IPv6 Fundamentals

August 2022

CMU

Software Engineering Institute
Carnegie Mellon University
Pittsburgh, PA 15213

Markings

This material is based upon work funded and supported by the Department of Defense under Contract No. FA8702-15-D-0002 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center.

DM22-0665

Introduction

Basic Concepts

Addressing

Dynamic Network Behavior

Integration and Transition Technologies

Design Principles: Why IPv6?

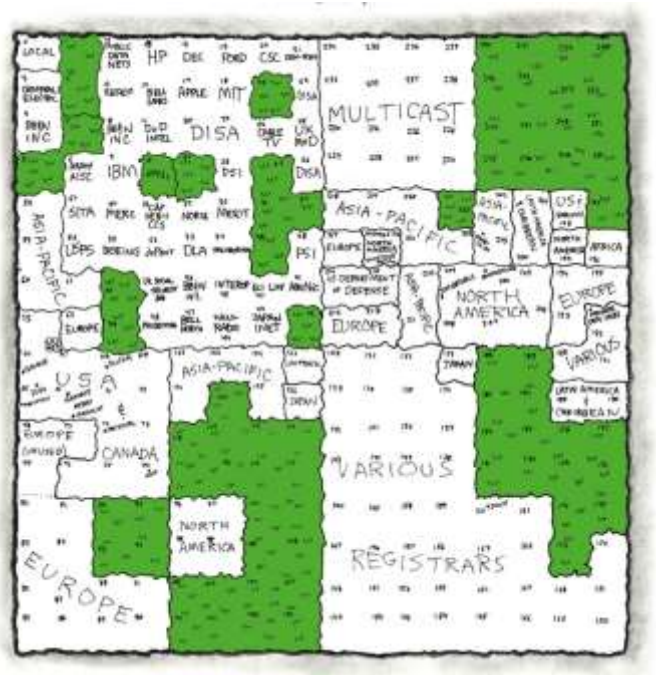

Source: https://xkcd.com/195/

"For the IPv6 map just imagine the XP default desktop picture."

The Windows XP default desktop picture

Design Principles: Why IPv6?

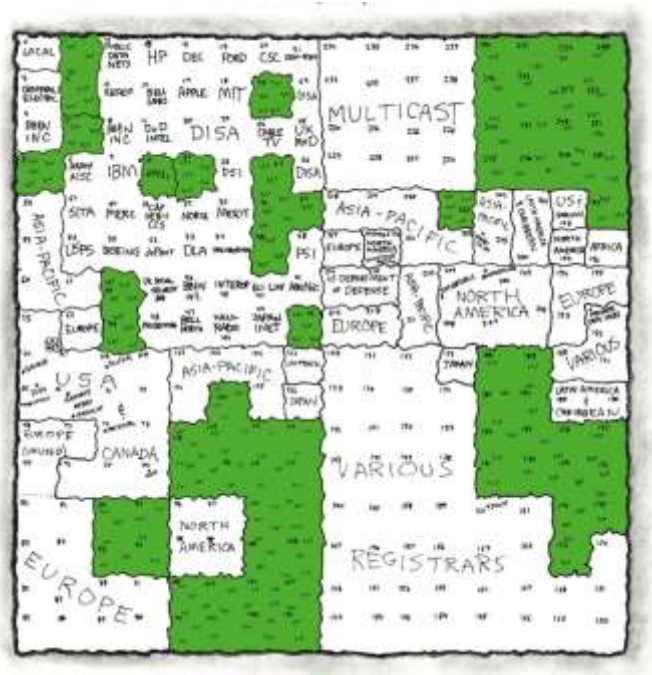

Source: https://xkcd.com/195/

1981: RFC 791

DARPA INTERNET PROGRAM PROTOCOL SPECIFICATION (IPv4)

1998: RFC 2460

Internet Protocol, Version 6 (IPv6) Specification

Concerns

- Address Exhaustion
- Transmission Efficiency
- Security and privacy
- Mobility

IPv6 is a mature protocol–and it's growing

The Good
- Well-tested
- Strong hardware/software support

The Bad
- "Evolutionary dead-ends"
  - NAT-PT

The Ugly
- Dual-stack
- Backporting (IPSec)
- Confusing similarities to IPv4

Per-Country IPv6 adoption

Source: https://www.google.com/intl/en/ipv6/statistics.html

Basic Concepts

IPv4 vs. IPv6 – The Headers

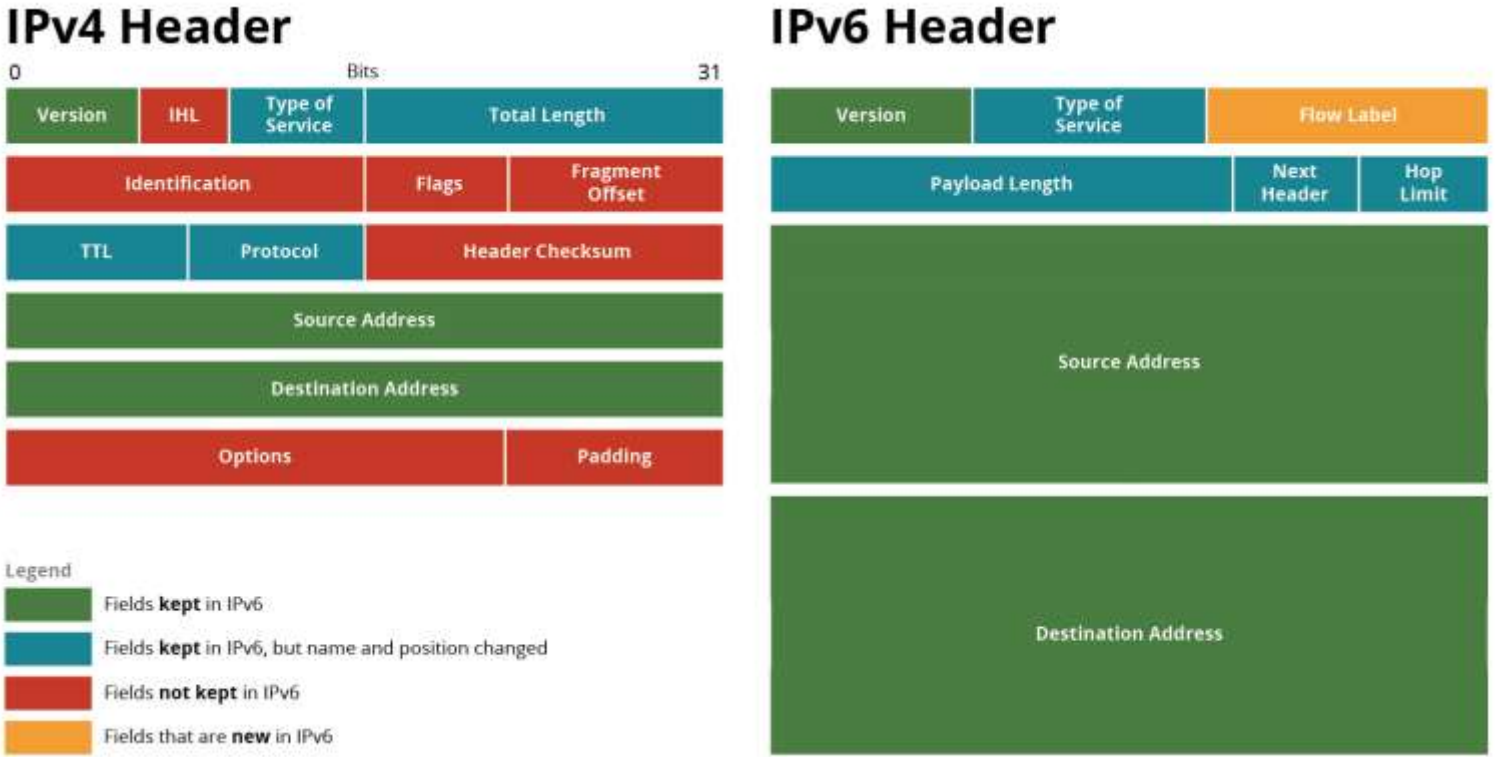

IPv4 vs. IPv6 – The Addresses

IPv4

- 32-bit address
  - **4.3 x 10^9** unique addresses
- Four 8-bit ranges (**"Octets"**)
- Decimal notation

192.168.1.1

IPv6

- 128-bit address
  - **3.4 x 10^{38}** unique addresses
- 8 16-bit ranges (**"Hextets"**)
- Hexadecimal notation w/abbreviations

1234:d3ad:ca01:fe71:b100:cafe:d3d1:82d8

Representing IPv6 Addresses

Lowercase is recommended:

1234:D3AD:… → *1234:d3ad:…*

Leading zeros in each hextet may be omitted:

2001:db8:0000:0000:0000:0000:0002:0001

→ *2001:db8:0:0:0:0:2:1*

Runs of zero-value hextets may be elided:

2001:db8:0:0:0:0:2:1 → *2001:db8::2:1*

IPv6 loopback address:

::1 (or *::1/128*)

ULA network range:

fc00::/7 (*broken into* *fc00::/8* and *fd00::/8*)

Special Cases

IPv4 embedded Address

- Used to represent IPv4 address inside an IPv6 address
- A special text representation may be applied ONLY when IPv4 is the last 32 bits of an IPv6 address. (i.e., network prefix is 96 bits)
- Examples
  - *::192.168.1.2* (Prefix is *::/96*; deprecated since RFC 4291)
  - *::ffff:192.168.1.2* (Prefix is *::ffff/96*; used as example in RFC 4291)
  - *64:ff9b::192.168.1.2* (Prefix is the “well-known” prefix ”for use in an algorithmic mapping”; see RFC 6052)
  - Other prefixes are possible

::/128 – Unspecified; cannot be assigned to an interface.

- Used to indicate absence of an address

CIDR and addressing

Global unicast address

- Global Routing prefix: 48 bits (/48)
  - Network portion of address- assigned by provider
- Subnet ID: 16 bits (/64)
- Interface ID: 64 bits (/128)
  - A host (or “node”) may have multiple interfaces.
  - All-0s and All-1s are legal interface addresses

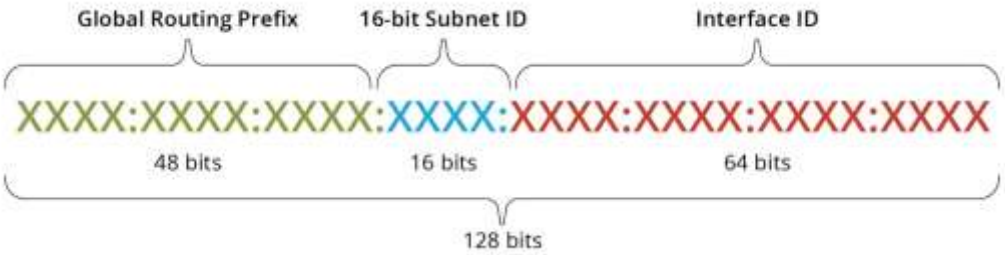

- 64-bit interface ID ≅ 18 quintillion ≅ 1.84 x 10^{19} (18,446,744,073,709,551,616) devices/subnet
- 16-bit-Subnet ID (initially recommended) = 65,536 subnets

Addressing

Unicast vs Multicast vs Broadcast vs Anycast

Unicast
- 1:1
- Identifies an interface on a host
- Global, unique local, link-local, unspecified, loopback

Multicast
- 1:many
- Addresses used as ”listening groups”

Broadcast
- 1:all
- Unavailable in IPv6 (use all-nodes multicast instead)

Anycast
- 1:1 (of many)
- Routed to the “nearest” device, determined by lowest number of network hops

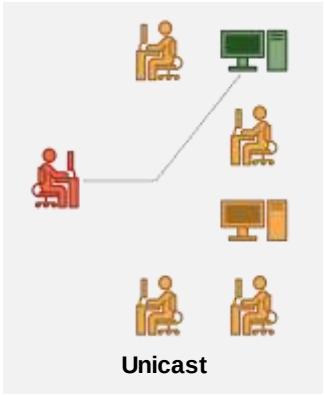

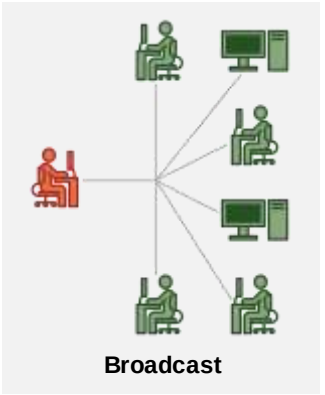

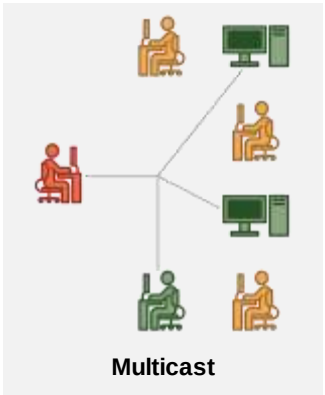

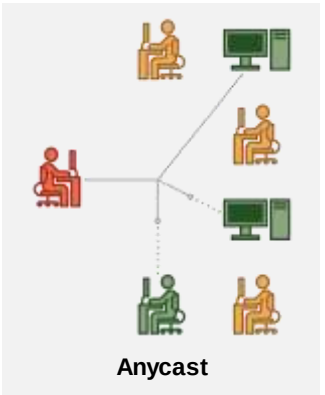

IPv6 Addressing Model

Addresses are assigned to **interfaces**

- An interface may have **more than one address**
- An interface for a host **MUST** have:
  - A *link-local* unicast address
  - The *loopback* address
  - Any other unicast or anycast addresses that are configured
  - Multicast addresses for groups to which it belongs
- An interface for a router must have additional addresses to support anycast/multicast

Prefix	First Hextet	Allocation
0000::/8	0000-00ff	Special Cases: Unspecified, Loopback, Embedded IPv4
2000::/3	2000-3fff	**Global Unicast (GUA)**–Publicly routable
fc00::/7	fc00-fdff	**Unique Local Unicast (ULA)**–Non-publicly routable
fe80::/10	fe80-febf	**Link-Local Unicast**–Not routable
ff00::/8	ff00-ffff	**Multicast**–Addresses for multicast groups

Global Unicast Addressing (GUA)

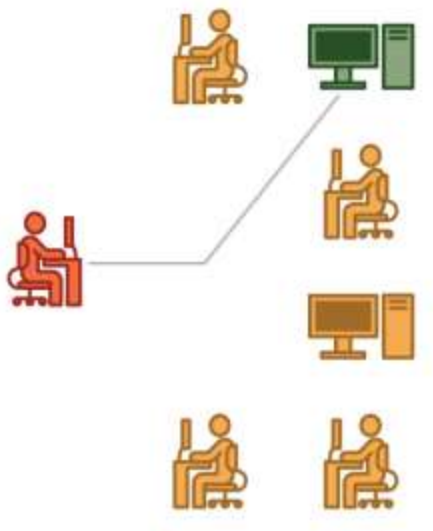

Expected to be globally unique and routable in the global internet

Begin with *2000::* through *3fff::*

Provider Assigned (PA)-ISP

- Provider assigned/aggregatable address space (PA) a block of IP address assigned through ISP and come from a large block allocated to a Provider by RIPE.

Provider Independent (PI)

- Addresses assigned by a regional Internet registry (RIR) directly to an end-user organization.
- End-users can
  - Change service providers without renumbering their networks
  - Use multiple access providers in a multi-homed configuration.

Unique Local Addressing (ULA)

Expected to be globally unique but should not be routable in the global internet

- Meant to assist with merging local networks

Begins with *fc00::* to *fdff::*

Analogous to RFC 1918 space in IPv4. **BUT:**

- Many-to-one NAT is unnecessary and problematic in IPv6
  - If "NAT-like" traffic patterns are desired, implement firewall policies
- Without NAT, there are very few use cases for ULA
  - One exception: Avoiding network renumbering when using PA GUA addresses and multiple ISPs

In dual-stack networks, ULA networks will be unused by default, per RFC 6724

Multicast Addressing

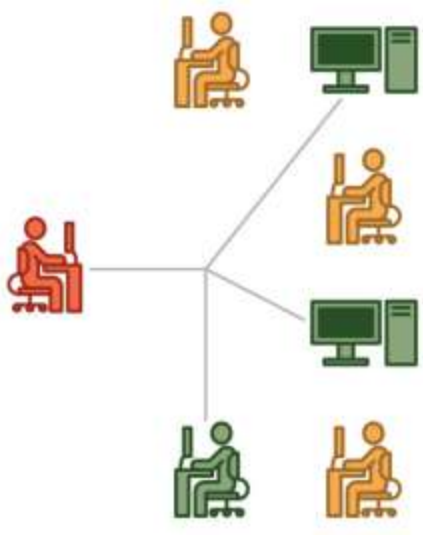

Multicast address defines a group of devices known as a multicast group

- Equivalent to *224.0.0.0/4*

Multicast Address format: *ff00::/8*

- Bits 8-11: Flags
  - O (unused), R (rendezvous), P (prefix) T (transient)
- Bits 12-15: Scope
  - Defines propagation range (interface local to global)
- Bits 13-127: Group ID

Some ”permanent” multicast addresses defined/reserved

Anycast

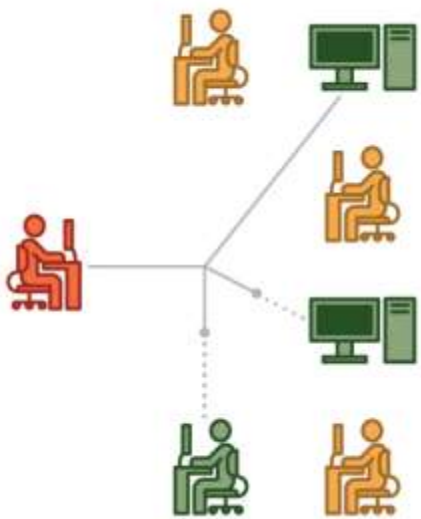

Communication between a unicast interface and “any” of a number of identically-provisioned interfaces providing the same service

No special prefix; logically, an anycast address is a unicast address assigned to multiple interfaces

- Routers must be configured to properly handle addresses as anycast

Required anycast address: subnet-routers

- Address format: prefix + all zeros (e.g., *1234:5678::)*
- Group of all routers that route that prefix

Dynamic Network Behavior

Address Initialization Options

Static Configuration

- Useful for network infrastructure

DHCP (DHCPv6)

- Useful for centralizing address assignment/network access
- Can be stateful or stateless

Stateless Automatic Address Configuration (SLAAC)

- Useful for highly decentralized/dynamic address assignment/network access

All (except static) utilize **Neighbor Discovery Protocol (NDP)**

- Based on ICMPv6 and multicast

Neighbor Discovery Protocol

ICMPv6 messages

- Router Solicitation (RS) messages
- Router Advertisement (RA) messages
- Neighbor Solicitation (NS) messages
- Neighbor Advertisement (NA) message
- Redirect Messages
- Duplicate Address Detection (DAD)-NS message

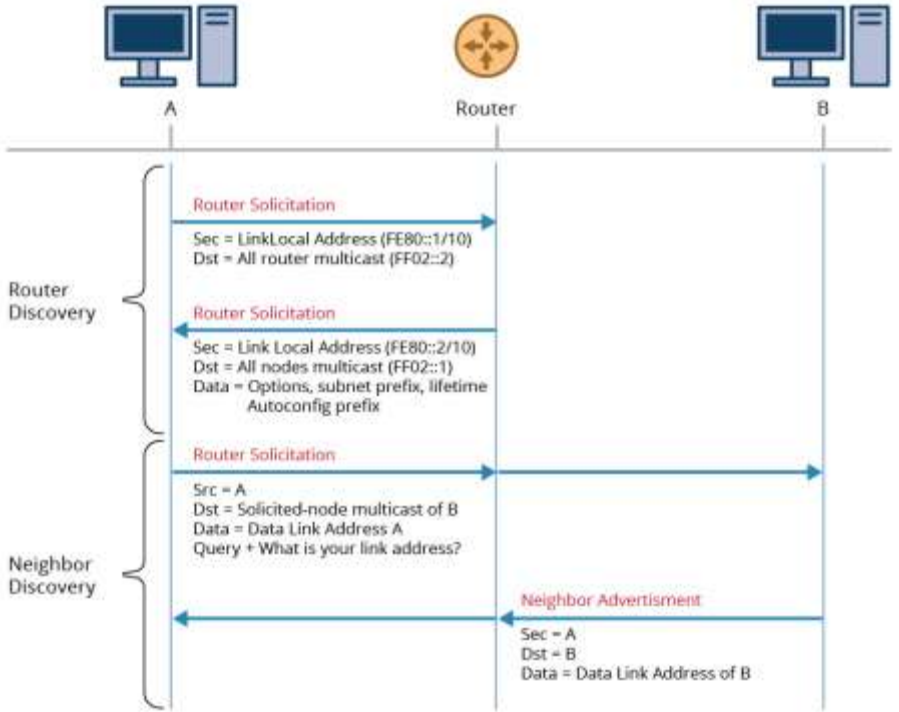

DHCPv6 and SLAAC – Initialization Sequence

SLAAC

- Discover routers and network prefix
- Compute Extended Unique Identifier (EUI-64) from MAC address
  - But see RFC 4941 for Privacy Extensions

DHCPv6 Stateful

- Full port of DHCPv4 to IPv6
- Location of DHCP server communicated to host in Router Advertisement

DHCPv6 Stateless

- After SLAAC, discover and communicate with DHCPv6 server
- For providing additional network information (e.g. DNS servers)

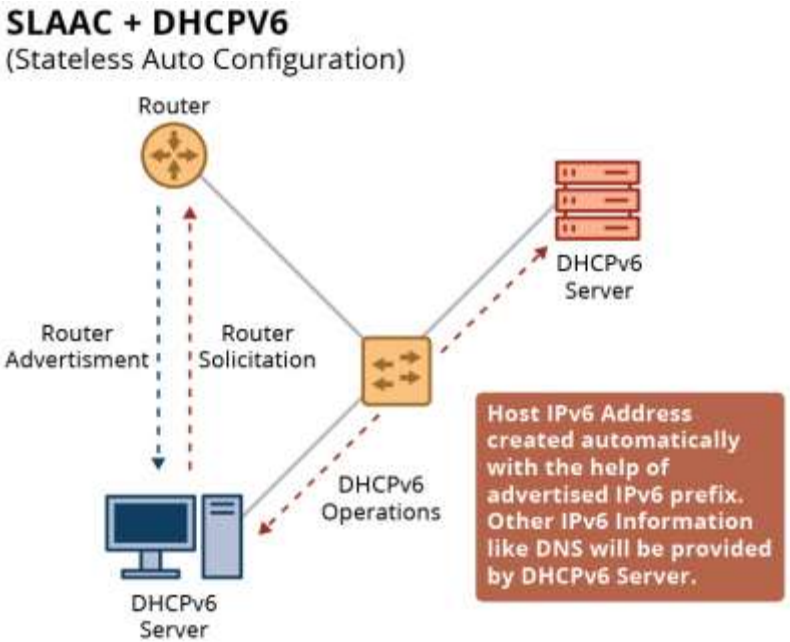

Mobile IPv6

RFC 6275 – “[A] protocol that allows nodes to remain reachable while moving around in the IPv6 Internet.”

- Mobile node has a “care of” address
  - Traffic to the “care of” address is sent to the mobile node’s current address (the “mobile address”)
- New headers (Mobility Header, Type 2 Routing Header)
- New ICMPv6 messages defining new protocol
  - Home Agent Address Discovery Request/Reply
  - Mobile Prefix Solicitation/Advertisement
- New Information Options describing home address in RA messages

Integration and Transition Technologies

DNS

A (address)- used in IPv4, used to map a fully qualified domain name (FQDN) to an IPv4 address and acts as a translator by converting domain names to IP addresses.

AAAA (quad-A) record, also known as "IPv6 address record", maps a hostname to a 128-bit IPv6 address in the Domain Name System (DNS) Internet Authentication Service .

DNS AAAA records are exactly like DNS A records, except that they store a domain's IPv6 address instead of its IPv4 address.

New domain: IP6.ARPA (to support reverse lookups)

Dual-Stack

Configure IPv4 networking in parallel on system with IPv6

- Best network chosen based on heuristics (e.g., RFC 6724)

Dual Stack Lite (DS-Lite)

- Tunnel-based IPv6 transition solution for ISPs with IPv6 infrastructure to connect IPv4 subscribers to the Internet.
- Unrelated to endpoint dual-stack

IPv6, NAT, and NPT

Traditional NAT is discouraged

- RFC 4864 – **Local Network Protection for IPv6**

  "[T]his document shows how Local Network Protection (LNP) using IPv6 can provide the same or more benefits without the need for address translation."

Network Prefix Translation (NPTv6)

- RFC 6296 (Experimental)
- Stateless 1:1 translation
- Use case: Insulating local network from external prefix changes

NAT64

- IPv4 transition mechanism to permit IPv4/6-only hosts to reach nodes on the "other" network

Questions

pgroce@sei.cmu.edu

References

RFCs:

RFC 2373: IP Version 6 Addressing Architecture

RFC 2526: Reserved IPv6 Subnet Anycast Addresses

RFC 3307: Allocation Guidelines for IPv6 Multicast Addresses

RFC 3596: DNS Extensions to Support IP Version 6

RFC 4291: IP Version 6 Addressing Architecture

RFC 4861: Neighbor Discovery for IP version 6 (IPv6)

RFC 5952: A Recommendation for IPv6 Address Text Representation

RFC 6052: IPv6 Addressing of IPv4/IPv6 Translators

RFC 6296 (Experimental): IPv6-to-IPv6 Network Prefix Translation

RFC 6437: IPv6 Flow Label Specification

RFC 6724: Default Address Selection for Internet Protocol Version 6 (IPv6)

RFC 7371: Updates to the IPv6 Multicast Addressing Architecture

Other References (order of appearance):

Naming IPv6 address parts

Statista: Number of connected devices worldwide in 2014 and 2020, by device

What are Provider Aggregatable (PA) addresses and Provider Independent (PI) addresses?

IPv6 Unique Local Addresses (ULA) Made Useless

Cisco: NPTv6 Support

Juniper: Stateless Source Network Prefix Translation for IPv6

Tunneling

GRE

- Developed by Cisco Systems that can encapsulate a wide variety of network layer protocols inside virtual point-to-point links or point-to-multipoint links over an Internet Protocol network.

ISATAP (Intra-Site Automatic Tunnel Addressing Protocol)

- An IPv6 transition mechanism meant to transmit IPv6 packets between dual-stack nodes on top of an IPv4 network
- Uses IPv4 as a virtual nonbroadcast multiple-access network (NBMA) data link layer, does not require the underlying IPv4 network infrastructure to support multicast.

6to4

- 6to4 is utilized during the initial phases of deployment to full, native IPv6 connectivity, since IPv6 is not required on nodes between the host and the destination.
- The first 16 bits of the prefix are always *2002:*, the next 32 bits are the IPv4 address, and the last 16 bits of the prefix are available for addressing multiple IPv6 subnets behind the same 6to4 router.

Tunneling detection

"41" is the protocol number assigned to be used for automatic IPv6 in IPv4 tunneling mechanisms such as 6to4 (RFC3056) or ISATAP (RFC5214).

Another automatic tunneling mechanism is Teredo (RFC4380). Teredo encapsulates IPv6 in IPv4 and uses UDP. IPv6 through Teredo can bypass NAT and firewall devices. The default UDP port number used by Teredo is 3544.