IPv6

Bitlair Workshop

May 6, 2017

Introductions

Introductions: Bitlair, me, you?

Before we begin

Drinks can be bought at the bar.

- Deposit 5 YOURNAME (to deposit 5 euro)
- Yes, you want to create an account
- To buy an item:
- SCAN ITEM
- YOURNAME

Table of contents

Let's start with some basics...

What happens when..

Browsing to https://bitlair.nl/Workshops

- Domain name system lookup bitlair.nl
- Get back IP 2001:470:d1e7:1337:5054:ff:fe49:a94
- Connect to 2001:470:d1e7:1337:5054:ff:fe49:a94
- Set up secure connection (S part of HTTPS)
- Do a GET request for the /Workshops page

Internet Protocol

We've all heard the term IP address

- An IP address is the location in the network to route packets to
- It is therefore (mostly) hierarchical outside the Default Free Zone

The IPv4 problem

The IPv4 problem

The IPv4 problem

- IPv4 was never meant to connect the whole planet
- Only 3.7 billion addresses available
- Lots of small IP blocks make for large routing tables
- We lost end-to-end connectivity because of NAT
- Small pool of private addresses leads to conflicts
- Importantly, though: There are no more addresses for new Internet Service Providers or new devices.

Lots of small IP blocks

- The Default Free Zone, where ISPs exchange routing information, currently has 668000 IPv4 blocks
- This is stored in routers that need to forward terabits of packets per second
- For 100Gbps, 6.4 nanoseconds total time per packet
- Fast memory is very expensive

NAT - The root of all evil

- To connect the western world, we could only allow 1 IP per household
- That means multiple devices have to share 1 address
- That means that there is no publicly reachable address for your IP phone
- ... but the IP phone needs to send and receive UDP packets
- ... to another phone behind NAT? Oh shit.

NAT - The root of all evil

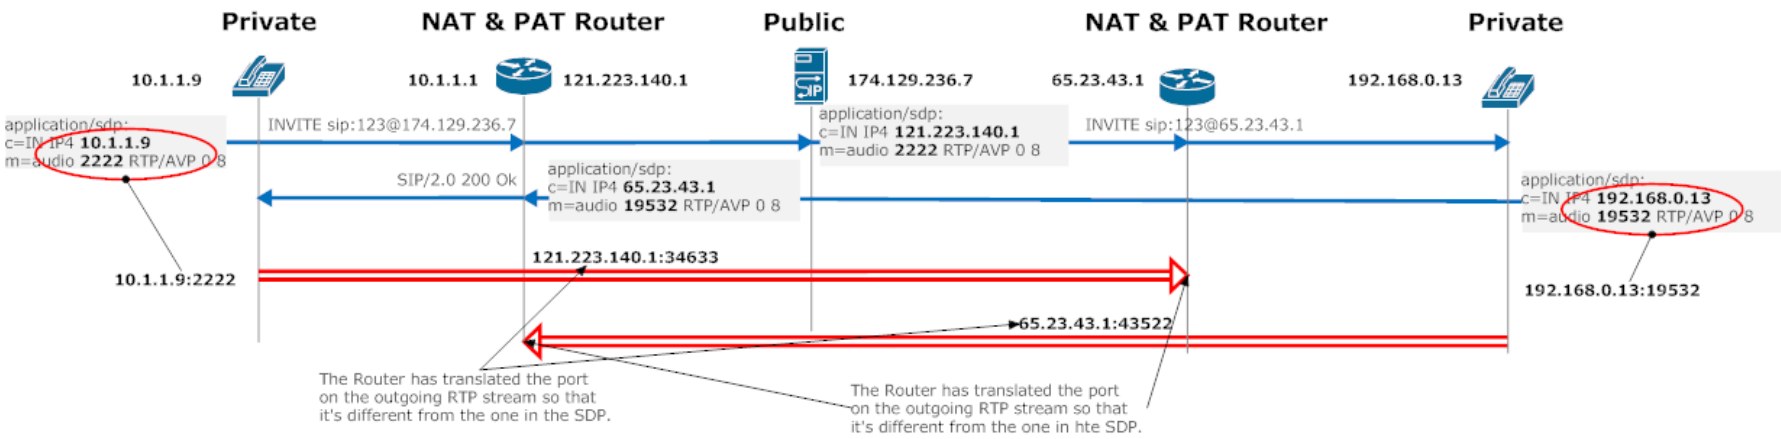

NAT - The root of all evil

But won't that be the same with firewalls without NAT?

- No! They can tell the firewall to allow the stream, by both initiating the UDP stream in opposite directions
- NAT prevents this, because there is no way to know the port to connect to.. router has the IP

Address space conflicts

- You run with 10.42.0.1/24 at home.. Oh, but your hackerspace also does this. You cannot connect your VPN.
- You want to use internet connection sharing in Network Manager, this uses 10.42.0.1/24.. but your hackerspace also does this. You can not reach the network anymore.
- Two companies both use 172.16.0.0/16, they need to merge.. Good luck.

Root of the problem

We are out of IPv4 addresses

Internet History

- 1969: ARPANET was created, ran the NCP protocol
- 1981: ARPANET expanded to Science community
- 1981: The internet protocol was developed
- 1983: First deployment of the Internet Protocol
- 1986: Foundation of the Internet Engineering Task Force
- 1988: First European Internet connection (CWI)
- 1993: First private citizens connect to the Internet

Internet History

- 1992: Address exhaustion predicted: Call for papers (IETF)
- 1994: Adoption of IPng model by the IETF
- 1996: IPng published as IPv6 RFCs by the IETF
- 2011: Top level IPv4 exhausted. Asia out of IPv4 addresses
- 2012: Europe out of IPv4 addresses
- 2014: Latin America and Caribbean out of IPv4 addresses
- 2015: North America out of IPv4 addresses

So IPv6

So... IPv6

The cause of and solution to all of life's problems

Enter IPv6

- 128-bit addresses ($2^{128} > 10^{38}$ addresses)
- End to end connectivity for the entire world
- Large blocks per ISP
- Smaller routing table for ISPs
- Enough addresses per household (at least 256 networks)

No more NAT!

- Yay!
- Torrents work!
- VoIP works!
- Bidirectional UDP works!
- Firewalls don't have to track sessions anymore

Addressing

IPv6 addressing

Addressing

- IPv6 uses hexadecimal notation
- IPv4 was annoying as hell with decimal notation and prefixes

Addressing

2001:0db8:0080:8001:0000:0000:0000:0001

2001:db8:80:8001:0:0:0:1

2001:db8:80:8001::1

Counting in HEX

With 4 bits you can get 1 HEX value
Thus 2F is 8 bits (1 byte), is in binary 0010 1111

Counting in Hex

Binary	Hex	Binary	Hex
0 0 0 0	0	1 0 0 0	8
0 0 0 1	1	1 0 0 1	9
0 0 1 0	2	1 0 1 0	A
0 0 1 1	3	1 0 1 1	B
0 1 0 0	4	1 1 0 0	C
0 1 0 1	5	1 1 0 1	D
0 1 1 0	6	1 1 1 0	E
0 1 1 1	7	1 1 1 1	F

Counting in binary, IPv4 style

Blue == 1 ('on'), white == 0
If the bit is 'on', just add them
8 bits (thus 1...128) = 1 byte

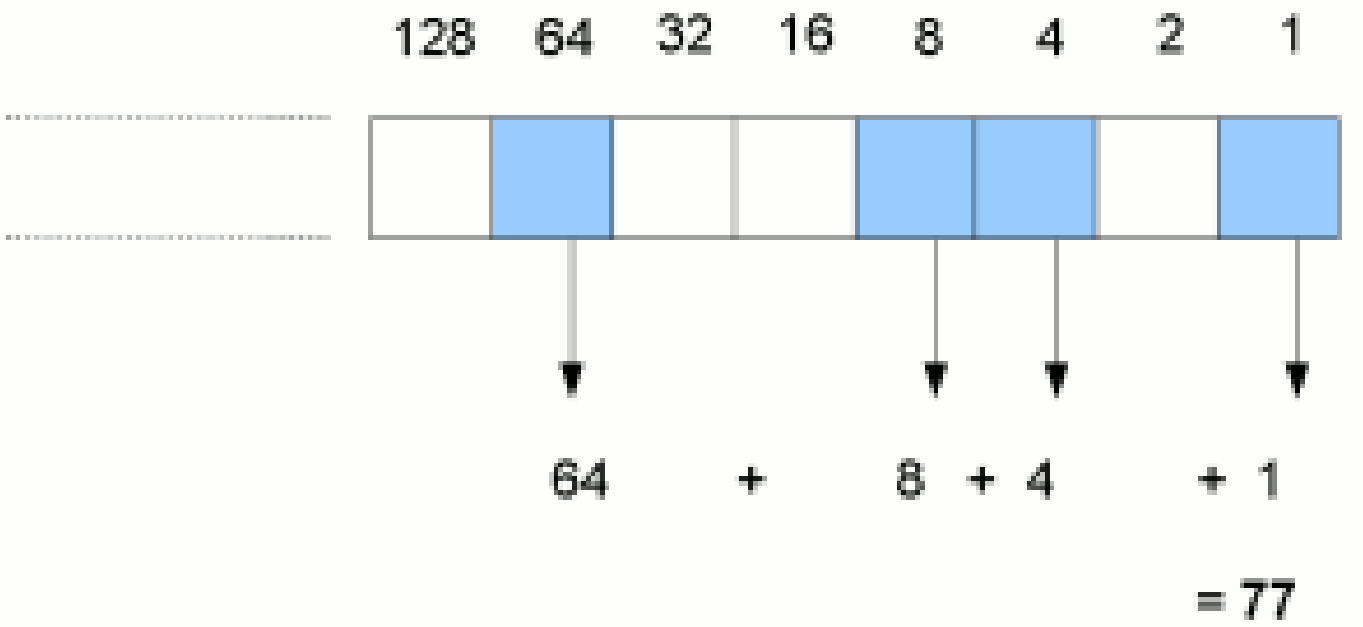

Hacking time

Your turn!

- What is hex 8E in binary?
- What are the bits in address 2001:db8::1
- Care to try IPv4: decimal 234 in binary?

Subnetting

Subnetting

Basically the same as with IPv4

Subnetting

- An IPv6 subnet has a prefix
- The prefix has a length, most often /64
- The prefix length denotes how many bits are fixed in the subnet
- Example: 2001:db8::/32 is the documentation prefix
- 2001:db8 is fixed, everything after this part of the subnet

Routing

Example topology

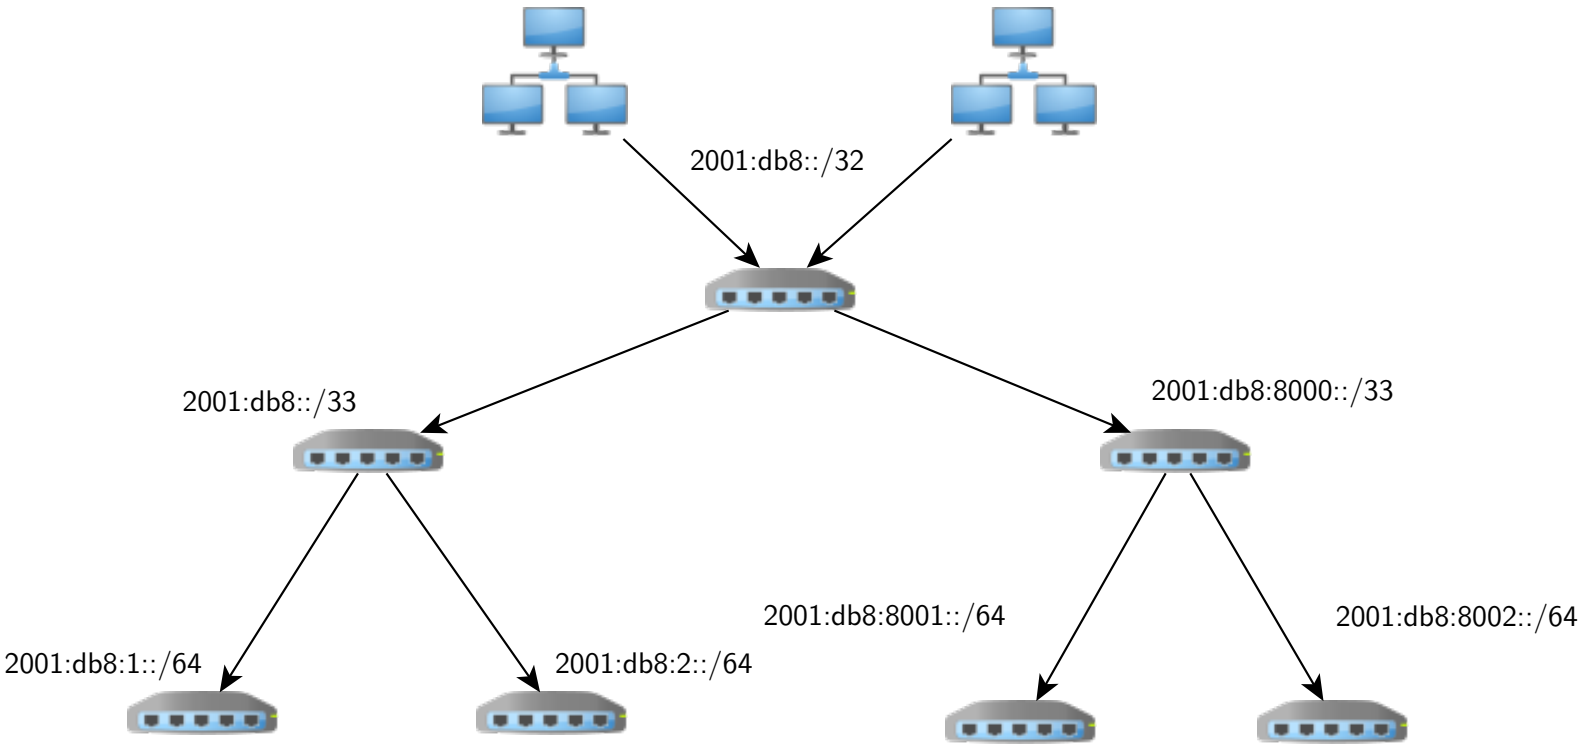

Hacking time

Your turn!

- What bits are fixed in 2001:db8::/64?
- How many addresses are in 2001:db8:3:3::/64?
- How many /64 subnets can we make in 2001:db8:8000::/48?
- Care to try IPv4? Which bits are fixed in 192.0.2.44/29?

Routing

Routing

Basically the same as with IPv4

Routing

A routing table is nothing more than a Longest Prefix Match table, containing the next hop and/or interface.

Prefix	Next hop	Interface
2001:db8:2000:1::/64	fe80::2	eth0
fe80::/64		eth0
2001:db8:2000::/48		blackhole
::/0	fe80::1	eth0

Address configuration

Address configuration

Address configuration

- Routers sends router advertisements to clients
- In the router advertisements, there are options
- In the options, there may be a prefix
  .. or the managed flag, stating that DHCPv6 must be used
- There is DNS information
  .. and/or the other configuration flag, for getting extra options (but not address) from DHCPv6

For ISPs

- ISP router sends a managed configuration flag
- Client router will do DHCPv6
- ISP router sends DHCPv6 Prefix Delegation option
- Client router now knows what prefix can be used for subnets in the home/office
- Client router assigns IP and announces some /64 prefixes in the home with SLAAC

Stateless address configuration (SLAAC)

The most common case

- Router sends a prefix (always /64 length on ethernet)
- Router sends DNS and/or other-config flag (both for best compatibility)
- Run DHCPv6 stateless, just for DNS options
- Client configures EUI-64 address and/or privacy address

EUI-64

- EUI-64 is a 64-bit derivation of the mac-address in a 64-bit prefix
- But.. bit 1 of byte 0 is flipped and ff:fe is added in the middle
- Example: 40:6A:3B:3F:4A:8C
- 2001:db8::426a:3bff:fe3f:4a8c

Privacy address

- Privacy addresses are 64-bit random to the 64-bit prefix
- So example: 2001:db8::c3e7:4fe4:b56d:a8b2:d30f
- Rotates every hour
- Up to 16 old addresses are kept active for long sessions

Getting your own IPv6

- Ask your ISP!
- In the meantime... go to tunnelbroker.net and request a free tunnel
- Configure it on your router

Example VyOS config

```
interfaces {
     ethernet eth0 {
          description "Uplink"
          address dhcp
          address dhcpv6
     }
}
```

Example VyOS config

```
interfaces {
    ethernet eth1 {
        description "LAN"
        address 10.0.0.1/24
        address 2001:db8:1f15:d04::1/64
        ipv6 {
            router-advert {
                prefix 2001:db8:1f15:d04::/64 {
                }
                name-server 2001:4860:4860::8888
                send-advert true
            }
```

Example VyOS config

Or tunneled...

Example VyOS config

```
interfaces {
     ethernet eth0 {
          description "Uplink"
          address dhcp
     }
}
```

Example VyOS config tunneled

```
interfaces {
      tunnel tun0 {
          address 2001:db8:1f14:d04::2/64
          description "HE.NET IPv6 Tunnel"
          encapsulation sit
          local-ip 192.0.2.2 # Your local IP
          remote-ip 216.66.84.46
      }
}
```

Example VyOS config tunneled

```
protocols {
      static {
          interface-route6 ::/0 {
              next-hop-interface tun0 {
              }
          }
      }
}
```