

Introduction to IPv6

by Chuck Anderson

Sr. Network Engineer, Worcester Polytechnic Institute

President, Worcester Linux Users' Group

cra@wpi.edu

Version 2, 2011/09/14

Why IPv6?

- Extended Address Space
 - 128-bits vs. IPv4's 32-bits
 - We really are running out of IPv4, no really we are this time for sure!
 - IANA allocated the last remaining IPv4 blocks to the 5 regional registries on February 3, 2011.
 - Regional registries will run out in the next year or so
- Auto-configuration (no server required)
- Elimination of NAT
- No fragmentation by routers, only end hosts
- Multiple Addresses per interface are the norm

IPv6 Addressing

- Full Notation

- 32 hexadecimal digits split into groups of 4 digits
- 8 groups of 4 digits separated by colons
 - 2001:0db8:0001:ffff:0000:0000:dead:beef

- Shorthand Notation

- Remove leading zeros from each group of 4
- Replace ONE set of consecutive zeros with ::
 - 2001:db8:1:ffff::dead:beef

- Perl Regular Expression

- \$aeron = qr/^\(((?=(?>.*?::)(?!.*:)))(::)?([0-9A-F]{1,4}::?){0,5}([0-9A-F]{1,4}:){6}(\2([0-9A-F]{1,4}(::?)\$)){0,2}((25[0-5]|(2[0-4]|1[0-9]|[1-9])?[0-9])\.\.(\$)){4}([0-9A-F]{1,4}:[0-9A-F]{1,4})(?<![^:]:)(?<!\.)\z/i;

IPv6 Addressing (cont'd)

- Network Prefix
 - Identifies the address range assigned to a specific site or network
 - n bits long, where $0 < n \leq 64$ (often $32 \leq n \leq 56$)
- Subnet ID
 - Identifies the link (subnet) within a site
 - $(64-n)$ bits long
- Host ID (Interface ID) - Fixed Size!
 - Identifies a specific host or interface on a host
 - 64 bits long, often in EUI-64 format based on MAC address

EUI-64 Addressing

- Extended Unique Identifier 64-bits
- RFC 4291 defines usage in IPv6
 - Take first 3 octets of MAC (Organizationally Unique Identifier, OUI part)
 - Insert FF:FE
 - Take last 3 octets of MAC (NIC-specific part)
 - Invert the Universal/Local (U/L) flag (bit 7)
- Example:
 - MAC Address: 00:12:7F:EB:6B:40
 - 00:12:7F (OUI) + FF:FE + EB:6B:40 (NIC)
 - IPv6 Interface ID: 0212:7FFF:FEEB:6B40

Privacy Addressing

- MAC address in EUI-64 can be tracked as you roam across the Internet
- Privacy Addressing solves this by generating an Interface ID differently:
 - Apply a cryptographic hash algorithm to an EUI-64 or random 64-bit number
 - Perform Duplicate Address Detection (DAD)
 - Change the Interface ID periodically
- Privacy Addressing can be enabled/disabled
 - On by default in newer Windows versions

IPv6 Address Types

- Unicast - Address a single node
- Multicast - Address a group of nodes
 - FF00::/8 (vs. IPv4 224.0.0.0/4)
- Anycast - Address the "nearest" single node
- Broadcast - Address all nodes on the local network
 - IPv4 255.255.255.255
 - Doesn't exist in IPv6
 - Use multicast instead

IPv6 Address Scopes

- Interface-local
 - Loopback `::1/128` (vs. `127.0.0.1` in IPv4)
- Link-local
 - `FE80::/10` (sort of like `169.254.0.0/16` in IPv4)
 - Always assigned to every interface
 - Used for control protocols like ND, DHCPv6
- Site-local (deprecated, replaced by ULA)
 - `FEC0::/10`
- Unique local unicast (ULA, RFC 4193)
 - `FC00::/7`, randomly generated by each site independently
 - Not coordinated with any central authority

IPv6 Address Scopes (cont'd)

- Global
 - Currently allocated out of 2000::
- Embedded IPv4 Unicast
 - ::ffff:192.168.1.1 or ::192.168.1.1
 - Mostly shouldn't appear on the wire
 - Used by the BSD sockets API and often shows up in log messages and maybe some configs
- Special Address Types
 - Unspecified Address all-zeros :: used before a node has an address (e.g. DHCPv6)

ICMPv6

- Similar features to IPv4 ICMP
 - Echo Request/Reply (ping)
 - Error Notification (Destination Unreachable, Packet Too Big, Time Exceeded, etc.)
- Do NOT block all ICMPv6 or you will break IPv6
- Neighbor Discovery (replaces IPv4 ARP)
 - Neighbor Unreachability Detection (NUD)
 - Duplicate Address Detection (DAD)
 - Router Advertisements (RA)
 - Auto-configuration (SLAAC)

ICMPv6 (cont'd)

- Path MTU Discovery, minimum MTU 1280
 - Automatically finds the smallest end-to-end MTU
 - Eliminates the need for routers to fragment packets in transit (which isn't allowed in IPv6)
 - End host can fragment packets if necessary (but should avoid it for performance reasons)
- Multicast Listener Discovery (MLD)
 - Replaces the function of IGMP in IPv4
 - Used to manage multicast group membership
 - Multicast (at least link-local multicast) is *required* for IPv6 to function at all! (NS, ND, RA, etc.)

Neighbor Discovery

- Relies on IPv6 Link-Local Multicast
- Neighbor Solicitation (NS, like ARP Request)
 - Sent to determine the MAC address of a neighbor, or for Duplicate Address Detection (DAD)
- Neighbor Advertisement (NA, like ARP Reply)
 - Response to NS, or sent unsolicited to announce a MAC address change
- Router Solicitation (RS)
 - Sent to ask any routers to send RAs immediately

Neighbor Discovery (cont'd)

- Router Advertisement (RA)
 - Sent in response to RS and also periodically
 - Contains various information:
 - Router Address for use as a default gateway
 - Router Priority (to determine which router to use as a default gateway)
 - Available prefixes/prefix lengths (like IPv4 Subnet Mask)
 - Hop-Limit (like IPv4 TTL), MTU, etc.
 - Auto-configuration Information
 - RFC 6106 adds a DNS server option
- Watch out for Rogue Routers sending RAs!
 - e.g. Windows Internet Connection Sharing

Router Advertisements

- Also contains several important flags:
 - Managed Address Configuration flag (M)
 - Indicates that Stateful DHCPv6 is available
 - Other Configuration flag (O)
 - Indicates that Stateless DHCPv6 is available
 - Meaningless when the M flag is set
 - Autonomous Address Configuration flag (A)
 - Indicates that SLAAC can be used to configure an address with this prefix
 - Usually disabled when using Stateful DHCPv6

Address Assignment

- Static Addressing
- Stateless Address Auto Configuration (SLAAC)
 - Cannot hand out other info like DNS servers (except for RFC 6106 but this isn't widely implemented)
- DHCPv6 (Stateful or Stateless)
 - Can't hand out a default router or prefix-length (yet?)
 - Must use in combination with RAs to get default router
 - Can't easily assign reserved address to a specific MAC, must use DUID instead

Stateless Address Auto Configuration (SLAAC)

- Host sends RS and listens for an RA
- Host uses prefix information from RA to form top part of its address (prefix)
- Host uses EUI-64 or one of the Privacy Addressing methods to pick the lower 64 bits (interface ID)
- Host combines the two parts and performs Duplicate Address Detection (DAD)
- The same process is followed to create a link-local address by combining with the well-known prefix FE80::/10, no router required

DHCPv6

- Stateless or Stateful
 - Stateless is combined with SLAAC in order to hand out additional parameters
 - such as DNS servers, NTP servers, Proxy, etc.
 - DHCPv6 not used to assign host addresses
 - Stateful assigns IPv6 addresses to hosts
 - plus any additional parameters
- Prefix Delegation
 - DHCPv6 also supports handing out entire prefixes (subnets) to downstream routers
 - Can be a hierarchy of chained routers

DHCPv6 (cont'd)

- DHCPv6 uses UDP ports 546 (clients) and 547 (servers & relay agents)
 - Don't forget to allow these through firewall!
- DHCPv6 uses Link-Local Multicasts
 - FF02::1:2 - All_DHCP_Relay_Agents_and_Servers
 - FF05::1:3 - All_DHCP_Servers
- Four Packet Handshake
 - SOLICIT, ADVERTISE, REQUEST, CONFIRM
- Rapid Commit Option, Two Packets
 - SOLICIT, REPLY

DHCP Unique Identifier

- DHCP Unique Identifier (DUID)
 - Uniquely identifies servers and clients
 - Generated once at OS installation, should not change thereafter
 - Used instead of Client Hardware Address (chaddr) to identify clients and their leases
 - Required, not optional like Client Identifier in DHCPv4
 - Can use DUID with DHCPv4 as well (RFC 4361)
 - Allows Dynamic DNS updates to inter-operate between IPv4 and IPv6 so the same DNS name can have 'A' and 'AAAA' records added

DUID Formats

- Link-Layer Address + Time (DUID-LLT)
 - Based on time-of-generation and MAC address
 - Recommended when writable non-volatile storage is available
- Enterprise Number (DUID-EN)
 - IANA assigned Enterprise Number
 - Vendor assigned unique identifier
- Link-Layer Address (DUID-LL)
 - Recommended for devices that have a permanently-connected network interface, but otherwise don't have non-volatile writable storage

DUID Formats (cont'd)

- Universally Unique Identifier (DUID-UUID, RFC 6355)
 - Embeds an existing UUID (RFC 4122) which is stored in system firmware
 - Allows multistage booting to use a stable, unchanging DUID
 - e.g. PXE bootloaders
 - Easier to predict ahead of time what the DUID will be
 - Assuming all software stages use the same DUID format!

Obsolete Linux Commands

- ifconfig is obsolete, use ip!
 - ifconfig <interface> inet6 add <ipv6address>/<prefixlength>
- route is obsolete, use ip!
 - route -A inet6 add default gw <ipv6address>
- netstat is obsolete, use ss or ip!
 - netstat -A inet6 -r
 - netstat -A inet6 -nlp
- "arp -a" only works for IPv4

Current Linux Commands

- `ss` for Socket Statistics
 - `ss -6 -nlp`
- `"/sbin/ip"` for all your IP configuration needs:
 - `ip -6 [link|addr|neigh|route] show [dev <interface>]`
 - `ip -6 addr add 2001:db8:1::dead:beef/64 dev eth0`
 - `ip -6 route add default via 2001:db8:1::1`
- `"ip -6 neigh"` shows the Neighbor Discovery cache
 - IPv6 equivalent to the IPv4 ARP cache
- All these commands work for IPv4 too!

Red Hat/Fedora Configuration

- `/etc/sysconfig/network`
 - NETWORKING_IPV6=yes
 - IPV6FORWARDING=no
 - IPV6_AUTOCONF=no
 - IPV6_ROUTER=no
 - IPV6_AUTOTUNNEL=no
- `/etc/sysconfig/network-scripts/ifcfg-eth0`
 - IPV6INIT=yes
 - IPV6ADDR=2001:db8:1::dead:beef/64
 - IPV6_DEFAULTGW=2001:db8:1::1
 - or: DHCPV6C=yes, IPV6_AUTOCONF=yes (need BOTH of these if using DHCPv6!)

Troubleshooting Commands

- ping6
 - ping6 2001:db8:1:aaaa::1
 - Must use -I <interface> (also known as a Scope ID) when pinging link-local addresses:
 - ping6 -I eth0 fe80::216:3eff:fe59:34:46
 - OR use %<interface> for the Scope ID after the IPv6 link-local address:
 - ping6 fe80::216:3eff:fe59:34:46%eth0
- traceroute6
 - traceroute6 2001:db8:1:eeee::10

IPv6 Literals

- If you want to type an IPv6 literal address rather than a DNS name into some applications (web browsers, etc.) you need to put it in [brackets]:
 - `elinks 'http://[2001:468:616:882:211:43ff:fe2b:897d]/'`
- For link-local addresses, add the %<interface> Scope ID syntax:
 - `firefox 'http://[fe80::211:43ff:fe2b:897d%eth0]'`
 - BUT this doesn't work in most software! Firefox works luckily, Chrome not so much.

Domain Name System (DNS)

- DNS Transport over IPv6
 - Whether the DNS server listens/responds/queries over TCP/IPv6 and UDP/IPv6
- DNS IPv6 Resource Records
 - What records are actually stored and served by the DNS Zone
 - AAAA a.k.a. "Quad-A record"
 - PTR for Reverse DNS
- These two things are completely independent!
 - Can query 'AAAA' over IPv4 transport!

DNS Examples

- Using IPv6 Transport

- /etc/resolv.conf

search example.net.

nameserver 2001:db8:1:ffff::10

nameserver 2001:db8:1:eeee::10

- Configuring IPv6 Resource Records

- www.example.net IN AAAA 2001:db8:1::20

- 0.2.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.1.0.0.0.8.b.d.
0.1.0.0.2.ip6.arpa. IN PTR www.example.net.

Address Selection

- If there are multiple DNS records for a domain name, which one should be used?
 - One or more IPv4 'A' record(s)
 - One or more IPv6 'AAAA' record(s)
- Which IPv6 Source Address should be used?
- RFC 3484 Address Selection
 - `getaddrinfo()` API should be used instead of `gethostbyname()`
 - A set of ordering rules is applied to returned results

Address Selection (cont'd)

- Address Selection Policy Table
 - Longest-matching-prefix lookup table
 - Address types/scopes compared
 - Prefer best match
 - Results ordered by preferences in table
- What does this all mean?
 - Most OSes are configured to prefer native IPv6 over native IPv4 if both 'A' and 'AAAA' records are present
 - On Linux, the rules can be changed in `/etc/gai.conf`

Further Topics

- Transition Mechanisms
 - Tunneling
 - 6to4, ISATAP, Teredo, 6rd
 - Carrier-Grade NAT (CGN), Dual-Stack Lite
 - Tunnel Brokers
 - Translation
 - Stateless IP/ICMP Translation Algorithm (SIIT) RFC 2765
 - NAT-PT (RFC 2766 and RFC 4966)
 - NAT64/DNS64 (RFC 6146)
 - Bump in the Stack (BIS) RFC 2767
 - Transport Relay Translator (TRT) RFC 3142
 - Bump in the API (BIA) RFC 3338

Further Topics (cont'd)

- Security
 - Internet Protocol Security (IPsec)
 - Secure Neighbor Discovery (SEND)
 - RA Guard - block Rogue RAs
 - DHCPv6 Snooping, ND Inspection, IPv6 Source Guard
 - IPv6 Firewalls
 - ip6tables (/etc/sysconfig/ip6tables)
 - Intrusion Detection/Prevention
 - Bandwidth Management

Further Topics (cont'd)

- Routing
 - OSPFv3, IS-IS, BGP
- Address Numbering Plans
 - How to divvy up all that address space
- Deployment Strategies
 - Native Dual-Stack vs. Tunneled vs. Translated
 - Core, External Services, Internal Services, Clients
 - Incremental Deployment
 - IPv6 ACLs on switch ports
 - IPv6 Protocol VLANs

References

- Guidelines for the Secure Deployment of IPv6, National Institute of Standards and Technology
<http://csrc.nist.gov/publications/nistpubs/800-119/sp800-119.pdf>
- IANA Free Pool of IPv4 Address Space Depleted
<http://www.nro.net/news/ipv4-free-pool-depleted>
- The IPv4 Depletion Site
<http://www.ipv4depletion.com/>
- A Regular Expression for IPv6 Addresses
<http://forums.dartware.com/viewtopic.php?t=452>
- EUI-64 in IPV6
<http://packetlife.net/blog/2008/aug/4/eui-64-ipv6/>

References (cont'd)

- IPv6 ULA (Unique Local Address) RFC4193 registration

<http://www.sixxs.net/tools/grh/ula/>

- Starting Over from the Top: Campus IPv6 Deployment and Security, Phil Deneault

<http://www.educause.edu/Resources/StartingOverfromtheTopCampusIP/203126>

- Linux IPv6 HOWTO

<http://www.tldp.org/HOWTO/Linux+IPv6-HOWTO/>

- 6 Surprising Facts about IPv6

<http://www.breakingpointsystems.com/community/blog/6-surprising-facts-about-ipv6/>

References (cont'd)

- Hurricane Electric Tunnel Broker

<http://tunnelbroker.net/>

- ARIN IPv6 Wiki

<http://www.getipv6.info/>