IPv6 Addressing

To be
Presented to
St. Louis Unix Users Group
June 13, 2018
By
David Forrest

Who am I?

David Forrest graduated from Oregon State University in BA, Finance Emphasis, Physics & Mathematics. Lifelong hobbyist in IT from IBM 1401, Model 20, Model 30, Model 85, Sigma 7, XDS 7, SWTP 6800, M6809, 8080, 80286, 80386, OS2, and on to currently running XP, CentOS6/7, Raspbian, Mint, and Chrome on various local and cloud machines.

Why is it NECESSARY

- IPv4 was established in 1981 and provided over four billion separate routable addresses to specific sites.
- Over the past 37 years all addressable addresses have been assigned.
- Network Address Translation “NATing” has allowed expansion for those fortunate sites that have been assigned, but requires severe complication to securely implement.

What is “NATing”

- This red slide and the following 15 to the next red are a prologue “background” and will be largely skipped in the actual presentation.
- This prologue to V6 addressing is an aside largely from Julia Evans site: https://drawings.jvns.ca/nat/ and wikipedia.
- Local addressing is routable only locally and can be “NATed” in the IPv4 system to use a public routable v4 site address (no longer available from the agreed regional international registries, RIRs,) to be shared by those lucky enough to own them.
- This prologue presentation will attempt to cover the local functions needed to implement the “sharing”.

NETWORK ADDRESS TRANSLATION

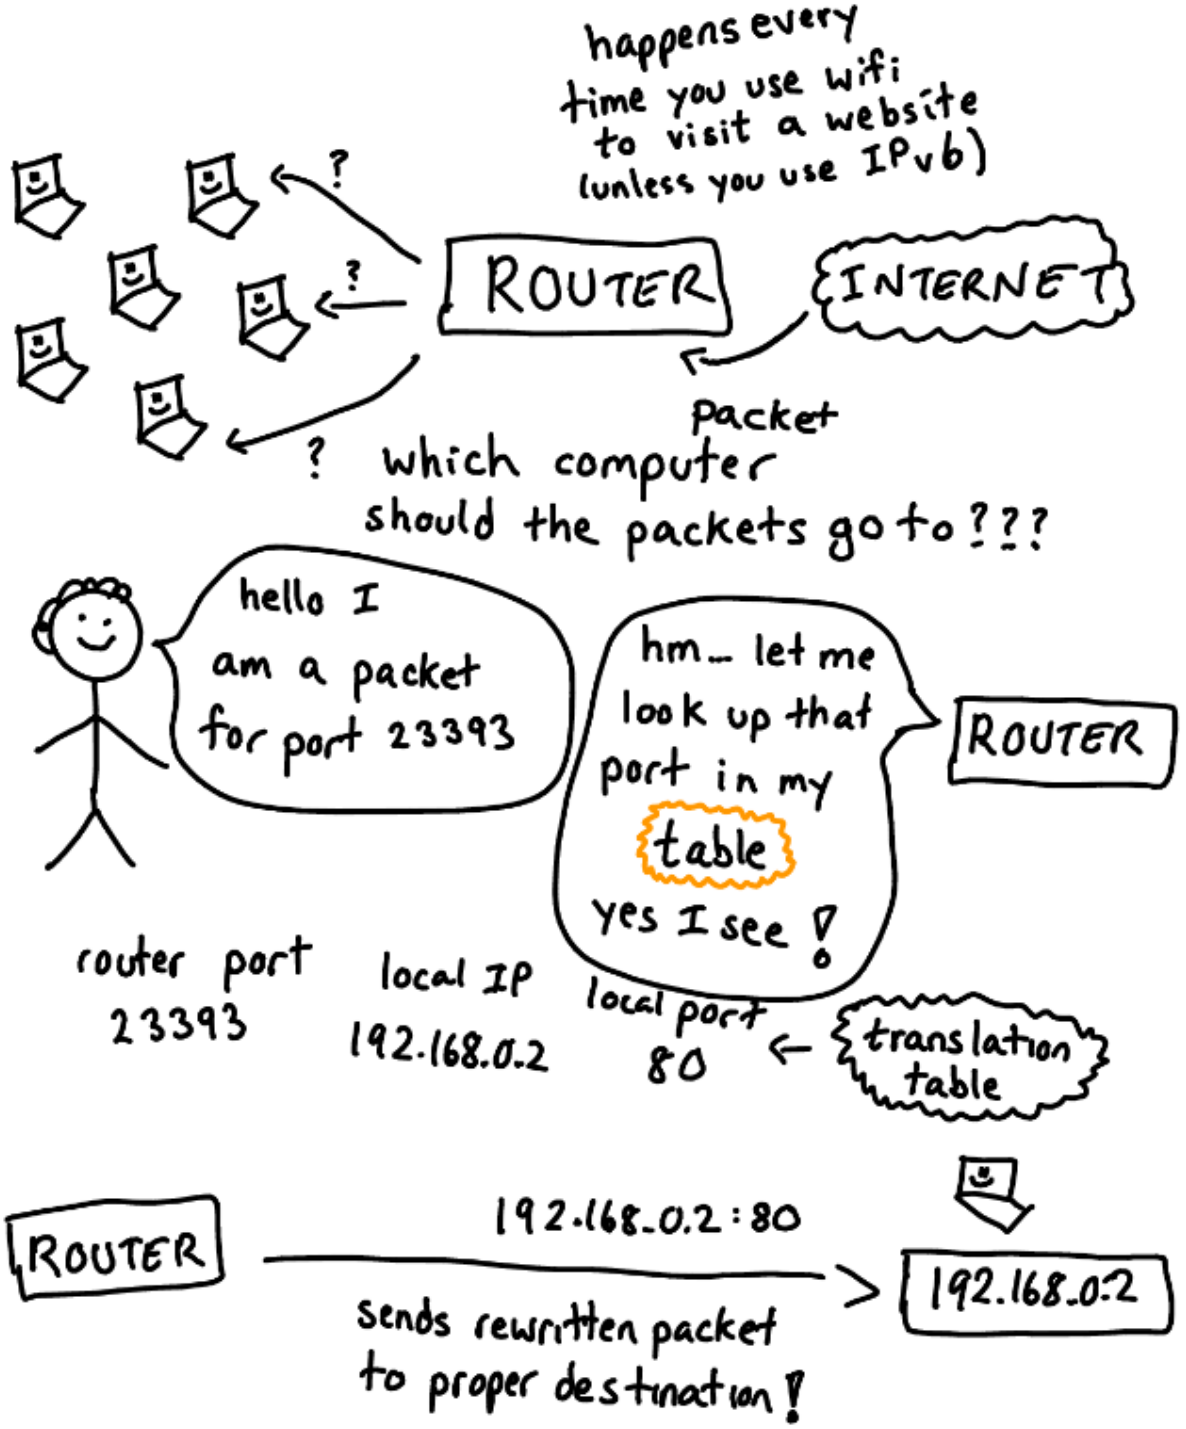

Why is NAT used?

- Because there are not enough addresses for everyone to get routable addresses in IPv4.
- IPv6 was established in 1999 to expand the address base to what was thought at that time to provide routing to all who might require it. And it works quite well, although as an entirely separate stack, was not often used by other than internet transit concerns.
- Some of those used “Carrier Grade NAT” which is similar but not entirely.

NAT Devices

NAT devices allow the use of private IP addresses on private networks behind routers with a single public IP address facing the Internet. The internal network devices communicate with hosts on the external network by changing the source address of outgoing requests to that of the NAT device and relaying replies back to the originating device.

Servers

This leaves the internal network ill-suited for hosting servers, as the NAT device has no automatic method of determining the internal host for which incoming packets are destined. This is not a problem for general web access and email. However, applications such as peer-to-peer file sharing, VoIP services, and video game consoles require clients to be servers as well. Incoming requests cannot be easily correlated to the proper internal host. Furthermore, many of these types of services carry IP address and port number information in the application data, potentially requiring substitution with deep packet inspection.

Standardization

Network address translation technologies are not standardized. As a result, the methods used for NAT traversal are often proprietary and poorly documented. Many traversal techniques require assistance from servers outside of the masqueraded network. Some methods use the server only when establishing the connection, while others are based on relaying all data through it, which increases the bandwidth requirements and latency, detrimental to real-time voice and video communications.

Tranversal

NAT traversal techniques usually bypass enterprise security policies. Enterprise security experts prefer techniques that explicitly cooperate with NAT and firewalls, allowing NAT traversal while still enabling marshalling at the NAT to enforce enterprise security policies. IETF standards based on this security model are Realm-Specific IP (RSIP) and middlebox communications (MIDCOM).

Techniques

Socket Secure (SOCKS) is a technology created in the early 1990s that uses proxy servers to relay traffic between networks or systems.

Traversal Using Relays around NAT (TURN) is a relay protocol designed specifically for NAT traversal.

NAT hole punching is a general technique that exploits how NATs handle some protocols (for example, UDP, TCP, or ICMP) to allow previously blocked packets through the NAT.

Techniques (more)

Session Traversal Utilities for NAT (STUN) is a standardized set of methods and a network protocol for NAT hole punching. It was designed for UDP but was also extended to TCP.

Interactive Connectivity Establishment (ICE) is a complete protocol for using STUN and/or TURN to do NAT traversal while picking the best network route available. It fills in some of the missing pieces and deficiencies that were not mentioned by STUN specification.

Techniques (and more yet)

UPnP Internet Gateway Device Protocol (IGDP) is supported by many small NAT gateways in home or small office settings. It allows a device on a network to ask the router to open a port.

Techniques (… more)

NAT-PMP is a protocol introduced by Apple as an alternative to IGDP.

PCP is a successor of NAT-PMP.

Application-level gateway (ALG) is a component of a firewall or NAT that allows for configuring NAT traversal filters. It is claimed by numerous people that this technique creates more problems than it solves.

Symmetric NATS

The recent proliferation of symmetric NATs has reduced NAT traversal success rates in many practical situations, such as for mobile and public WiFi connections. Hole punching techniques, such as STUN and ICE, fail in traversing symmetric NATs without the help of a relay server, as is practiced in TURN. Techniques that traverse symmetric NATs by attempting to predict the next port to be opened by each NAT device were discovered in 2003 by Yutaka Takeda at Panasonic Communications Research Laboratory and in 2008 by researchers at Waseda University. Port prediction techniques are only effective with NAT devices that use known deterministic algorithms for port selection. This predictable yet non-static port allocation scheme is uncommon in large scale NATs such as those used in 4G LTE networks and therefore port prediction is largely ineffective on those mobile broadband networks.

IPsec

IPsec virtual private network clients use NAT traversal in order to have Encapsulating Security Payload packets traverse NAT. IPsec uses several protocols in its operation which must be enabled to traverse firewalls and network address translators:

Internet Key Exchange (IKE) – User Datagram Protocol (UDP) port 500

Encapsulating Security Payload (ESP) – IP protocol number 50

Authentication Header (AH) – IP protocol number 51

IPsec NAT traversal – UDP port 4500, when NAT traversal is in use

Many routers provide explicit features, often called IPsec Passthrough.

Some Windows Issues

In Windows XP, NAT traversal is enabled by default, but in Windows XP with Service Pack 2 it has been disabled by default for the case when the VPN server is also behind a NAT device, because of a rare and controversial security issue. IPsec NAT-T patches are also available for Windows 2000, Windows NT and Windows 98.

NAT traversal and IPsec may be used to enable opportunistic encryption of traffic between systems. NAT traversal allows systems behind NATs to request and establish secure connections on demand.

Hosted NAT Transversal

Hosted NAT traversal (HNT) is a set of mechanisms, including media relaying and latching, used by intermediaries.

The IETF advises against using latching over the Internet and recommends ICE for security reasons.

My Thoughts on NAT

All of the confusion/complexities on NAT (although some still believe it gives them “safety through routing”) leads me to conclude that the new paradigm in IPv6 properly excludes such except in prefix NATing necessary for those using the “Unique (Universal) Local Address” scheme available in IPv6. IPv6 allows enough addresses to give safe and effective firewall routing control throughout our assigned internet addresses.

And maybe not NPT

IPv6-to-IPv6 Network Prefix Translation (**NPTv6**) is an experimental specification for IPv6 to achieve the address-independence at the network edge, given by network address translation (NAT) in IPv4. It has fewer architectural problems than traditional IPv4 NAT; it is for example stateless and preserves the reachability attributed to the end-to-end principal. However, the method still lacks solutions to translate embedded IPv6 addresses, for example in IPsec, and requires a more complex nameserver setup (split-horizon DNS).

So what?

So what is NPTv6? NPTv6 is simply rewriting IPv6 prefixes. If your current IPv6 prefix is 2001:db8:cafe::/48 then using NPTv6 it would allow you to change it to 2001:db8:fea7::/48 - that is it. It is a one for one prefix rewrite - you can't overload it, have mismatching prefix allocations sizes, re-write ports or anything else. Importantly, it doesn't touch anything other than the prefix. Your subnetwork/host portion remains intact with no changes.

NAT Summary

- NAT is not needed for address sharing, but maybe for other situations.
- In IPv4 networks, we solved the shortage of addresses by using NAT to share one public IP address between many hosts. In IPv6, we have no address shortage and do not need to share IP addresses any more. For other uses of NAT work is still going on to figure out how to solve these – but we will propably end up using NAT66 in some situations in our networks. By learning how the use of NAT and private address space breaks the network architecture and adds costs to projects like VoIP and causes additional delays in the network we will not add these by default when building IPv6 networks.

Neelie Kroes
Vice-President of the European Commission
responsible for the Digital Agenda

“We must make this transition. The alternative is that the Internet will begin to suffer; and innovation and economic growth will feel the consequences. These are not things we can afford at the moment. To all of you out there doing business on the Internet: governments, content providers, service providers, my message is clear. Switch to IPv6 as soon as possible. And we can start enjoying the amazing opportunities of the future Internet.”

IPv6 Site Assignment

- My router's current assignment from Hurricane Electric is 2001:470:c416:0001::1/48
- That's 48 bits of my 64 bit network leaving me the ability to filter the last 16 bits, 65,536 separate networks, as best suits my organizational needs.
- $ host erl.maplepark.com
- 1.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0. \ (# local aka fe80::1) 1.0.0.0.6.1.4.c.0.7.4.0.1.0.0.2.ip6.arpa
- domain name pointer erl.maplepark.com.

What is it?

IPv4 stack	IPv6 stack
1981	1999
32 bit	128 bit
192.168.0.0/32 (4 Octets)	C0A8:0::/32 (8 Hex Nibbles)
2^32 Addresses	2^128 (64 Prefix & 64 Link)

And WHY!


Need more sales? If so, you need more printing!
Let our full-color business printing help BRAND your business!
Fast & Affordable
Digital Printing!
Rose Brothers
757-898-3051
www.Rose-Brothers.com
$10 OFF REPAIR
thank you
BRAND: the marketing practice of using a logo, name, or design that identifies and differentiates your business from others. From stationary and brochures to envelopes and presentation folders, let Minuteman Press - Des Peres help brand your business!
CALL 636-256-2475 TO SAVE MONEY ON YOUR NEXT PROJECT!
Stationary • Note Cards • Full Color Envelopes • Business Cards
Note Pads • Carbonless Forms • Presentation Folders • Custom Binders
MP

The world's a'changing

- IPv6 is an entirely different stack and is handled internally completely differently from IPv4.
- Current thought seems to favor change in parallelism in lieu of meshing – new projects in dual stack and conversion of older into dual as maintenance requires.
- My adoption in 2006 was two steps: first – get IPv6 access publicly and second to add the address to my httpd.conf file.

IPv6 Address Planning

- My plan was simple: get one and use it. I only needed one and only to serve a web page or two.
- Much of what I want to talk about is here: https://www.internetsociety.org/policybriefs/ipv6

Various Plans

- Surf-net
  https://ipv6forum.com/dl/presentations/IPv6-addressing-plan-howto.pdf
- Infoblox
  https://www.infoblox.com/wp-content/uploads/2016/04/infoblox-whitepaper-ipv6-addressing-plan-basics_1.pdf
- Ipspace
  http://blog.ipspace.net/2015/04/how-do-i-start-my-ipv6-addressing-plan.html

ARIN


ARIN
American Registry for Internet Numbers
RIPE NCC
RIPE NETWORK COORDINATION CENTRE
APNIC
lacnic
AFRINIC

American Registry of Internet Numbers

- About IPv6:
  https://www.arin.net/knowledge/ipv6_info_center.html
- Everything you wanted to know about IPv6
  https://youtu.be/rWJZfShWE6g
- A reserved IPv4 /24 for deployment of IPv6!

ARIN on IPv6 Anatomy

- How is IPv6 Different than IPv4?
- IPv6 differs from IPv4 in many ways, including address size, format, notation, and possible combinations. We've created a quick video to highlight some of the differences.
-
- An IPv6 address consists of 128 bits (as opposed to the 32-bit size of IPv4 addresses) and is expressed in hexadecimal notation. The IPv6 anatomy graphic below represents just one possible configuration of an IPv6 address, although there are many different possibilities.

IPv6 Anatomy

DREN

- Defense Research and Engineering Network
- This presentation is from them in 2011 – old, old, and older – but says a lot!

http://maplepark.com/~drf/RemoteReads/DREN_AddressingPlans.pdf

Fin

- And that's it, folks!
- I'm open and interested in questions (?)

Thanks,