The Art of Running Out of IPv6 Addresses

Benedikt Stockebrand
Stepladder IT Training+Consulting GmbH

RIPE 77
October 2018
Amsterdam, The Netherlands

Copyright © 2018 Benedikt Stockebrand

NOTICE

If you find this presentation on the Web, be aware it only contains the backdrop of the talk. Please check

https://ripe77.ripe.net/archives/

for a video recording of the actual talk to get the whole story.

Free of charge, no registration required :-)

About Me

- Long term Unix, TCP/IP, Internet, architecture, operations, major cleanup guy
- Got serious with IPv6 in 2003
- Wrote a book about it in 2006
- Co-authored a study on secure IPv6 deployment for the German Federal Office for IT Security (BSI) in 2010
- Became one of your friendly IPv6 WG co-chairs in 2014
- Currently do a contract as "program manager" for IPv6 deployment at the datacenter operator for the state of Hessen (Germany)

Copyright © 2018 Benedikt Stockebrand

Part I

Science

What I Tell People in Trainings

Number of IPv6 addresses:

$$2^{128} = 340\,282\,366\,920\,938\,463\,463\,374\,607\,431\,768\,211\,456 \approx 3.4 \times 10^{38}$$

Copyright © 2018 Benedikt Stockebrand

What I Tell People in Trainings

Number of IPv6 addresses:

$$2^{128} = 340\,282\,366\,920\,938\,463\,463\,374\,607\,431\,768\,211\,456 \approx 3.4 \times 10^{38}$$

Ratio of IPv6 to IPv4 addresses:

$$2^{96} : 1 = 79\,228\,162\,514\,264\,337\,593\,543\,950\,336 \approx 8 \times 10^{28}$$

Copyright © 2018 Benedikt Stockebrand

What I Tell People in Trainings

Number of IPv6 addresses:

$$2^{128} = 340\,282\,366\,920\,938\,463\,463\,374\,607\,431\,768\,211\,456 \approx 3.4 \times 10^{38}$$

Ratio of IPv6 to IPv4 addresses:

$$2^{96} : 1 = 79\,228\,162\,514\,264\,337\,593\,543\,950\,336 \approx 8 \times 10^{28}$$

Ratio of IPv6 site prefixes (/48) to IPv4 /24 prefixes:

$$2^{24} : 1 = 16\,777\,216 \approx 1.7 \times 10^{7}$$

Copyright © 2018 Benedikt Stockebrand

What I Tell People in Trainings

Number of IPv6 addresses:

$$2^{128} = 340\,282\,366\,920\,938\,463\,463\,374\,607\,431\,768\,211\,456 \approx 3.4 \times 10^{38}$$

Ratio of IPv6 to IPv4 addresses:

$$2^{96} : 1 = 79\,228\,162\,514\,264\,337\,593\,543\,950\,336 \approx 8 \times 10^{28}$$

Ratio of IPv6 site prefixes (/48) to IPv4 /24 prefixes:

$$2^{24} : 1 = 16\,777\,216 \approx 1.7 \times 10^{7}$$

Ratio of IPv6 site prefixes (/56) to IPv4 /24 prefixes:

$$2^{32} : 1 = 4\,294\,967\,296 \approx 4.3 \times 10^{9}$$

Copyright © 2018 Benedikt Stockebrand

What I Tell People in Trainings

Number of IPv6 addresses:

$$2^{128} = 340\,282\,366\,920\,938\,463\,463\,374\,607\,431\,768\,211\,456 \approx 3.4 \times 10^{38}$$

Ratio of IPv6 to IPv4 addresses:

$$2^{96} : 1 = 79\,228\,162\,514\,264\,337\,593\,543\,950\,336 \approx 8 \times 10^{28}$$

Ratio of IPv6 site prefixes (/48) to IPv4 /24 prefixes:

$$2^{24} : 1 = 16\,777\,216 \approx 1.7 \times 10^{7}$$

Ratio of IPv6 site prefixes (/56) to IPv4 /24 prefixes:

$$2^{32} : 1 = 4\,294\,967\,296 \approx 4.3 \times 10^{9}$$

There is no shortage of addresses with IPv6!

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k, n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k,n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k,n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.
- For k addresses we need $\log_2 k$ bits (theoretical optimum).

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k, n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.
- For k addresses we need $\log_2 k$ bits (theoretical optimum).
- E.g. for 5 addresses we need 2.321928... bits.

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k,n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.
- For k addresses we need $\log_2 k$ bits (theoretical optimum).
- E.g. for 5 addresses we need 2.321928. . . bits.
- Fractions make sense; see Shannon (1948/49).

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k,n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.
- For k addresses we need $\log_2 k$ bits (theoretical optimum).
- E.g. for 5 addresses we need 2.321928... bits.
- Fractions make sense; see Shannon (1948/49).
- Address space utilization is measured logarithmically,
  i.e. in bits needed/allocated.

Copyright © 2018 Benedikt Stockebrand

The HD Ratio (RFC 3194) and IPv6

$$HD(k,n) = \frac{\log_2 k}{\log_2 n}$$

with k the number of addresses used and n the number of addresses available.

- $\log_2 n$ is the number of address bits available.
- For k addresses we need $\log_2 k$ bits (theoretical optimum).
- E.g. for 5 addresses we need 2.321928... bits.
- Fractions make sense; see Shannon (1948/49).
- Address space utilization is measured logarithmically, i.e. in bits needed/allocated.

- Think in bit utilization, not addresses.

Copyright © 2018 Benedikt Stockebrand

Part II

Real World Strategies and Tactics

Waste Bits, not Addresses

Waste Bits, not Addresses

- Wasting addresses is *way* too tedious

Waste Bits, not Addresses

- Wasting addresses is *way* too tedious
- Wasting address *bits* is our only chance. . .

Copyright © 2018 Benedikt Stockebrand

Waste Bits, not Addresses

- Wasting addresses is *way* too tedious
- Wasting address *bits* is our only chance...
- ...but in doing so we achieve the inconceivable!

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[. . .]

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[. . .]
- Combine with Microsegmentation

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[. . .]
- Combine with Microsegmentation
- 99.999 999 999 999 999 994 58% of the addresses are already gone

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[. . .]
- Combine with Microsegmentation
- 99.999 999 999 999 999 994 58% of the addresses are already gone

- 18 446 744 073 709 551 616 subnet prefixes to go. . .

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[. . .]
- Combine with Microsegmentation
- 99.999 999 999 999 999 994 58% of the addresses are already gone

- 18 446 744 073 709 551 616 subnet prefixes to go. . .
- . . . or 64 bits. . .

Copyright © 2018 Benedikt Stockebrand

Half the Work Done: RFC 4291

- RFC 4291, section 2.5.1:
  For all unicast addresses, except those that start with the binary value 000, Interface IDs are required to be 64 bits long[...]
- Combine with Microsegmentation
- 99.999 999 999 999 999 994 58% of the addresses are already gone

- 18 446 744 073 709 551 616 subnet prefixes to go...
- ... or 64 bits...
- ... or 8 ~~bytes~~ octets

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

- The Internet grows, and exponentially so.

The Temporal Dimension

- The Internet grows, and exponentially so.
- Every time the Internet doubles in "size" we need another bit in our addresses.

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

- The Internet grows, and exponentially so.
- Every time the Internet doubles in “size” we need another bit in our addresses.
- The Internet grows at a constant bit rate.

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

- The Internet grows, and exponentially so.
- Every time the Internet doubles in “size” we need another bit in our addresses.
- The Internet grows at a constant bit rate.
- *Suitable* measurements of Internet growth are difficult to get hold of right now.

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

- The Internet grows, and exponentially so.
- Every time the Internet doubles in “size” we need another bit in our addresses.
- The Internet grows at a constant bit rate.
- *Suitable* measurements of Internet growth are difficult to get hold of right now.
- Once the effects of DS-Lite etc. wear off we can hope for reasonable predictions again.

Copyright © 2018 Benedikt Stockebrand

The Temporal Dimension

- The Internet grows, and exponentially so.
- Every time the Internet doubles in “size” we need another bit in our addresses.
- The Internet grows at a constant bit rate.
- *Suitable* measurements of Internet growth are difficult to get hold of right now.
- Once the effects of DS-Lite etc. wear off we can hope for reasonable predictions again.
- Rough estimate: In 50–100 years we will outgrow IPv6.

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
    - IANA

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses
  - National offices of the businesses

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses
  - National offices of the businesses
  - Local offices

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses
  - National offices of the businesses
  - Local offices
  - Local teams

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses
  - National offices of the businesses
  - Local offices
  - Local teams
  - Personal fiefdoms

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
    - IANA
    - RIPE NCC
    - Major international conglomerate
    - Individual businesses
    - National offices of the businesses
    - Local offices
    - Local teams
    - Personal fiefdoms
- Now let every layer stack some “few percent extra”. . .

Copyright © 2018 Benedikt Stockebrand

Multi Level Delegation

- For most practical purposes, we can only delegate at bit boundaries
- But we can do so multiple times, burning up to a bit every time, e.g.
  - IANA
  - RIPE NCC
  - Major international conglomerate
  - Individual businesses
  - National offices of the businesses
  - Local offices
  - Local teams
  - Personal fiefdoms
- Now let every layer stack some "few percent extra"...
- ... or even allocate at nibble boundaries.

Bit Repurposing

Bit Repurposing

- What are IP addresses there for?

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask)
    (to simplify packet filter configurations)

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask) (to simplify packet filter configurations)
  - geographic coordinates

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask) (to simplify packet filter configurations)
  - geographic coordinates
  - some "magic" IoT stuff (see RIPE-76)

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask)
    (to simplify packet filter configurations)
  - geographic coordinates
  - some "magic" IoT stuff (see RIPE-76)
  - replace "IoT" with another buzzword, rinse and repeat

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask) (to simplify packet filter configurations)
  - geographic coordinates
  - some "magic" IoT stuff (see RIPE-76)
  - replace "IoT" with another buzzword, rinse and repeat
  - E.164 telephone numbers

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask) (to simplify packet filter configurations)
  - geographic coordinates
  - some "magic" IoT stuff (see RIPE-76)
  - replace "IoT" with another buzzword, rinse and repeat
  - E.164 telephone numbers
  - ...

Copyright © 2018 Benedikt Stockebrand

Bit Repurposing

- What are IP addresses there for?
- To hold *routing information*.

- So let's put other things in there:
  - Organizational structure
  - port numbers served from the subnet (as bitmask)
    (to simplify packet filter configurations)
  - geographic coordinates
  - some "magic" IoT stuff (see RIPE-76)
  - replace "IoT" with another buzzword, rinse and repeat
  - E.164 telephone numbers
  - ...
- What about a Tech-C mail address?

Copyright © 2018 Benedikt Stockebrand

~~Alpha Male Behaviour~~ Politics

Copyright © 2018 Benedikt Stockebrand

~~Alpha Male Behaviour~~ Politics

- Always demand more than your peer got.

Copyright © 2018 Benedikt Stockebrand

~~Alpha Male Behaviour~~ Politics

- Always demand more than your peer got.
- Never mind evaluating your actual requirements.

Copyright © 2018 Benedikt Stockebrand

~~Alpha Male Behaviour~~ Politics

- Always demand more than your peer got.
- Never mind evaluating your actual requirements.
- Facts are for wimps.

Copyright © 2018 Benedikt Stockebrand

Profit Optimization

Copyright © 2018 Benedikt Stockebrand

Profit Optimization

- As an ISP, get an allocation large enough to provide every customer with a /48...

Copyright © 2018 Benedikt Stockebrand

Profit Optimization

- As an ISP, get an allocation large enough to provide every customer with a /48...
- ... and then give out /62s.

Copyright © 2018 Benedikt Stockebrand

Profit Optimization

- As an ISP, get an allocation large enough to provide every customer with a /48...
- ... and then give out /62s.
- Keep the rest "for future use".

- It works with IPv4, why shouldn't it work with IPv6, too?

Copyright © 2018 Benedikt Stockebrand

Things We Learn From History

Things We Learn From History

Picture: Wikipedia

Things We Learn From History

Picture: Wikipedia

Georg Wilhelm Friedrich Hegel

Things We Learn From History

Picture: Wikipedia

Georg Wilhelm Friedrich Hegel
German Philosopher

Things We Learn From History

Picture: Wikipedia

Georg Wilhelm Friedrich Hegel
German Philosopher

The only thing we learn from history

Things We Learn From History

Picture: Wikipedia

Georg Wilhelm Friedrich Hegel
German Philosopher

The only thing we learn from history is that we learn nothing from history.

Contact Information

Stepladder IT
Training+Consulting GmbH
Benedikt Stockebrand

Fichardstr. 38
D-60322 Frankfurt/Main
Germany

`contact@stepladder-it.com`

Web pages:
`http://www.stepladder-it.com/`
`http://www.benedikt-stockebrand.de/`

Video blog:
`http://www.stepladder-it.com/bivblog/`