IPv6 at Penn

Shumon Huque
Information Systems & Computing

September 26th 2013
University of Pennsylvania
Philadelphia, PA

twitter
@shuque

Who am I?

- An I.T. Director at the University of Pennsylvania
- Have also been:
  - Programmer (C, Perl, Python, Lisp)
  - UNIX Systems Administrator
  - Network Engineer
- Education: B.S. and M.S. (Computer Science) from Penn
- Also teach a Lab course on Network Protocols at Penn's School of Engineering & Applied Science

Who am I?

- Website: http://www.huque.com/~shuque/
- Blog: http://blog.huque.com/
- Twitter: https://twitter.com/shuque

twitter
@shuque

- Google Plus:
  - https://plus.google.com/105308234918217701741/posts

IPv6 Motivation

World IPv6 Launch

WORLD IPv6 LAUNCH
LAUNCH INTO THE FUTURE 06.06.12

- http://www.worldipv6launch.org/

6 JUNE 2012

Major Internet service providers (ISPs), home networking equipment manufacturers, and web companies around the world are coming together to permanently enable IPv6 for their products and services by 6 June 2012.

- Google, Facebook, Netflix, Yahoo!, MS Bing, ...
- ISPs: Comcast, AT&T, Free Telecom, Time Warner, ...
- CDNs: Akamai, Limelight, ...
- Some universities, corporations, government agencies, ...

http://www.google.com/intl/en/ipv6/statistics.html

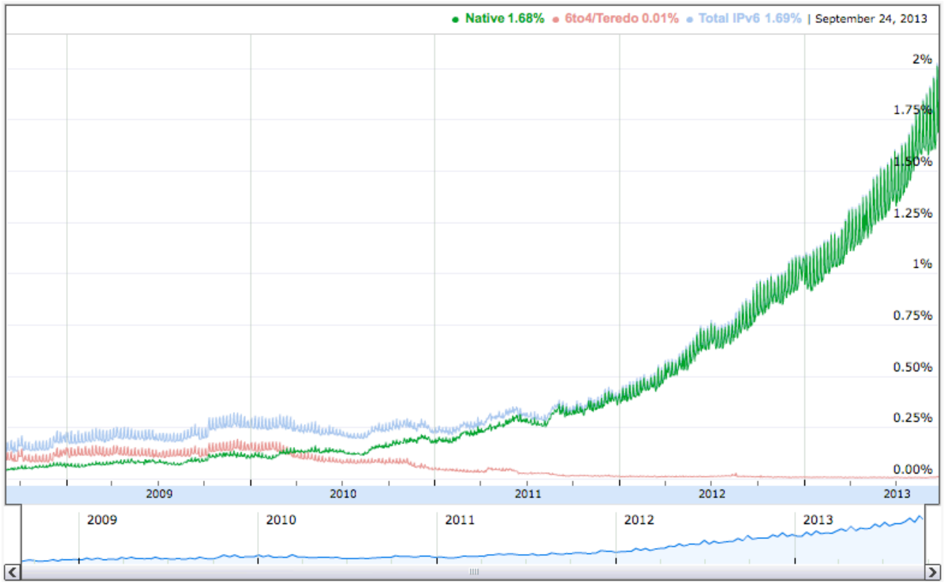

Also see http://www.worldipv6launch.org/measurements/
And http://www.vyncke.org/ipv6status/

IPv6: Internet Protocol v6

- Version 6: The next generation Internet Protocol
- Much larger address space: 128 bits vs 32 bits
  - (Note: not 4x larger, but 2^{96} times larger!)
- No NAT (goal: restore end-to-end architectural model)
- Scalable routing (we'll talk about multihoming later)
- Other: header simplification, NDP (a better version of ARP), auto-configuration, flow labelling, and more ..
- Note: *IPv6 is not backwards compatible with IPv4*

IPv6: Internet Protocol v6

- But primary impetus is the much larger address space
- Impending exhaustion of IPv4 addresses
- But Internet continues to grow
  - Not only in terms of the number of users, but also in the number and range of devices being connected to the network
  - The "*Internet of Things*"

Adverse Consequences of not deploying IPv6?

IPv4 Transfer Markets

- IPv4 transfer markets (sanctioned or unsanctioned)
  - March 2011: Microsoft acquired block of 600,000 addresses from Nortel for $7.5 million ($11.25/address)
  - December 2011: Borders books sold a /16 to Cerna for $786,432 ($12.00/address)
  - Rise of brokering companies: Addrex, Kalorama, Hilco streambank etc

More NAT

- More NAT, and more layers of NAT
  - Carrier Grade NATs (CGN), Large Scale NATs (LSN)
- Damaging impacts on applications
- Other implications of large scale address sharing
  - single points of failure, performance bottleneck, easy DoS target, inability to do geo-location, impacts on ACLs, blocklists, port space rationing, resource management, NAT traversal method reliability, ALG complexity, and more ..

Balkanization

- Islands of IPv4 and IPv6
- and resulting disruption of universal Internet connectivity

Transition vs Co-existence

- IPv4 isn't going away anytime soon, possibly not for many decades
- So, for most folks, already connected to the IPv4 Internet, ***we are not (yet) transitioning to IPv6***
- We are ***deploying IPv6 to co-exist with IPv4***
- To allow us to communicate with both the IPv4 and IPv6 Internet
- Note: some people in the near future will be moving directly to IPv6 though, due to IPv4 depletion

IPv6: Brief History

- Design work began by IETF in 1993, to deal with projected depletion of IPv4 addresses (then ~ 2010-2017)
- Completed in ~1999
  - RFC 1883: first version of IPv6 specification (Dec 1995)
  - RFC 2460: Internet Protocol version 6 specification (Dec 1998)
- April 1999: first RIR allocation of IPv6 address space
- By now hundreds of RFCs exist, describing various aspects of IPv6 and its applications

from http://ipv4.potaroo.net (Geoff Huston, APNIC)

(from September 26th 2013)

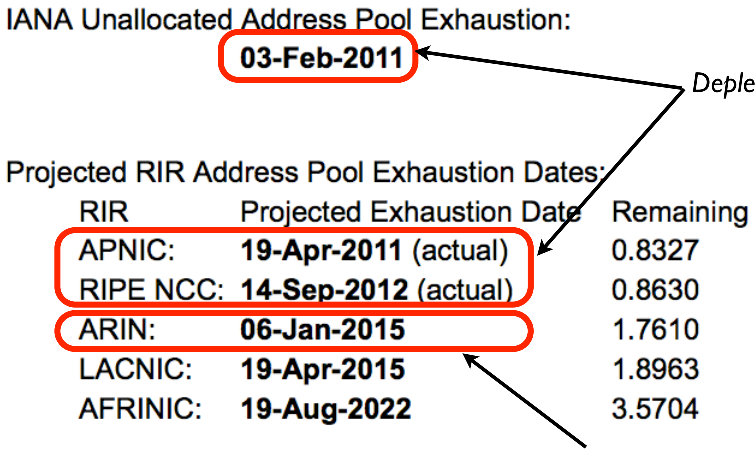

IANA Unallocated Address Pool Exhaustion:

03-Feb-2011

Projected RIR Address Pool Exhaustion Dates:

RIR	Projected Exhaustion Date	Remaining
APNIC:	**19-Apr-2011** (actual)	0.8327
RIPE NCC:	**14-Sep-2012** (actual)	0.8630
ARIN:	**06-Jan-2015**	1.7610
LACNIC:	**19-Apr-2015**	1.8963
AFRINIC:	**19-Aug-2022**	3.5704

Depleted already!

See https://www.arin.net/resources/request/ipv4_countdown.html

Some other projections for ARIN are more pessimistic. See for example Tony Hain's: http://tndh.net/~tony/ietf/ARIN-runout-projection.pdf

IPv6 addressing and protocol details

IPv4 addresses

- Example: 192.168.7.13
- 32 bits
- “Dotted Quad notation”
- Four 8-bit numbers (“octets”) in range 0..255, separated by dots
- 2^{32} = 4.3 billion (approximate) possible addresses
  - *(Usable number of addresses much lower though: routing & subnet hierarchies - see RFC 3194 - Host Density ratio)*

IPv6 addresses

- 128-bits (four times as large)
- 8 fields of 16 bits each (4 hex digits) separated by colons (:)
- [Hex digits are: 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, a, b, c, d, e, f]
- 2^{128} possible addresses (an incomprehensibly large number)

```
2001:0db8:3902:00c2:0000:0000:0000:fe04
```

(2^{128} = 340,282,366,920,938,463,463,374,607,431,768,211,456)

IPv6 addresses

- Zero suppression & compression for more compact format
  - Suppress (omit) leading zeros in each field
  - Replace consecutive fields of all zeros with a double colon (::) - only one sequence of zero fields can be compressed this way

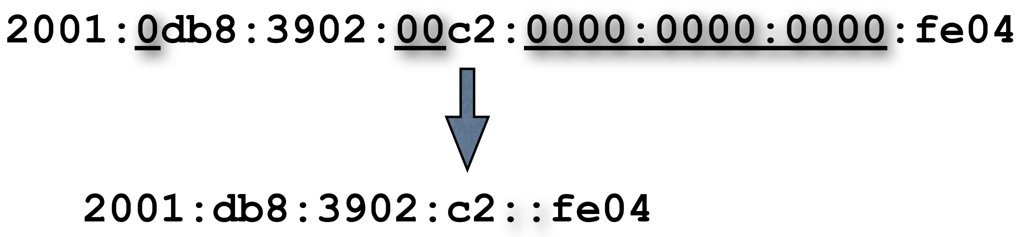

```
2001:0db8:3902:00c2:0000:0000:0000:fe04
```

```
2001:db8:3902:c2::fe04
```

IPv6 canonical form

- RFC 5952: A recommendation for IPv6 Text Representation
- Same address can be represented many ways in IPv6, making it more challenging to do some tasks (searching, pattern matching, programmatic processing of the text forms, etc)
- Define a (recommended) canonical text representation
  - must suppress leading zeroes in a field
  - Use :: to compress only the longest sequence of zero fields, and only the first one if there are multiple equal length sequences
  - Compression of a single zero field is not allowed
  - a, b, c, d, e, f must be in lower case

IPv4 mapped IPv6 address

```
Uses prefix ::ffff:0:0/96
          ( 0:0:0:0:0:ffff:0:0/96 )
Example     ::ffff:192.0.2.124
```

- Used for handling IPv4 connections on an IPv6 socket
- Note slightly different text representation to make it easier to embed 32-bit IPv4 address in the IPv6 address
- See RFC 4038 for details (“Application aspects of IPv6 transition”)
- Not normally seen on wire (only IPv4 packets seen)

IPv6 in URLs

- To represent literal IPv6 addresses in Uniform Resource Locators (URL), enclose the address in square braces:
  - `http://[2001:db8:ab:cd::3]:8080/index.html`
  - `ldap://[2001:db8:ab:cd::4]/`
  - `ftp://[2001:db8:ab:cd::5]/blah.txt`
- See RFC 3986 for details [URI: Generic Syntax]
- (This is generally only needed for debugging and diagnostic work)

IPv6 network prefixes

- Format: IPv6-Address / prefix-length

- `2001:db8::/32`
- `2001:db8:ab23::/48          (typical org assignment)`
- `2001:db8:ab23:74::/64       (most subnets)`
- `2001:db8:ab23:74::2/64`
- `2001:db8:ab23:75::1/127     (p2p links by some)`
- `2001:db8:ab23:76::a/128     (loopback)`

IPv6 DNS records

- **AAAA** ("**Quad-A**") DNS record type is used to map domain names to IPv6 addresses
- IPv4 uses the "**A**" record
- There was another record called **A6**, which didn't catch on (and now declared historic by RFC 6563)

```
www.ietf.org.  1800 IN  A     12.22.58.30
www.ietf.org.  1800 IN  AAAA  2001:1890:123a::1:1e
```

IPv6 Reverse DNS

- As in IPv4, PTR records are used for reverse DNS
- Uses "**ip6.arpa**" subtree (IPv4 uses "in-addr.arpa")
- The LHS of the PTR record ("owner name") is constructed by the following method:
  - Expand all the zeros in the IPv6 address
  - Reverse all the hex digits
  - Make each hex digit a DNS label
  - Append "ip6.arpa." to the domain name (note: the older "ip6.int" was formally deprecated in 2005, RFC 4159)

IPv6 reverse DNS example

```
host1.example.com.  IN AAAA 2001:db8:3902:7b2::fe04
```

```
2001:db8:3902:7b2::fe04                         (orig IPv6 address)

2001:0db8:3902:07b2:0000:0000:0000:fe04         (expand zeros)

20010db8390207b2000000000000fe04                (delete colons)

40ef0000000000002b7020938bd01002                (reverse digits)

4.0.e.f.0.0.0.0.0.0.0.0.0.0.0.0.2.b.7.0.2.0.9.3.8.b.d.
0.1.0.0.2                                       (make DNS labels)

4.0.e.f.0.0.0.0.0.0.0.0.0.0.0.0.2.b.7.0.2.0.9.3.8.b.d.
0.1.0.0.2.ip6.arpa.                             (append ip6.arpa.)
```

```
4.0.e.f.0.0.0.0.0.0.0.0.0.0.0.0.2.b.7.0.2.0.9.3.8.b.d.
0.1.0.0.2.ip6.arpa.  IN PTR host1.example.com.
```

IPv4 Header

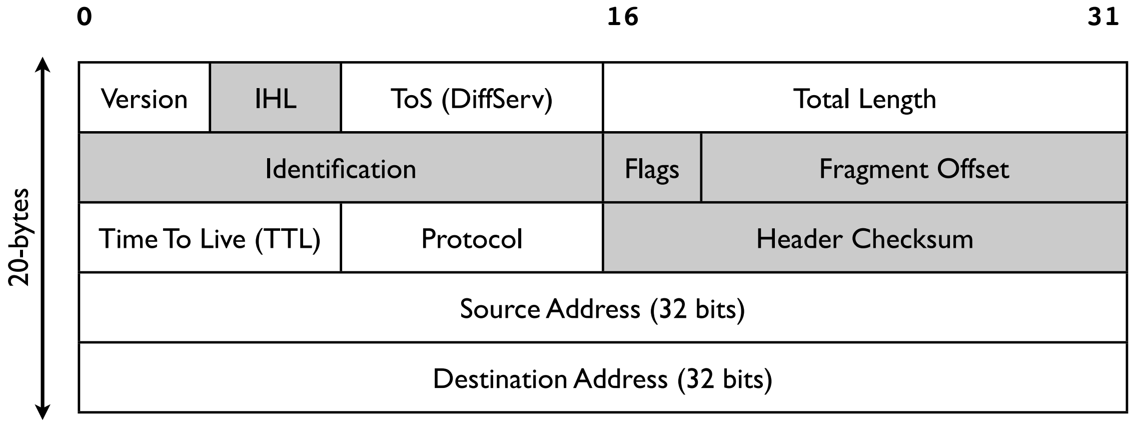

[Followed optionally by Options and Padding]

IPv6 Header

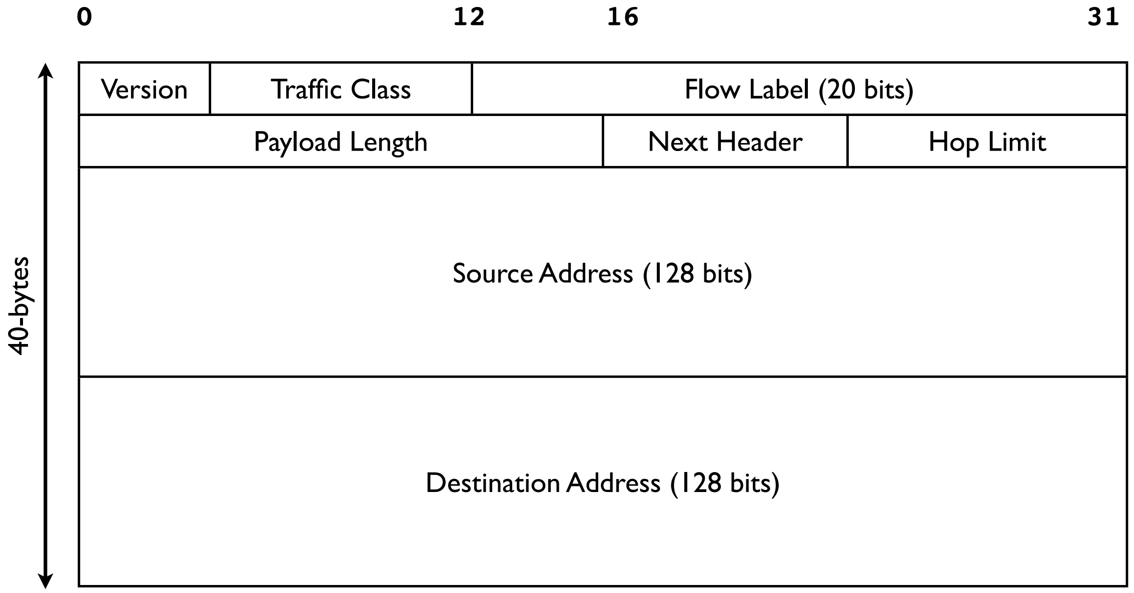

[Followed optionally by Extension Headers & Upper Layer payload]

Upper Layer payload

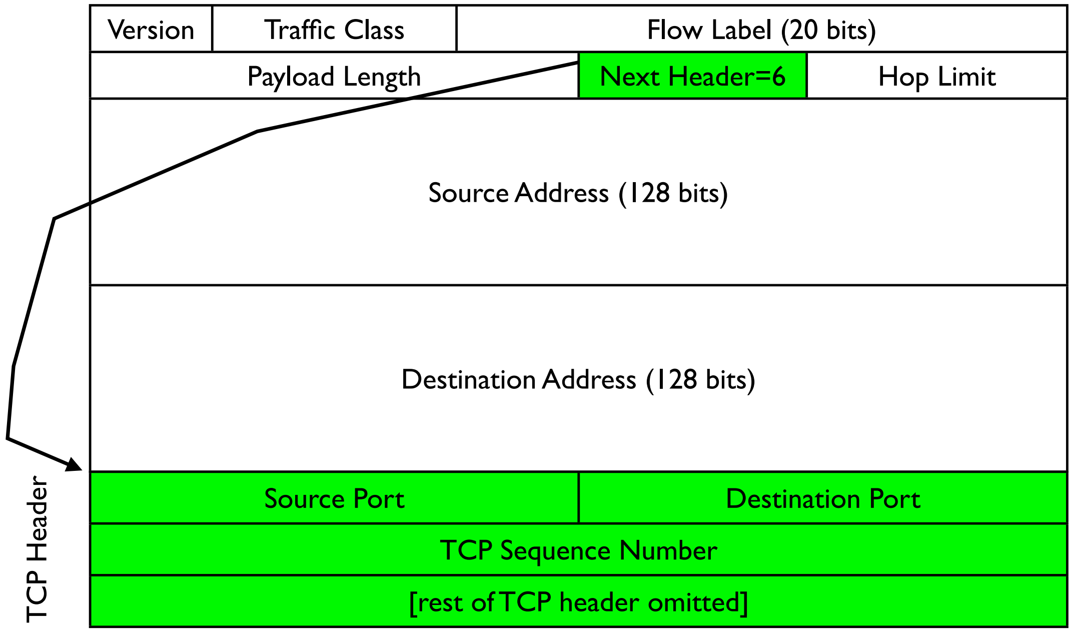

Version	Traffic Class	Flow Label (20 bits)

Payload Length	Next Header=6	Hop Limit

Source Address (128 bits)

Destination Address (128 bits)

Source Port	Destination Port

TCP Sequence Number

[rest of TCP header omitted]

Extension Headers

IPv6 Header Next Hdr = 6 (TCP)	TCP Header & Payload

IPv6 Header Next Hdr = 43 (Routing)	**Routing Hdr** Next Hdr = 6 (TCP)	TCP Header & Payload

IPv6 Header Next Hdr = 43 (Routing)	**Routing Hdr** Next Hdr = 51 (AH)	**AH Hdr** Next Hdr = 6 (TCP)	TCP Header & Payload

Extension Headers

- Hop-by-Hop (must be examined by all routers along path: eg. router alert)
- Destination Options (can appear after hop-by-hop when RH used, or at end)
- Routing (Note: RFC 5095, Dec 2007, deprecated RH type 0)
- Fragment (fragments less common in v6 because of path MTU discovery)
- Authentication (IPsec AH)
- Encapsulating Security Payload (IPsec ESP)
- Others: MIPv6, HIP, SHIM6, ...

[See also RFC 6564 - A Uniform Format for IPv6 Extension Headers]

IPv6 Address Types

- **Unicast**
- **Multicast**
- **Anycast**

- Note: there is no "**broadcast**" in IPv6

Unicast Address Types

- **Global** Unicast Addresses
  - Static, Stateless Address Autoconfiguration, DHCP assigned
  - Tunneled (6to4, Teredo, ISATAP, ...)
  - Others (CGA, HIP, ...)
- **Link Local** Addresses
- **Unique Local Addresses** (ULA)
- **Loopback** (::1)
- **Unspecified** (::)

Also see RFC 6890: Special Purpose IP address registries
and RFC 6666: IPv6 Discard Prefix

Link Local Addresses

- All IPv6 network interfaces have a Link Local address
- Special address used for communication on local subnet
- Self assigned in the range fe80::/10 (actually the subset fe80::/64)
- Last 64-bits derived from MAC address (EUI-64)
- Could be the same on multiple physical interfaces
- Often written with scope-id to differentiate interface
- fe80::21b:63ff:fe94:9d73%en0

↑ **Modified EUI-64** ↑ **scope-id**

Global IPv6 address form

48-bits	16-bits	64-bits
Global Routing prefix	SubnetID	Interface ID (host part)
001 + 45-bits	SubnetID	Interface ID (host part)

- Prefix 2000::/3 (address starts with bits 001)
- 45-bits: global routing prefix (IANA->RIR->LIR/ISP)
- 16-bits Subnet ID -- can number 65,536 subnets!
- 64-bits Interface ID

Multicast addresses

- Multicast: an efficient one-to-many form of communication
- A special IPv6 address prefix, **`ff00::/8`**, identifies multicast group addresses
- Hosts that wish to receive multicast traffic "join" the associated multicast group
- Have scopes (link local, site, global etc)
- In IPv4, the group joining and leaving protocol is IGMP
- In IPv6, the protocol is **MLD** (Multicast Listener Discovery)

Address Configuration

- Servers: usually have statically configured IPv6 addresses (and associated DNS records)
- Client computers: can automatically configure themselves an address (“Stateless Address Autoconfiguration”)
  - typically don’t have associated DNS records
- Managed address allocation can be performed with DHCPv6 (Dynamic Host Configuration Protocol for IPv6)
  - DNS can be pre-populated for DHCPv6 address pools

IPv6 Subnets

- Usually fixed size: 64-bits long (p2p links are often exceptions)
- First 4 fields defined the network portion of the address
- How many hosts can such a subnet accommodate?

264 =** 18,446,744,073,709,551,616 (or approx 18.5 quintillion)

```
eg. for a subnet: 2001:db8:ab23:74::/64

start :  2001:db8:ab23:74:0000:0000:0000:0000
end   :  2001:db8:ab23:74:ffff:ffff:ffff:ffff
```

IPv6 Subnets

- IPv6 Addressing Architecture (RFC 4291) requires the host portion of the address (or the "Interface Identifier") to be 64-bits long
- To accommodate a method that allows hosts to uniquely construct that portion: Modified EUI-64 format
- Generates unique 64-bit identifier from MAC address
- This is used by Stateless Address Autoconfiguration (to be described shortly)

Neighbor Discovery

- RFC 4861
- Analog of ARP in IPv4 but provides many other capabilities
- Stateless Address Autoconfiguration (RFC 4862)
- Managed configuration indication (address configuration policy)
- Router discovery
- Subnet Prefix discovery
- Duplicate address detection (DAD)
- Neighbor unreachability detection (NUD)

Neighbor discovery messages

- Uses 5 ICMPv6 message types:

- Router Solicitation
- Router Advertisement
- Neighbor Solicitation (like ARP Request)
- Neighbor Advertisement (like ARP Response)
- Redirect

[RFC 4443: ICMPv6 Specification]

Solicited node multicast

- Neighbor discovery involves finding other hosts & routers on the local subnet, but recall there is no broadcast in IPv6
- ND uses solicited node multicast addresses, which partition hosts on a subnet into distinct groups, each corresponding to a distinct multicast addresses associated with sets of IPv6 addresses
- For every IPv6 address a host has, it joins the corresponding solicited node multicast address
- Address contains last 24 bits of the IPv6 address
- First 104 bits are the well defined prefix
  - `ff02:0:0:0:0:1:ff00::/104`

Solicited node multicast

- If target address is: 2001:db8:123::ce97:7fce
- Last 24 bits are: 97:7f:ce. Prepend ff02::1:ff00:/104
- So solicited node multicast address is: ff02::1:ff97:7fce
- If Ethernet is the link layer, the corresponding ethernet multicast address: 33-33 + last 32bits of the solicited node address = 33-33-ff-97-7f-ce

- Main takeaway: In IPv6, neighbor discovery involves host sending packet to the solicited node multicast address associated with the target (in contrast to IPv4’s ARP, where we send to the broadcast address)

Modified EUI-64

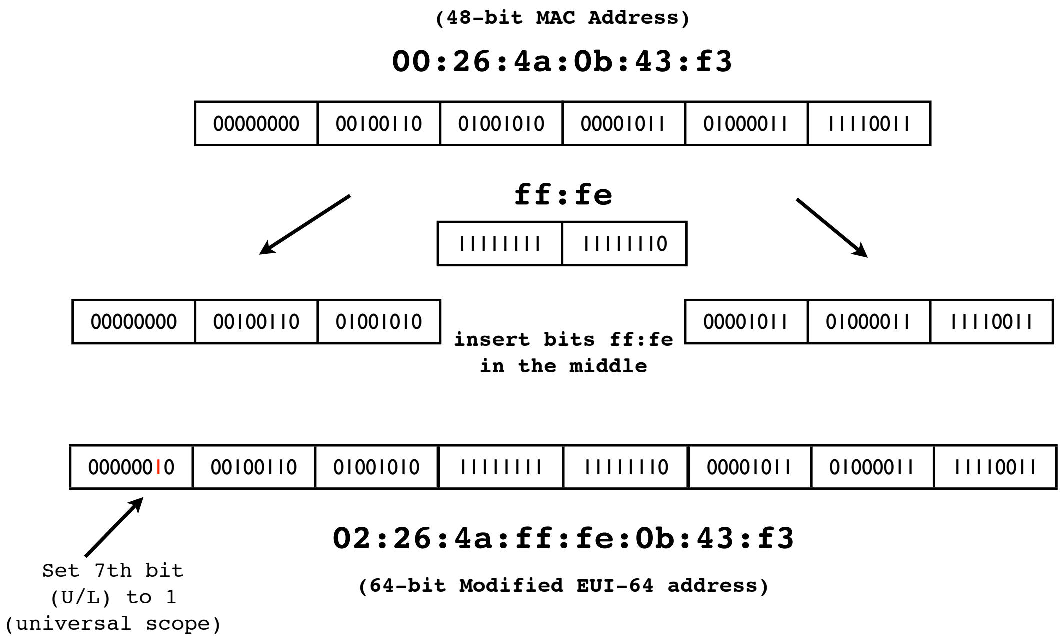

(48-bit MAC Address)

00:26:4a:0b:43:f3

00000000	00100110	01001010	00001011	01000011	11110011

ff:fe

11111111	11111110

00000000	00100110	01001010

insert bits ff:fe in the middle

00001011	01000011	11110011

00000010	00100110	01001010	11111111	11111110	00001011	01000011	11110011

Set 7th bit (U/L) to 1 (universal scope)

02:26:4a:ff:fe:0b:43:f3

(64-bit Modified EUI-64 address)

(details: see RFC 3513, Appendix A)

Autoconfiguration

- RFC 4862: Stateless Address Autoconfiguration (SLAAC)
- Host listens to Router Advertisements (RA) on local subnet
- Obtains 64-bit subnet prefix from RA (and perhaps other parameters)
- Computes modified EUI-64 from its MAC address and concatenates it to 64-bit subnet prefix to form IPv6 address

```
Link prefix from RA: 2001:db8:abcd:1234::/64
Host MAC address: 00:1b:63:94:9d:73
EUI-64 address:   021b:63ff:fe94:9d73
Resulting IPv6 address:
   2001:db8:abcd:1234:021b:63ff:fe94:9d73
```

Router Advertisement

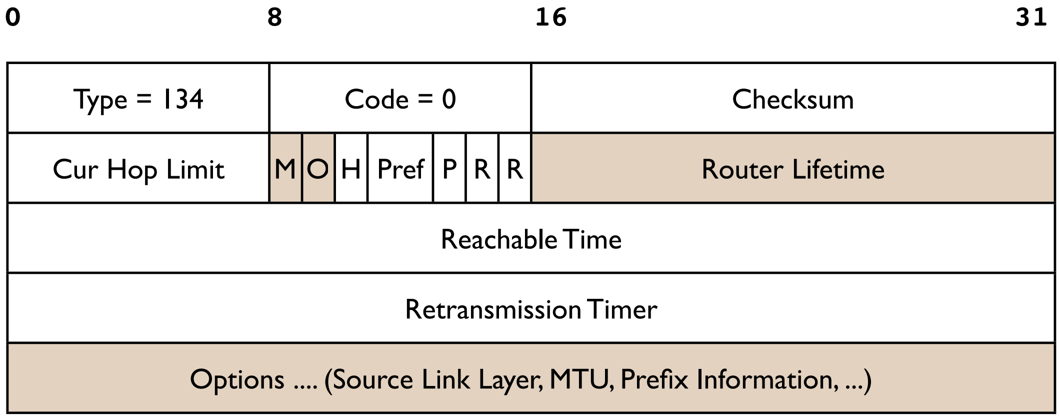

<table>
<tr><th>0</th><th colspan="7">8</th><th>16 … 31</th></tr>
<tr><td>Type = 134</td><td colspan="7">Code = 0</td><td>Checksum</td></tr>
<tr><td>Cur Hop Limit</td><td>M</td><td>O</td><td>H</td><td>Pref</td><td>P</td><td>R</td><td>R</td><td>Router Lifetime</td></tr>
<tr><td colspan="9">Reachable Time</td></tr>
<tr><td colspan="9">Retransmission Timer</td></tr>
<tr><td colspan="9">Options .... (Source Link Layer, MTU, Prefix Information, ...)</td></tr>
</table>

```
M = managed config flag: "use stateful DHCPv6"
O = other config flag: get "other" params via stateless DHCPv6
Pref = Default Router Preference (Hi, Lo, Med) - RFC 4191
The most common option is the Prefix Information option
```

RA: Prefix Info option

0	8	16			31
Type=3	Length	Prefix Length	L	A	Reserved1
Valid Lifetime					
Preferred Lifetime					
Reserved2					
Prefix					

L = "on link" prefix indicator
A = this prefix can be used for auto-configuration

Router Discovery

```
Router Solicitation Message ->
Src: fe80::c072:7a5f:c1b5:24d1
Dst: ff02::2 (all routers multicast)
ICMPv6 Type 133 (RS)
Option:
  Src Link Layer Addr (my MAC addr)
```

```
<- Router Advertisement Message
Src: router's link local addr
Dst: ff02::1 (all nodes or solicitor)
ICMPv6 Type 134 (RA)
Flags (M=0, O=0, pref=0)
Router Lifetime: 1800
Reachable time: 0
Retrans time: 0
Options:
  Src Link Layer Addr (my Mac)
  MTU: 1500
  Prefix Info
    prefix: 2001:db8:ab:cd::/64
    valid life: 2592000
    preferred lifetime: 604800
```

(Routers also periodically send out unsolicited router advertisements.)

Neighbor Discovery

```
Neighbor Solicitation Message ->
Src: A's IPv6 address
Dst: Solicited-node multicast of B
ICMPv6 Type 135 (NS)
Target: B's IPv6 address
Options:
  Src Link Layer Addr (A's MAC addr)
```

```
<- Neighbor Advertisement Message
Src: B's IPv6 address
Dst: A's IPv6 address
ICMPv6 Type 135 (NA)
Target: B's IPv6 address
Options:
  Src Link Layer Addr (B's MAC addr)
```

(Summary: A is asking: what is the link layer address associated with B's IPv6 address?)

SLAAC & Privacy?

- SLAAC exposes MAC address of a host to the world
- In IPv4, MAC was exposed to local subnet only
- Does this have privacy implications?
- Remote sites may be able to track & correlate your network activities by examining a constant portion of your address
- How serious are these compared to other highly privacy invasive mechanisms already in use at higher layers?
  - think of things like web cookies that track/expose user identity, often across sites

Temporary addresses

- RFC 4941: Privacy extensions for Stateless Address Auto-configuration
- Pool of "Temporary addresses" or "Privacy addresses"
- Derived from MAC initially, ala SLAAC, but then passed through a 1-way hash algorithm
- Designed to change over time; duration configurable or based on policy; hours, days, on reboot, or different addresses for different applications or endpoints
- Cons: complicate network debugging, security/audit implications (see proposal for "*stable privacy addresses*")

Temporary addresses

- On by default in Windows (since XP), Mac OS X 10.7, Open Suse 12.1, Ubuntu Linux 12.04, ..
- Also on in Apple iOS devices (iPhone, iPad etc)
- Android 4.0 uses and prefers privacy addresses (on wifi)
- Off by default in others, but easily turned on via configuration knobs in the operating system (eg. sysctl on Linux and *BSD)

DHCPv6

- Stateless DHCPv6 (RFC 3736)
  - No IPv6 address assignment (“stateless”); assumed that SLAAC or other method will be used for address configuration
  - Other network configuration parameters are provided, eg. DNS servers, NTP servers etc
- Stateful DHCPv6 (RFC 3315)
  - Managed address allocation analogous to DHCP in IPv4
  - Easy to populate DNS & reverse DNS (compared to autoconfig)

Stateful DHCPv6

- Stateful DHCPv6 (RFC 3315) - more details
- Conceptually similar to IPv4 DHCP
- Uses RA’s **M** (managed configuration) flag
- Requires DHCPv6 server, which assigns IPv6 leases to clients
- And provides other configuration info (DNS, NTP, ... etc)

Differences with IPv4 DHCP

- Uses UDP ports 546 (server) and 547 (client)
- Clients use autoconfigured link-local addresses as source
- Clients send messages to multicast group address ff02::1:2 (“all dhcp servers and relay agents group”); IPv4 uses broadcast
- **Does not assign default gateway** - use Router Advertisement
- DHCP servers can send “reconfigure” messages to clients
- Rapid Commit option (reduce exchange from 4 to 2 messages)
- DUID (DHCP Unique IDentifiers)
- Provision for temporary and non-temporary addresses

IPv4 v IPv6 DHCP messages

DHCP v4 (rfc 2131)	DHCP v6 (rfc 3315)
`C -> broadcast: DISCOVER`	`C -> multicast: SOLICIT`
`S -> C: OFFER`	`S -> C: ADVERTISE`
`C -> S: REQUEST`	`C -> S: REQUEST`
`S -> C: ACK`	`S -> C: REPLY`

IPv4 v IPv6 DHCP messages

with rapid commit option

DHCP v4 (rfc 2131)	DHCP v6 (rfc 3315)
`C -> broadcast: DISCOVER`	`C -> multicast: SOLICIT`
`S -> C: OFFER`	`S -> C: REPLY`
`C -> S: REQUEST`	
`S -> C: ACK`	

DHCPv6 DUID

- Clients no longer use hardware address to identify themselves
  - Issues: multiple interfaces, mobility, virtual interfaces & VMs etc
  - DUID: DHCP Unique IDentifier - use long lived unique id instead
  - Used by both clients and servers
  - Number of methods to initialize a DUID (based on link layer address, time, enterprise numbers etc): DUID-LLT/ET/LT

DHCPv6 DUID

- DUID construction methods:
  - DUID-LLT: constructed from link-layer address of one of the system interfaces (ie. from hardware address), hardware type, and timestamp
  - DUID-EN: using enterprise number of device manufacturer and an ID number
  - DUID-LL: constructed from link-layer address and hardware type
- Challenges with DUIDs:
  - when we want to obtain MACs; correlating IPv4/IPv6 addresses; persistent storage on some devices, etc

DHCPv6 Leases & Lifetimes

- Leases (bindings) as in IPv4
- Lifetimes: Offered addresses have preferred and Valid lifetimes as in stateless autoconfiguration

Stateless DHCPv6

- Triggered by “O (other config) flag” in RA messages
- INFORMATION_REQUEST message:
- To request other configuration parameters
  - C -> multicast: INFORMATION_REQUEST
  - S -> C: REPLY
- Conceptually similar to the DHCPINFORM message in IPv4

DHCPv6 options

- Used by both stateful and stateless DHCPv6
- Some common options for configuration information:
  - DNS Recursive Nameservers
  - DNS Search List
  - NTP servers
  - SIP servers
  - Prefix Delegation (RFC 3633) - eg. delegating prefix to a home router
  - and many more ...

DHCPv6 Other

- Other messages: RENEW, REBIND, CONFIRM, RELEASE, DECLINE, RECONFIGURE
- Relay Agents supported as in IPv4 (RELAY_FORW, RELAY_REPL)
- ServerFailover protocol
  - So far missing in v6, but development work in progress.
  - Less important for IPv6 (use multiple independent servers offering disjoint address pools), but there are some uses cases.
- Prefix delegation
- Client Link Layer Address Option (coming)
  - http://tools.ietf.org/html/draft-ietf-dhc-dhcpv6-client-link-layer-addr-opt-05

DHCPv6 with Relay Agent

Client	Relay	Server
`->Solicit`		
	`->Relay-forw(Solicit)`	
		`<-Relay-repl(Advertise)`
	`<- Advertise`	
`->Request`		
	`->Relay-forw(Request)`	
		`<-Relay-repl(Reply)`
	`<-Reply`	

Other config possibilities

- New Router Advertisement options
  - RFC 6106: RA options for DNS configuration
  - Allows transmission of DNS server and related info via RA (obviating need to deliver this via other means, eg. stateless DHCPv6)
  - Very few implementations to date though ..
- In the opposing camp, there is (was?) also a proposal to extend DHCPv6 to provide default gateway options, obviating the need to use Router Advertisements
  - http://tools.ietf.org/html/draft-droms-dhc-dhcpv6-default-router-00
  - http://tools.ietf.org/html/draft-ietf-mif-dhcpv6-route-option-05

Unique Local Address (ULA)

- RFC 4193, Prefix fc00::/7
- Replacement for IPv4 Private Addresses (RFC 1918)
- Note: the older *site local prefix* (fec0::/10) was deprecated
- Intended for local use within a site or group of sites
- Globally unique, but not routable on the global Internet
- Addresses some operational issues seen with IPv4 and RFC 1918 addresses

Unique Local Address (ULA)

```
    | 7 bits |1|  40 bits   |  16 bits  |          64 bits           |
    +--------+-+------------+-----------+----------------------------+
    | Prefix |L| Global ID  | Subnet ID |        Interface ID        |
    +--------+-+------------+-----------+----------------------------+

 Where:

    Prefix            FC00::/7 prefix to identify Local IPv6 unicast
                      addresses.

    L                 Set to 1 if the prefix is locally assigned.
                      Set to 0 may be defined in the future.  See
                      Section 3.2 for additional information.

    Global ID         40-bit global identifier used to create a
                      globally unique prefix.

    Subnet ID         16-bit Subnet ID is an identifier of a subnet
                      within the site.
```

[RFC 4193 excerpt]

IPv6 addresses on a Mac

```
$ ifconfig -a

lo0: flags=8049<UP,LOOPBACK,RUNNING,MULTICAST> mtu 16384
     options=3<RXCSUM,TXCSUM>
     inet6 fe80::1%lo0 prefixlen 64 scopeid 0x1
     inet 127.0.0.1 netmask 0xff000000
     inet6 ::1 prefixlen 128
en1: flags=8863<UP,BROADCAST,SMART,RUNNING,SIMPLEX,MULTICAST> mtu 1500
     ether e4:ce:8f:07:b6:13
     inet6 fe80::e6ce:8fff:fe07:b613%en1 prefixlen 64 scopeid 0x5
     inet6 2607:f470:6:3:e6ce:8fff:fe07:b613 prefixlen 64 autoconf
     inet6 2607:f470:6:3:3947:98a5:68f6:2ef1 prefixlen 64 autoconf temporary
     inet 165.123.70.49 netmask 0xffffff00 broadcast 165.123.70.255
     media: autoselect
     status: active
```

IPv6 addresses on Windows

```
C:>ipconfig

Windows IP Configuration

[...]

Ethernet adapter Local Area Connection:

   Connection-specific DNS Suffix  . :

   IPv6 Address. . . . . . . . . . . : 2607:f470:2f:1:2dde:6914:cafe:15fe
   Temporary IPv6 Address. . . . . . : 2607:f470:2f:1:806c:86ee:b372:47b2
   Link-local IPv6 Address . . . . . : fe80::2dde:6914:cafe:15fe%10
   IPv4 Address. . . . . . . . . . . : 128.91.196.91
   Subnet Mask . . . . . . . . . . . : 255.255.254.0
   Default Gateway . . . . . . . . . : fe80::216:9cff:fe6f:5dc0%10
                                       128.91.196.1
```

IPv6 addresses on Linux

```
$ ifconfig

lo        Link encap:Local Loopback
          inet addr:127.0.0.1  Mask:255.0.0.0
          inet6 addr: ::1/128 Scope:Host
          UP LOOPBACK RUNNING  MTU:16436  Metric:1
          RX packets:544285 errors:0 dropped:0 overruns:0 frame:0
          TX packets:544285 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:0
          RX bytes:355551886 (339.0 MiB)  TX bytes:355551886 (339.0 MiB)

eth0      Link encap:Ethernet  HWaddr 00:14:4F:01:31:F8
          inet addr:128.91.XXX.68  Bcast:128.91.255.255  Mask:255.255.254.0
          inet6 addr: 2607:f470:2a:1::a:2/64 Scope:Global
          inet6 addr: 2607:f470:2a:1::a:1/64 Scope:Global
          inet6 addr: 2607:f470:2a:1:214:4fff:fe01:34f7/64 Scope:Global
          inet6 addr: fe80::214:4fff:fe01:34f7/64 Scope:Link
          UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
          RX packets:9228907 errors:0 dropped:0 overruns:0 frame:0
          TX packets:3889095 errors:0 dropped:0 overruns:0 carrier:0
          collisions:0 txqueuelen:1000
          RX bytes:1686780678 (1.5 GiB)  TX bytes:1997866418 (1.8 GiB)
```

Linux RA example

Example of RA info seen on a Linux machine. This host has a static address, and 2 autoconfigured addresses, one deprecated because its preferred lifetime has expired.

```
$ /sbin/ip -6 addr show dev eth0

eth0: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1500 qlen 1000
    inet6 2607:f470:1001::1:12/64 scope global
       valid_lft forever preferred_lft forever
    inet6 2607:f470:1001:0:214:4fff:fee6:b650/64 scope global
dynamic
       valid_lft 2591957sec preferred_lft 604757sec
    inet6 2001:468:1802:101:214:4fff:fee6:b650/64 scope global
deprecated dynamic
       valid_lft 6308sec preferred_lft -892sec
    inet6 fe80::214:4fff:fee6:b650/64 scope link
       valid_lft forever preferred_lft forever
```

Common IPv6 assignments

< /32	RIRs and large ISPs
/32	Typically to LIRs and ISPs. Allows 65,536 /48 assignments, or 4 billion /64 subnets
/48	Most enterprises and endsites. Allows deployment of 65,536 /64 subnets
/56	Small sites; Residential service. Allows deployment of 256 /64 subnets
/64	Residential service. Allows one /64 subnet

- See RFC 6177 for latest thinking on endsite assignments

NRO

IPv6 ADDRESS SPACE

How much has been allocated to the RIRs?

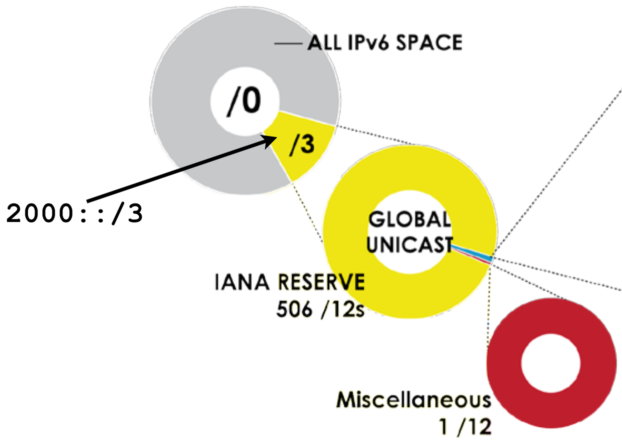

RIRs 5 /12s (October 2006)

RIR	IPv6 ADDRESS
AfriNIC	2C00:0000::/12
APNIC	2400:0000::/12
ARIN	2600:0000::/12
LACNIC	2800:0000::/12
RIPE NCC	2A00:0000::/12

Number Resource Organization

September 2011 Internet Number Resource Report

PA vs PI address space

- Provider Assigned (PA)
  - Usually assigned by your ISP, and suballocated by the ISP from a larger block of addresses the ISP has
  - ISP aggregates the announcement upstream
  - Customer usually obtains one PA block from each ISP
- Provider Independent (PI)
  - Sometimes called "Portable" address space
  - Not aggregated by upstream ISPs/Peers and appears as a distinct prefix in the global Internet routing table (***scalability issues***!)
  - Needed for multihoming (pending a better scalable solution)

Provider Aggregation eg.

A real example ...

`2001:468::/32`	`Internet2: PI block`
`2001:468:1800::/40`	`MAGPI GigaPop: PA block`
`2001:468:1802::/48`	`University of Pennsylvania: PA block`

Internet2 suballocates the /40 block from its own PI block to MAGPI (a regional ISP), and MAGPI suballocates a /48 from that to its downstream connector UPenn. Internet2 only sends the aggregate /32 announcement to its peers (other large ISPs and organizations), and only that /32 prefix is seen in the global Internet2 routing table.

Provider Aggregation eg.

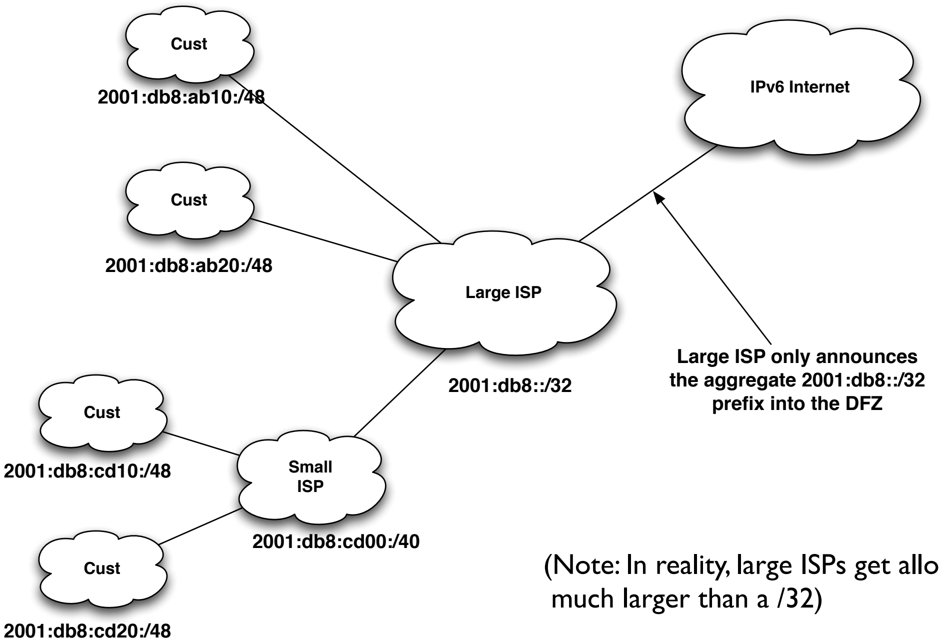

Large ISP only announces the aggregate 2001:db8::/32 prefix into the DFZ

(Note: In reality, large ISPs get allocations much larger than a /32)

Multihoming

- Not fully solved; an area of active research & protocol design
- Initial model: everything is provider assigned and aggregatable
- Now **PI** (Provider Independent) address space common
- Future possibilities:
  - **SHIM6** - RFC 5533, 5534, 5535
  - **LISP** - Locator/Identifier Separation Protocol - see IETF wg
  - IRTF routing research group
    - **RFC 6115**: Recommendation for a Routing Architecture
    - **ILNP**: Identifier-Locator Network Protocol (RFC 6740-6748)

IPv6 support in network service providers

ISPs offering IPv6

- Some: NTT/Verio, Global Crossing, Level 3, Cogent, Cable & Wireless, Reliance, Tata Communications, TeliaSonera, Hurricane Electric, Comcast Business, ... (growing list)
- http://www.sixxs.net/faq/connectivity/?faq=ipv6transit
- http://en.wikipedia.org/wiki/Comparison_of_IPv6_support_by_major_transit_providers
- Mixture of native and tunneled IPv6 service
- If you're a US edu, you might be able get IPv6 connectivity via the Internet2 R&E network
  - Equivalent opportunities with other national or continental RENs (JANET, SURFNet, GEANT, APAN etc)

Mobile Carriers

- An increasing number of mobile carriers now support IPv6 (e.g. on their 4G/LTE networks)
  - Verizon LTE, AT&T, T-Mobile
  - May need specific type of phones/hardware to take advantage of this though

T-Mobile deploys IPv6

January 2012

http://www.androidpolice.com/2012/01/29/t-mobile-usa-testing-ipv6-on-select-devices-here-is-what-it-all-means-and-yes-no-more-nat/

"T-Mobile USA is running an open beta for enabling IPv6 address assignment to some devices on its network in place of the traditional IPv4 addresses."

April 2012

http://www.extremetech.com/mobile/127213-ipv6-now-deployed-across-entire-t-mobile-us-network

“T-Mobile has completed the deployment of IPv6 services across its entire network. This isn’t the first IPv6 network, but it is the largest wireless IPv6 deployment in the world.”

https://sites.google.com/site/tmoipv6/lg-mytouch

Residential Service

- Not as encouraging, but ...
- Comcast is leading the charge
  - http://www.comcast6.net/
  - 50% of broadband footprint enabled; 2.5% customers are using IPv6 now; commercial metro ethernet IPv6 enabled too
- Who else?
  - AT&T some; Time Warner has limited trials
  - Verizon FIOS has no announced plans yet

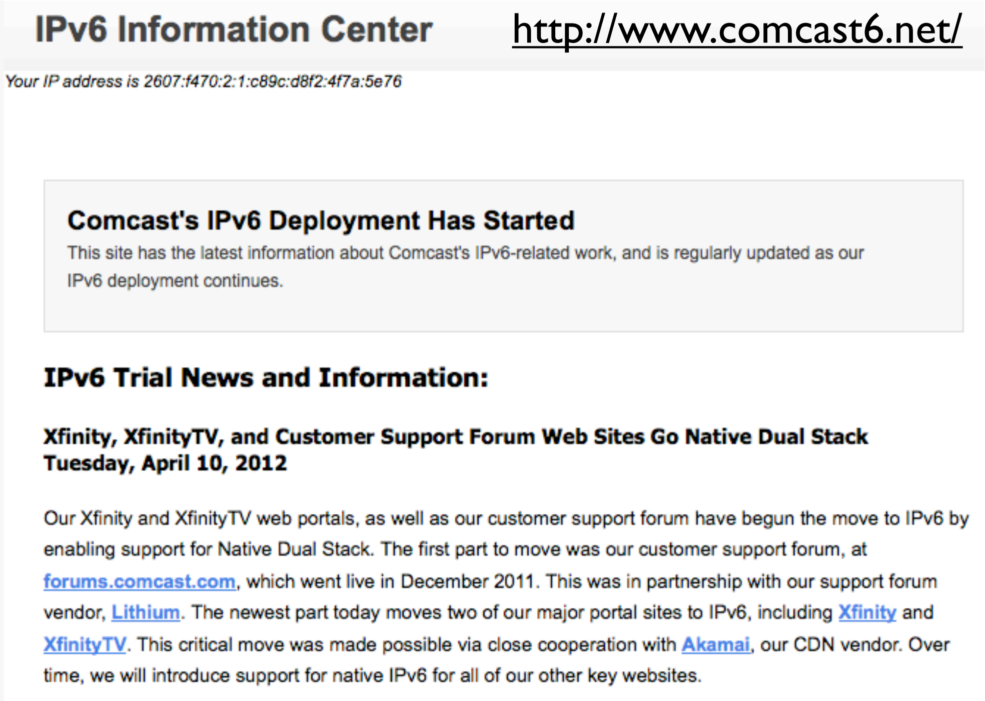

IPv6 Information Center http://www.comcast6.net/

Your IP address is 2607:f470:2:1:c89c:d8f2:4f7a:5e76

Comcast's IPv6 Deployment Has Started

This site has the latest information about Comcast's IPv6-related work, and is regularly updated as our IPv6 deployment continues.

IPv6 Trial News and Information:

Xfinity, XfinityTV, and Customer Support Forum Web Sites Go Native Dual Stack
Tuesday, April 10, 2012

Our Xfinity and XfinityTV web portals, as well as our customer support forum have begun the move to IPv6 by enabling support for Native Dual Stack. The first part to move was our customer support forum, at forums.comcast.com, which went live in December 2011. This was in partnership with our support forum vendor, Lithium. The newest part today moves two of our major portal sites to IPv6, including Xfinity and XfinityTV. This critical move was made possible via close cooperation with Akamai, our CDN vendor. Over time, we will introduce support for native IPv6 for all of our other key websites.

Also see http://www.cio.com/article/704136/Comcast_is_First_U.S._ISP_to_Offer_IPv6_to_Home_Gateway_Users

Time Warner Cable IPv6 trials

Date: September 27, 2011 8:35:42 AM CDT

To: "nanog@nanog.org" <nanog@nanog.org>

Subject: Volunteers needed for TWC IPv6 trial

Time Warner Cable is expanding our residential IPv6 trials in severalmarkets, and we need more people. If you're a Time Warner Cable High Speed Internet subscriber, and are interested in participating in our IPv6 trials, please let us know! We have a short form at

http://www.timewarnercable.com/Corporate/support/IPv6_volunteerform.html

that will help us find the right mix of people, equipment, and locations, toget the most out of our trials.

Thanks in advance for participating!

Time Warner Cable

Also see: http://www.macobserver.com/tmo/article/time_warner_cable_talks_about_ipv6_launch/

Content Delivery Networks

- **LimeLight Networks** supports IPv6
  - claims to be first IPv6 CDN
- **Akamai** announced production IPv6 support in April 2012
  - http://www.akamai.com/ipv6
- **Cloudflare** and **Edgecast** too
- See ISOC's deploy360 page for more:
  - http://www.internetsociety.org/deploy360/resources/ipv6-and-content-delivery-networks-cdns/

Other Cloud

- Amazon Web Services
  - http://aws.amazon.com/about-aws/whats-new/2011/05/24/elb-ipv6-zoneapex-securitygroups/
  - No internal infra support, but dualstack on outside facing possible
- Other Cloud providers: Cloudflare, Softlayer, Arpnetworks, Linode

IPv6 Support in Operating Systems & Applications

Operating System Support

- Most modern operating system support IPv6 out of the box
- Microsoft Windows, Apple Mac OS X, Linux, *BSD, Solaris, Tru64 UNIX, IBM AIX, etc
- Mobile OSes like iOS, Android do also
- They generally use autoconfiguration or DHCPv6 to configure IPv6 addresses
- For servers, it’s advisable to configure static addresses

Application Services

- Recall: IPv6 is not backwards compatible with IPv4
- Applications need to be modified to support IPv6
- Many open source & commercial applications already do!
- Don't forget to consider home grown, and locally developed applications also!

IPv6 ready applications

- Webservers: Apache, IIS
- E-mail: Sendmail, Postfix, UW IMAP, Cyrus, MS Exchange, Exim, Qmail, Dovecot, Courier
- DNS: BIND, NSD, PowerDNS, Microsoft DNS
- LDAP: OpenLDAP, Active Directory
- Kerberos: MIT, Heimdal, Active Directory
- More comprehensive lists:
  - http://www.ipv6-to-standard.org/
  - http://www.deepspace6.net/docs/ipv6_status_page_apps.html

IPv6 client software

- Browsers: Firefox, Internet Explorer, Safari, Chrome, Opera
- E-mail: Apple Mail, Thunderbird, MS Outlook
- [others to be added ...]

IPv6 Tunneling

Automatic Tunneling

- Even without IPv6 deployed in your network, computers may be using IPv6
- Via automatic tunneling mechanisms. Two popular ones are **6to4** and **Teredo**
- These work by **encapsulating** IPv6 packets inside IPv4 packets and sending them to a relay router that is connected to both the IPv4 and IPv6 Internet
- **Tunnels sometimes cause connectivity and performance problems. Native IPv6 deployment usually fixes all of them**

6to4

- A transition method for IPv6 capable hosts or networks that don’t have native IPv6 network connectivity to use tunneling to communicate with other IPv6 islands and/or the IPv6 Internet
- Does not involve explicit setup of the tunnels.
- 6to4 hosts and networks are numbered in the **2002::/16** prefix
- **6to4 routers** sit at the edge of an IPv6 site and the IPv4 Internet
- The most common deployment model of 6to4 involves using 6to4 anycast addresses to reach **6to4 relay routers**
  - **192.88.99.1** and **2002:c058:6301::**

6to4

- Site constructs a /48 IPv6 prefix by concatenating 6to4 router's IPv4 address to 2002::/16, and tunnels IPv6 packets from the 6to4 router to a 6to4 relay router that is connected to both the IPv4 and IPv6 Internet.
- A site could be a single computer, in which case it is itself the 6to4 router

```
References:
RFC 3056: Connection of IPv6 domains via IPv4 clouds
RFC 3068: An anycast prefix for 6to4 relay routers
RFC 6343: Advisory Guidelines for 6to4 deployment
```

6to4 Diagram

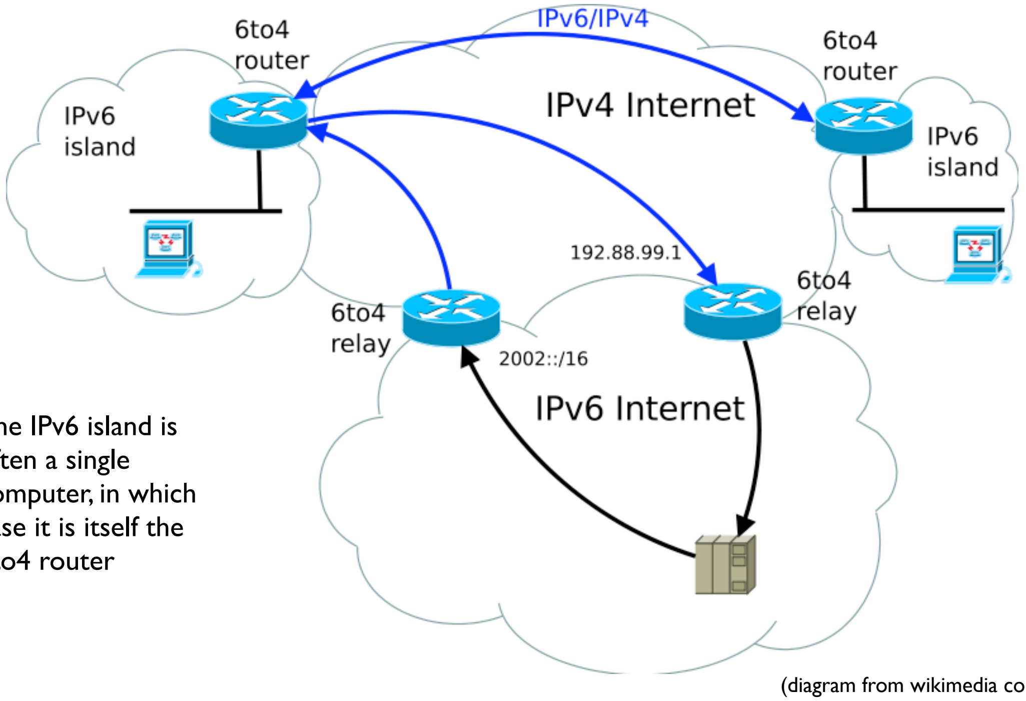

The IPv6 island is often a single computer, in which case it is itself the 6to4 router

(diagram from wikimedia commons)

6to4 Issues

- 6to4 can fail or perform poorly due to a variety of reasons:
  - Inbound/outbound blackholes (routers or firewalls filtering protocol 41, ICMP etc)
  - Lack of working return 6to4 relay
  - Circuitous/Asymmetric path with large round trip time
  - PMTU failures due to encapsulation overhead etc
- Privacy concerns with 3rd party relay routers
- See RFC 6343: Advisory Guidelines for 6to4 Deployment

Teredo

- Encapsulates IPv6 in UDP in IPv4 (see RFC 4380 for details)
- Works through NATs
- Special IPv6 prefix 2001::/32 (ie. 2001:0000::/32)
- Uses Teredo Servers and Teredo Relays

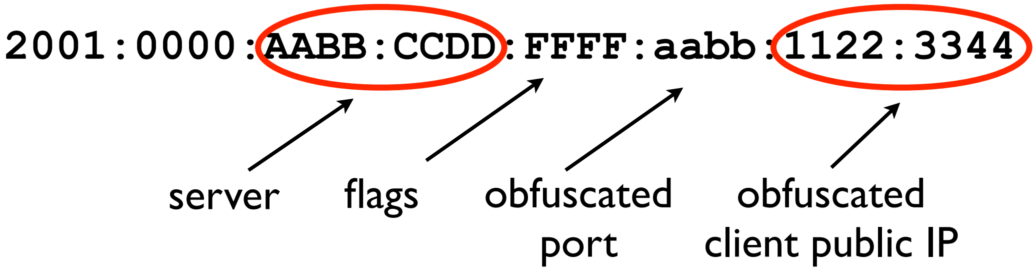

```
2001:0000:AABB:CCDD:FFFF:aabb:1122:3344
```

server

flags

obfuscated port

obfuscated client public IP

Teredo

- **Teredo Servers** are used for initialization, testing type of NAT, determining client's externally routable address, and for periodically maintaining state in NATs/firewalls
- **Teredo Relays** are used for relaying tunneled traffic to and from the IPv6 Internet

Teredo Diagram

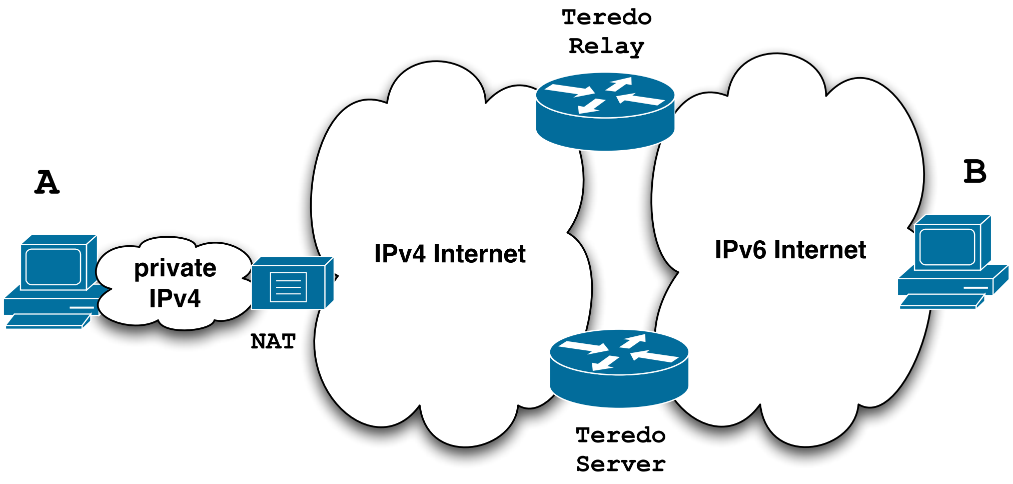

Teredo Issues

- Cannot work through some types of NAT (eg. Symmetric)
- NAT detection and traversal mechanisms employed have a significant impact on network performance
- Possible issues with inoperable Teredo servers and relays
- Privacy concerns with 3rd party servers and relays
- Security concerns have been expressed:
  - http://tools.ietf.org/html/draft-ietf-v6ops-teredo-security-concerns-02

Identifying tunneled traffic

- 6to4 uses well known prefix 2002::/16
- Teredo uses 2001::/32
- Both use value 41 (IPv6 encapsulation) in the IPv4 protocol field
- 6to4 encapsulates IPv6 packets directly in IPv4
- Teredo is encapsulated in UDP inside IPv4
- 6to4 commonly uses well-known anycast relay routers (192.88.99.0/24)
- There are also public Teredo servers and relays
- *Note: blindly blocking tunneled traffic may cause more harm than good*

Managed tunnels

- Statically configured, managed, IPv6 in IPv4 tunnels usually provide more predictable and more reliable service. A few managed tunnel providers

- Hurricane Electric: www.tunnelbroker.net
- Freenet6: www.hexago.com
- Consulintel: tb.consulintel.euro6ix.org
- Sixxs: www.sixxs.net

IPv6 at Penn

IPv6 at Penn

- 2002: Deployment in the MAGPI GigaPoP
- 2005: Initial deployment in PennNet
  - External connection & Core network infra.
  - Selected departmental subnets and server networks
- 2007: School of Engineering & Applied Science
- 2009: Annenberg School (for videoconferencing w/ China)
- 2011: Larger deployment throughout wired network
- 2012: Initial deployment in AirPennNet (but backed out)

IPv6 on wireless?

- 2012 summer: Initial deployment in AirPennNet
- Backed out at end of summer because of a specific deficiency
- IP mobility feature (wasn't implemented for IPv6 by our wireless equipment vendor, Aruba Networks; also broke IPv4 mobility)
- This summer: testing new IPv6 mobility feature for Aruba

IPv6 at Penn

- DNS
- NTP
- SSH (server administration)
- Jabber (XMPP)
- Assignments
- Penn Web server
- Departmental deployments (SEAS)

Assignments

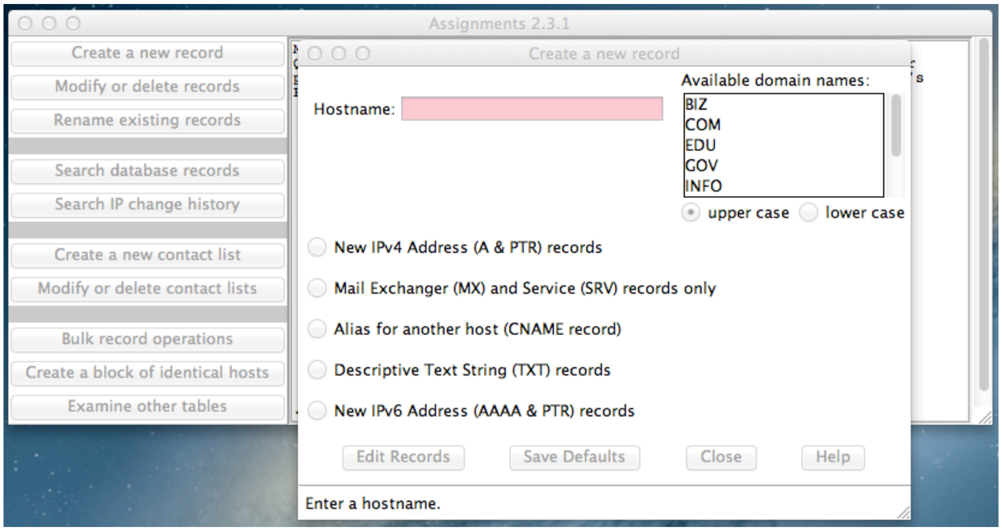

Hostname:

Available domain names:
BIZ
COM
EDU
GOV
INFO

● upper case ○ lower case

○ New IPv4 Address (A & PTR) records
○ Mail Exchanger (MX) and Service (SRV) records only
○ Alias for another host (CNAME record)
○ Descriptive Text String (TXT) records
○ New IPv6 Address (AAAA & PTR) records

Enter a hostname.

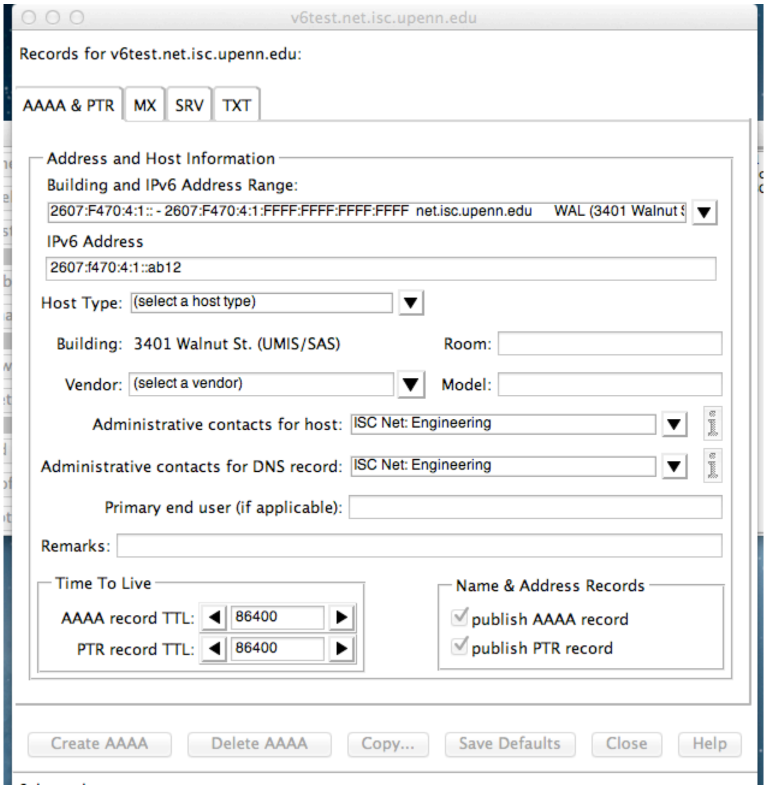

Records for v6test.net.isc.upenn.edu:

Address and Host Information

Building and IPv6 Address Range: 2607:F470:4:1:: - 2607:F470:4:1:FFFF:FFFF:FFFF:FFFF net.isc.upenn.edu WAL (3401 Walnut S

IPv6 Address: 2607:f470:4:1::ab12

Host Type: (select a host type)

Building: 3401 Walnut St. (UMIS/SAS)

Room:

Vendor: (select a vendor)

Model:

Administrative contacts for host: ISC Net: Engineering

Administrative contacts for DNS record: ISC Net: Engineering

Primary end user (if applicable):

Remarks:

Time To Live

AAAA record TTL: 86400

PTR record TTL: 86400

Name & Address Records

☑ publish AAAA record

☑ publish PTR record

Demo

Demos of using IPv6 enabled applications!

Address Selection

DualStack Address Selection

- I'm a dual stack (IPv4/IPv6) client
- I lookup "www.example.com" eg. using **getaddrinfo**()
  - Performs both A and AAAA DNS queries and may return a list of various IPv4 and IPv6 addresses
  - Which should I try connecting to? In what order?

DualStack Address Selection

- RFC 6724: Default Address Selection Algorithm
  - Updated from the original RFC 3484
- Many rules, but one effect is to generally prefer IPv6 over IPv4

Type	Prefix	Precedence	Label
Loopback	::1/128	50	0
IPv6	::/0	40	1
IPv4	::ffff:0:0/96	35	4
6to4	2001::/16	30	2
Teredo	2001::/32	5	5
ULA	fc00::/7	3	13
Site Local	fec0::/10	1	11
6Bone	3ffe::/16	1	12

Happy Eyeballs

- RFC 6555, 6556: Happy Eyeballs: Success with Dual Stack Hosts
  - Parallel connections to v4 & v6 destinations, but give v6 a small headstart or pref. Use first connection that succeeds & cache results; tunable knobs
- Apple Mac OS X Lion:
  - Not quite Happy Eyeballs: no preference for IPv6 over IPv4; use what seems to work best, leading to more non-deterministic behavior
- Windows: http://blogs.msdn.com/b/b8/archive/2012/06/05/with-ipv6-in-windows-8.aspx
- Survey of what various OS and apps used to do/currently do (G. Huston, RIPE64): https://ripe64.ripe.net/presentations/78-2012-04-16-ripe64.pdf
- Traditional resolver vs "Connect-by-Name" APIs

Migration strategies for IPv6 services

- DualStack migration is the ideal, but has possible issues if broken IPv6 client connectivity is widespread
- An overview of some alternate strategies given here:
  - RFC 6589: Considerations for Transitioning content to IPv6
  - DNS Resolver Whitelisting; Resolver Blacklisting; IPv6 specific service names, etc

IPv6 and Security

IPv6 Security issues

- IPsec myth (IPv6 is automatically more secure because of IPsec)
- Code and implementations may not be as well tested in production and at scale, leading to bugs and possible security issues
- Lack of maturity of IPv6 support in (some) firewalls,VPNs, IDS, IPS
- Lack of DNS Block Lists, geolocation, reputation services
- Attack tools beginning to emerge
- Defensive (or offensive) network scanning: see RFC 5157
- State of support of PCI and other regulatory requirements

IPv6 Security issues

- How to correlate network addresses with users, in the face of auto-configuration, temporary addresses, larger address space per subnet
- Local subnet attacks - these are not qualitatively different from what we have in IPv4 today. See RFC 3756 for IPv6 ND based threats.
- Potential covert channel concerns
- Network scanning and router ND queue saturation (DoS)
  - See RFC 6583: Operational problems with neighbor discovery
- Good general discussion of issues and available solutions:
  - https://wikispaces.psu.edu/display/ipv6/IPv6+security

IPv6 Security issues

- Operational security considerations for IPv6 Networks:
  - http://tools.ietf.org/html/draft-ietf-opsec-v6-00
- Security concerns with native and tunneled traffic:
  - http://tools.ietf.org/html/draft-ietf-opsec-ipv6-implications-on-ipv4-nets-00
- Security implications of IPv6 fragmentation and ND:
  - http://tools.ietf.org/html/draft-ietf-6man-nd-extension-headers-01

ICMPv6 filtering

- ICMPv6 is critical to the operation of IPv6 networks
- Used for many functions: Neighbor discovery, router discovery, Path MTU discovery, multicast group membership management (MLD), Mobile IPv6, and more
- Don’t blindly block ICMPv6
- RFC 4890: Recommendations for Filtering ICMPv6 Messages in Firewalls

Rogue RA issue

- Frequently observed phenomenon at some sites
- Most incidents appear to be unintentional misconfiguration rather than malicious
- Appears to be associated with Internet Connection Sharing features in some operating systems
- RFC 6104: Rogue RA problem statement
- Defenses: ACLs, RAGuard (RFC 6105), tweak default router preferences (RFC 4191)
- SeND (cryptographic protocol - challenging to deploy)

Rogue RA vs Rogue DHCP

- IPv4 has to deal with rogue DHCP servers
- Is the situation worse or better with IPv6?
- IPv6 has to deal with both rogue RA and rogue DHCP
- RAs can impact a larger number of hosts faster
- DHCP clients generally have to wait for lease timers to expire
- But, recovery/mitigation can be faster with RA

IPv6 Firewalls

- Stateful Firewalls
- Network vs host based firewalls
- RFC 6092: simple security in IPv6 residential CPE
  - by default block unsolicited incoming except IPsec
- Advanced security CPE?
  - http://tools.ietf.org/html/draft-vyncke-advanced-ipv6-security-03

IPv6 Firewalls

- Status of open source and commercial firewall implementations (Sep 2009, European Conference on Applied IPv6):
  - www.guug.de/veranstaltungen/ecai6-2007/slides/2007-ECAI6-Status-IPv6-Firewalling-PeterBieringer-Talk.pdf
- Survey of IPv6 Availability on Commercial Firewalls (ICANN, March 2010)
  - http://www.icann.org/en/announcements/announcement-2-01mar10-en.htm
- NSA Firewall Design Considerations (July 2010)
  - www.nsa.gov/ia/_files/ipv6/I733-041R-2007.pdf

Microsoft recommendations

- IPv6 security considerations & recommendations (Aug 2011)
- http://technet.microsoft.com/en-us/library/bb726956.aspx
- Discusses SeND and DHCP Authentication, but states Microsoft doesn't support either
- Recommends IPsec: limited support in windows for IPv6 IPsec, but could protect tunneled IPv6 traffic with IPv4 + IPsec
- Recommends IPv6 capable firewalls, IDS, etc

~~Attack~~ Tools

Vulnerability Assessment

- THC-IPv6: http://freeworld.thc.org/thc-ipv6/
- IPv6 Toolkit (SI Networks) http://www.si6networks.com/tools/ipv6toolkit
- scapy - packet manipulation tool
  - http://www.secdev.org/conf/scapy-IPv6_HITB06.pdf
- Note: attacks using IPv6 are already going on today; even on networks that haven't yet deployed IPv6
  - http://tools.ietf.org/html/draft-gont-opsec-ipv6-implications-on-ipv4-nets-00
  - RFC 6169: Security concerns with IPv6 tunneling

Attacks are happening

- IPv6 DDoS attacks observed on the Internet
  - 2012-02-22 Arbor: IPv6 sees first DDoS attacks
  - http://www.h-online.com/security/news/item/Report-IPv6-sees-first-DDoS-attacks-1440502.html
  - http://www.zdnet.com/blog/networking/first-ipv6-distributed-denial-of-service-internet-attacks-seen/2039
- Various forms of IPv6 malware
  - Using IPv6 as covert channel to communicate with botnet controller
  - including one that advertises a host as an IPv6 router and uses v4-v6 transition mechanisms to hijack both IPv4 and IPv6 traffic through it!

References

References

- http://www.upenn.edu/computing/ipv6/
- http://www.upenn.edu/computing/ipv6/strategy.html

References

- http://www.internetsociety.org/deploy360/ipv6/
- http://www.getipv6.info/index.php/Main_Page
- http://www.ietf.org/ (hundreds of protocol specs!)
- http://ipv6.com/
- https://www.arin.net/resources/request/ipv4_depletion.html
- https://www.arin.net/knowledge/v4-v6.html

Questions?

Shumon Huque
shuque -@- upenn.edu

twitter
@shuque