coexistence

- Hard to predict when (or if) IPv6 completely replaces IPv4
- Currently the vast majority of the internet is IPv4-only
- IPv6 hosts can talk to IPv4 via tunneling
 - Encapsulate IPv6 packets inside IPv4 packet data
 - Tunnel endpoints pack/unpack IPv6 packets
 - Several tunnel types: 6to4, Teredo, 6in4, ...
 - Teredo can work behind NAT (IPv4)

Tunnel brokers

- Various free (beer) services exist
- Hurricane Electric - <http://tunnelbroker.net/>
- SixXS - <http://www.sixxs.net/>
- Freenet6 - <http://gogonet.gogo6.com/>
- Other regional brokers exist
- Set up router to send IPv6 traffic via the tunnel

World IPv6 Day

- Was held on June 8
- Tested the readiness of IPv6 world-wide
- Major websites, including the top four, offered content over IPv6
- Went pretty smoothly, no major issues reported
- IPv6 traffic increased – still a blip against IPv4

More Information

- Can test your readiness and what needs to be done locally and at ISP, etc.:
 - <http://test-ipv6.com/>
- Books – no real recent ones focused on Linux
 - *IPv6 in Practice* – Benedikt Stockebrand
 - Debian sarge (2.6.8 kernel – a bit outdated)
 - *Running IPv6* – Iljitsch van Beijnum
 - RH 9 and RHEL 4

Web Sites

- <http://ipv6.com> - lots of IPv6 information
- <http://tldp.org/HOWTO/Linux+IPv6-HOWTO/>
 - has some holes, blank topics, but lots of good info
- http://www.deepspace6.net/docs/ipv6_status_page_apps.html
 - status of application support for IPv6
- Two Linux.com IPv6 “crash course” articles
 - lots of good information in those
- Wikipedia
- Lots more