Where did all those IPv6 addresses go?

David Conrad
channelling
Geoff Huston

It seems rather odd…

- To be considering address exhaustion issues in a technology that is really only ramping up
  - “If the earth were made entirely out of 1 cubic millimetre grains of sand, then you could give a unique address to each grain in 300 million planets the size of the earth.” -- Wikipedia
- And this is a highly speculative exercise…

IETF IPv6 Address Structure

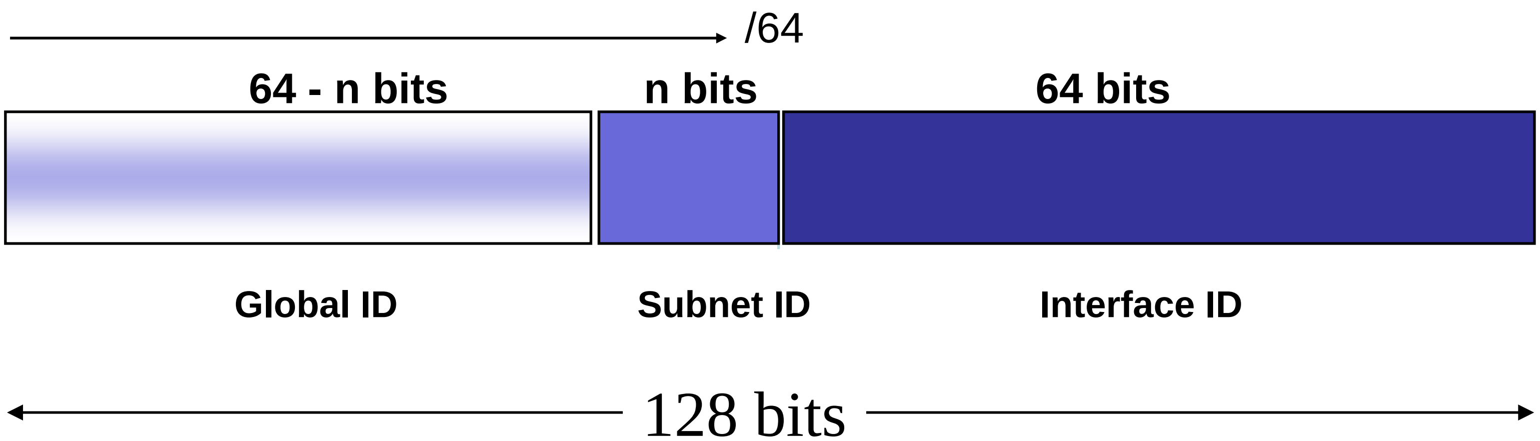

RIR IPv6 Address Structure

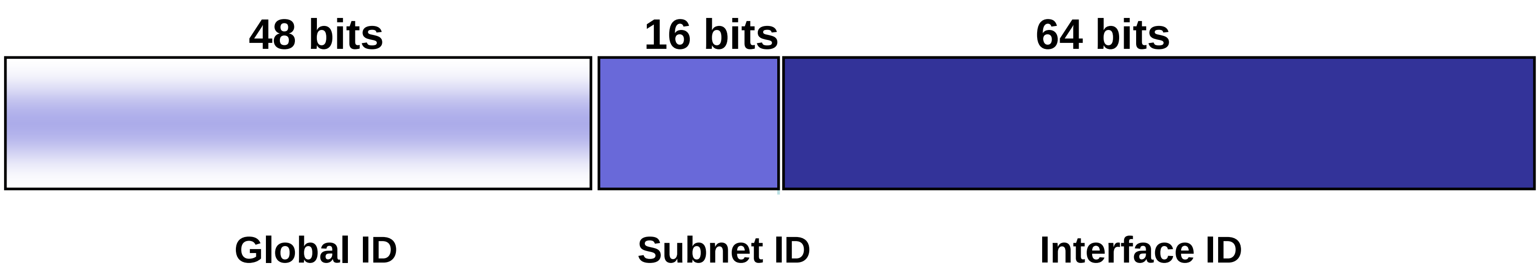

Current Policy

- RIR to ISP(LIR):
  - Initial allocation: /32 (minimum)
  - Subsequent allocation: /32 (minimum)
- ISP(LIR) to customer:
  - Only 1 interface ever: /128
  - Only 1 subnet ever: /64
  - Everything else: /48 (minimum)
- ISP(LIR) to each POP:
  - /48

The HD-Ratio of 0.8

Host Count	80%	HD = 0.8
End Customer Size	**IPv4 Allocation**	**IPv6 Allocation**
205	**/24**	**/32**
410	**/23**	**/32**
819	**/22**	**/32**
1638	**/21**	**/32**
3277	**/20**	**/32**
7131	**/18**	**/32**
12416	**/18**	**/31**
21618	**/17**	**/30**
37640	**/16**	**/29**
65536	**/15**	**/28**
114104	**/14**	**/27**
198668	**/14**	**/26**
345901	**/13**	**/25**
602248	**/12**	**/24**
1048576	**/11**	**/23**
1825676	**/10**	**/22**
3178688	**/10**	**/21**
5534417	**/9**	**/20**
9635980	**/8**	**/19**
16777216	**/7**	**/18**

Google(subscribers millions)

- Broadband
  - 150 million total globally
    - 85 Million DSL Globally
      - 12 Million in US Today
      - 58 Million in US in 2008
- Cellular
  - Cingular: 50 Million
  - Verizon: 43 Million
  - Korea: 37 Million
  - Russia: 20 Million
  - Asia: 560 Million
    - China: 580 million subscribers by 2009

QuickTime™ and a
TIFF (Uncompressed) decompressor
are needed to see this picture.

Squeezing in Bigger Numbers for Longer Timeframes

- The demand - global populations:
  - Households, Workplaces, Devices, Manufacturers, Public agencies
  - Thousands of service enterprises serving millions of end sites in commodity communications services
  - Addressing technology to last for decades
  - Total end-site populations of tens of billions of end sites

  i.e. the total is Order (10**11)?

- The supply – inter-domain routing
  - We really may be stuck with BGP
  - 200,000 entries today
  - A billion entries ?

  i.e. a total is Order (10**7)

- The shoe horn
  - Aggregation and hierarchies in the address plan

Putting it together

Aggregation and hierarchies are not highly efficient addressing structures

The addressing plan needs to accommodate both large and small

The addressing plan needs to be simple

16 bit subnets + HD = 0.8 + global populations + 60 years =?

HD Ratio for Bigger Networks

Prefix	/48 count	end-site count
/21	134,217,728	3,178,688
/20	268,435,456	5,534,417
/19	536,870,912	9,635,980
/18	1,073,741,824	16,777,216
/17	2,147,483,648	29,210,830
/16	4,294,967,296	50,859,008
/15	8,589,934,592	88,550,677
/14	17,179,869,184	154,175,683
/13	34,359,738,368	268,435,456
/12	68,719,476,736	467,373,275
/11	137,438,953,472	813,744,135
/10	274,877,906,944	1,416,810,831
/9	549,755,813,888	2,466,810,934
/8	1,099,511,627,776	4,294,967,296
/7	2,199,023,255,552	7,477,972,398
/6	4,398,046,511,104	13,019,906,166
/5	8,796,093,022,208	22,668,973,294

Multiplying it out

A possible consumption total:

very simple address plan (16 bit subnets)
x aggregation factor (HD = 0.8)
x global populations (10**11)
x 60 years time frame

= /1 -- /4 range

Where's the wriggle room?

The HD ratio

- 0.8 = 1
- 0.87 = 0.5
- 0.94 = 0.1
- i.e. moving to a higher HD ratio will recover 3 bits here

The subnet field

- /56 will recover 8 bits
- Variable subnetting can recover between 0 and 16 bits

= /10 -- /17 range total

Varying the HD Ratio

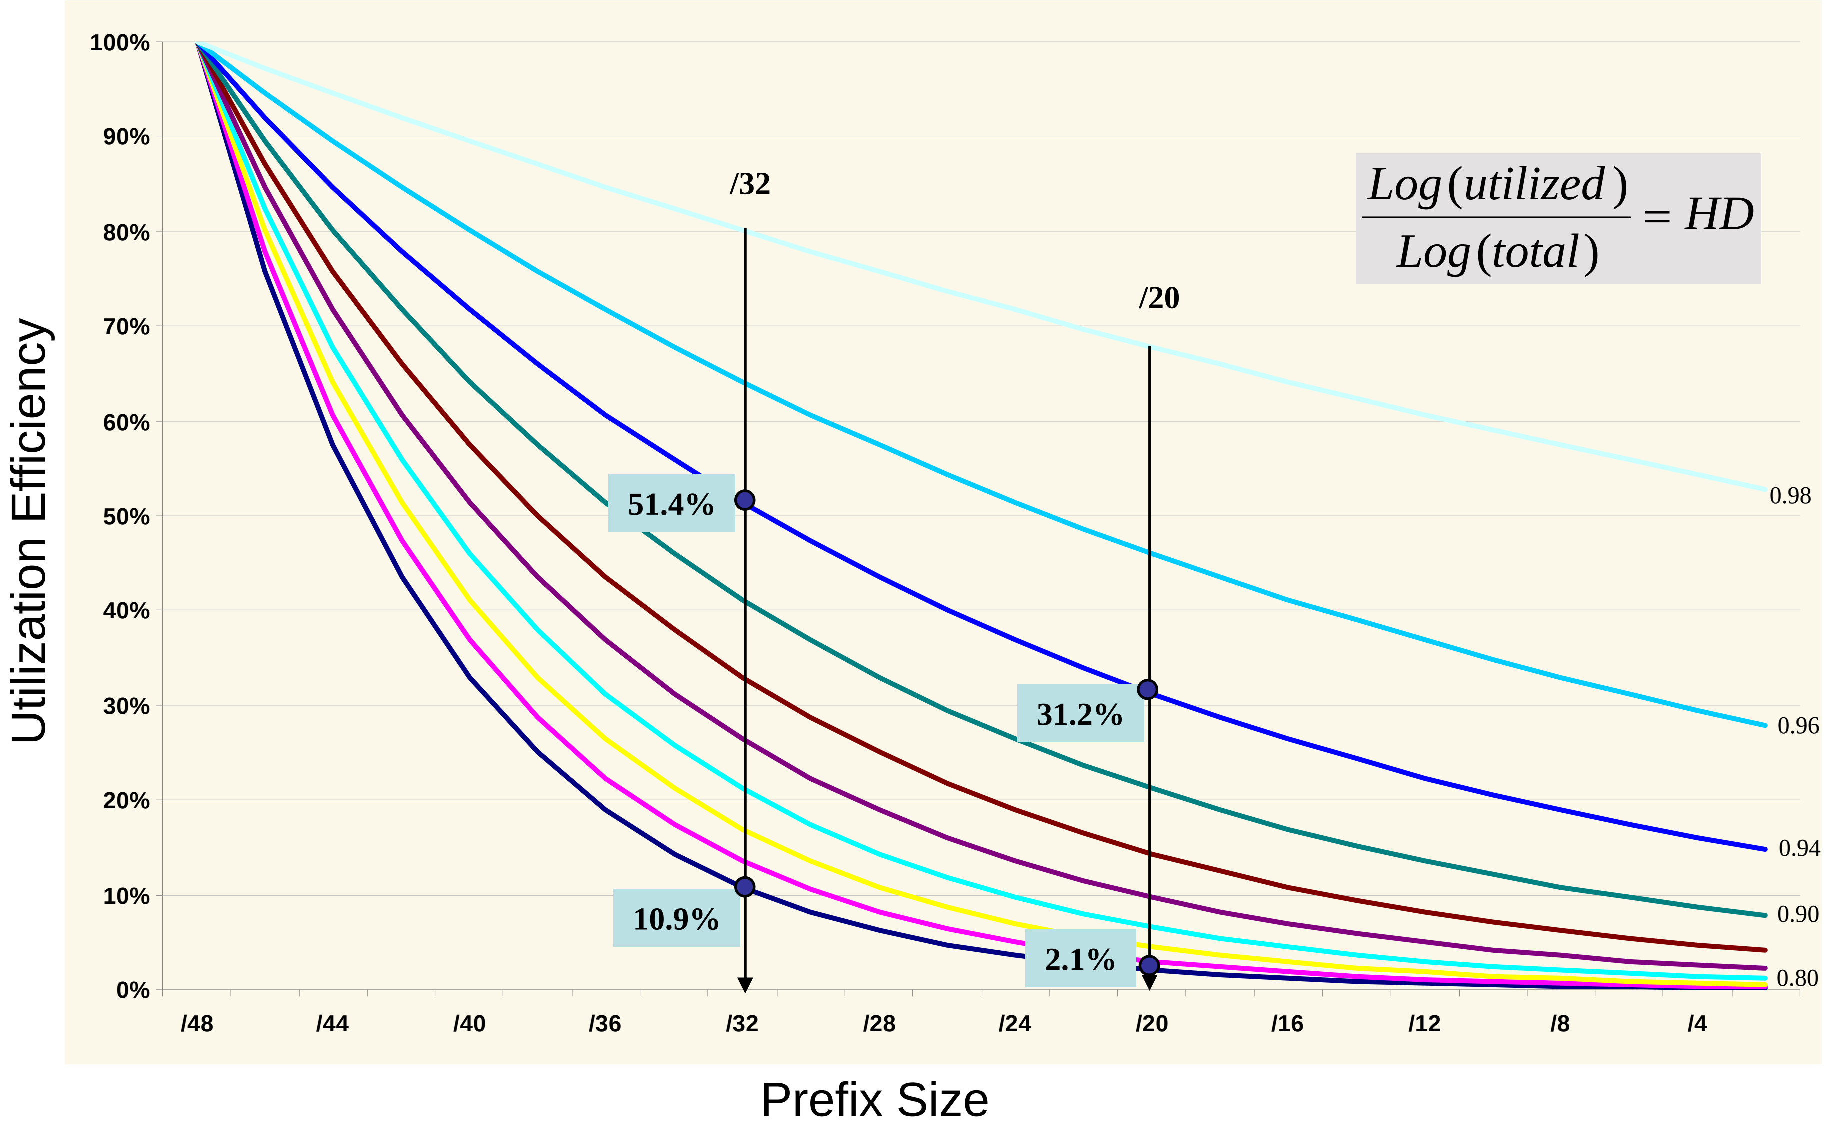

$$\frac{Log(utilized)}{Log(total)} = HD$$

Comparison of prefix size distributions from V6 registry

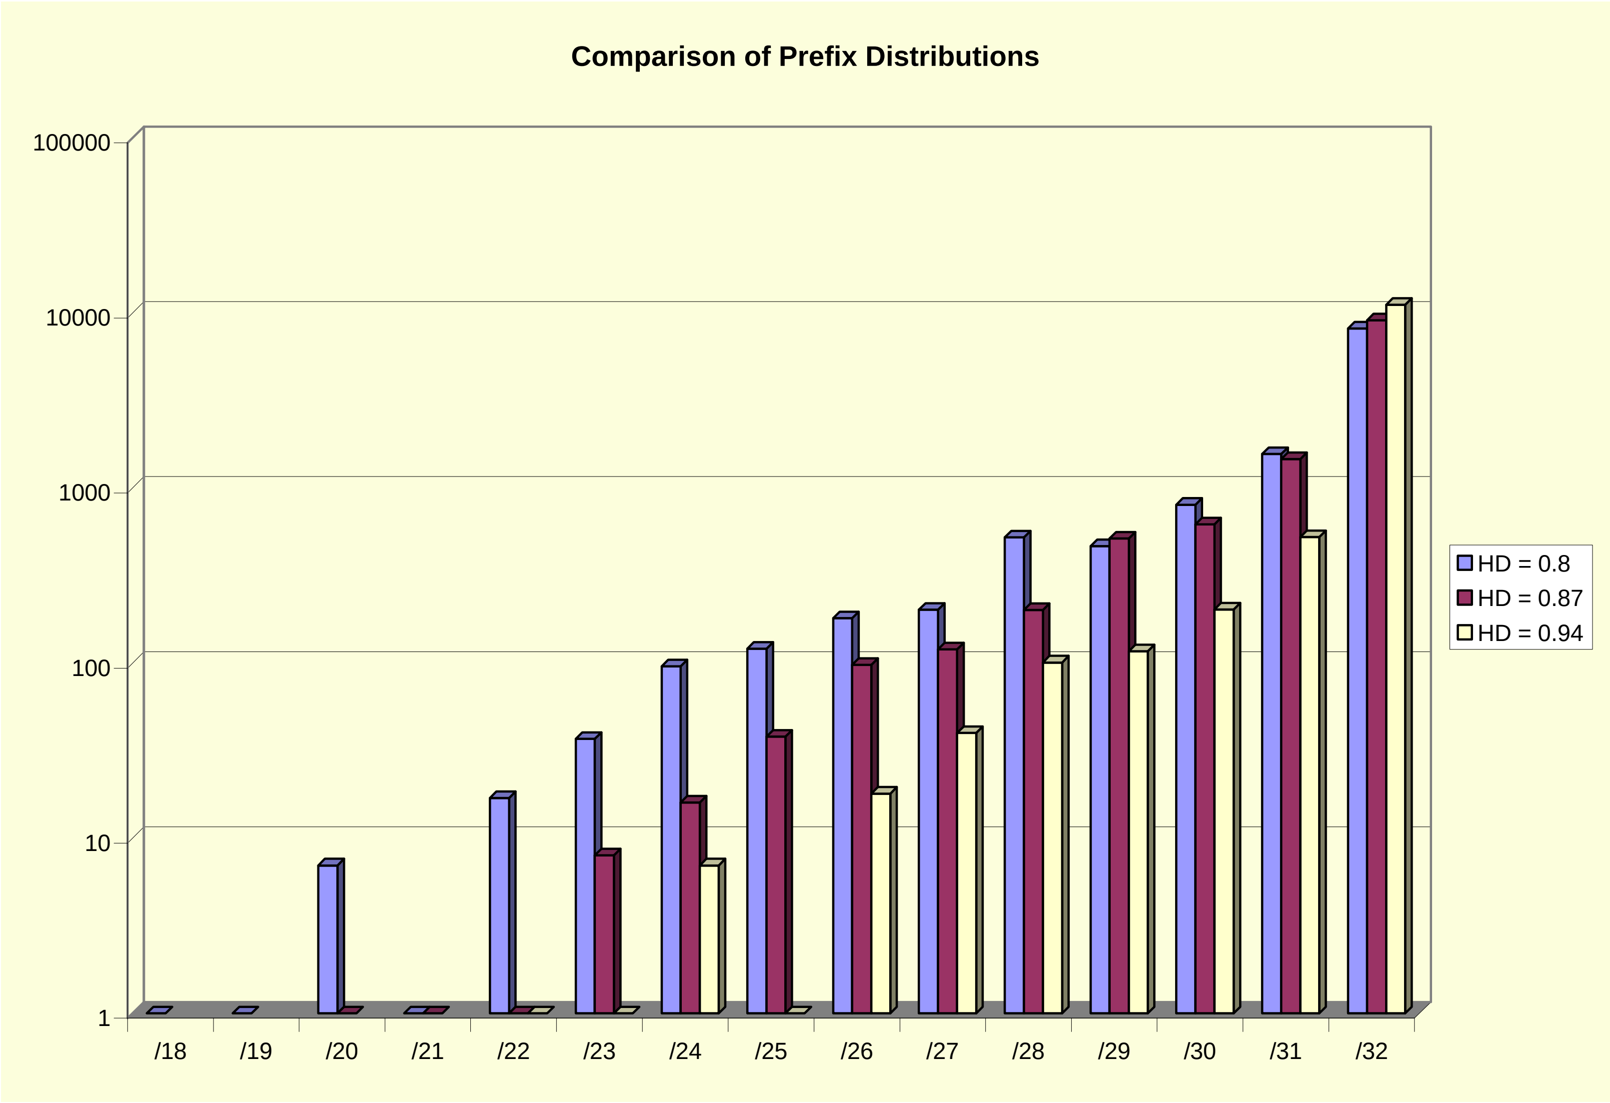

Observations

- **80% of all allocations are /31, /32 for HD ratio of 0.8 or higher**
  - Changing the HD ratio will not impact most allocations in a steady state registry function
- **Only 2% of all allocations are larger than a /27**
  - For these larger allocations the target efficiency is lifted from 4% to 25% by changing the HD Ratio from 0.8 to 0.94
- **Total 3 year address consumption is reduced by a factor of 10 in changing the HD ratio from 0.8 to 0.94**

What is a “good” HD Ratio to use?

- Consider **<u>what is common practice</u>** in today’s network in terms of internal architecture
  - APNIC is conducting a survey of ISPs in the region on network structure and internal levels of address hierarchy and will present the findings at APNIC 20
- Define a ***<u>common ‘baseline’ efficiency level</u>*** rather than an average attainable level
  - What value would be readily achievable by large and small networks without resorting to renumbering or unacceptable internal route fragmentation?
- Consider overall **<u>longer term objectives</u>**
  - Anticipated address pool lifetime
  - Anticipated impact on the routing space

Revisiting the /48 mantra?

- Concerned that current policy is quite wasteful
- NRPM 6.9 says: more than one subnet gets /48!!
  - Home users “should receive a /48”
  - This is from RFC3177
- Proposed feedback to IETF:
  - At least allow /56 assignments for SOHO
    - Change of two orders of magnitude in consumption
  - Consider /60 also…

IPv6 Address Reservations

Thomas Narten

narten@us.ibm.com

April 20, 2005

IPv6 Address Reservations

- When assigning prefix to LIR, RIR maintains “reservation” for future growth
- Goal: want subsequent assignment to be adjacent so a single aggregate covers old and subsequent allocation.
- Could be explicit reservation
- Could be implicit (e.g., in sparse allocation)
- Could be dynamic, based on growth prediction, etc.

IPv6 Address Reservations (cont.)

- Key issue: if insufficient room held for growth, address space will eventually fragment - we want to avoid this
- Time frame over which we want to preserve aggregation is O(decades) not O(months)
- Current proposal (2004-8) mentions reservations, but details completely unspecified
- Need to develop consensus recommendations that all RIRs can support and that achieve shared goal of preserving long-term

Questions?

ITU-T V6 Proposal

- Allocate each nation a contiguous V6 address block
- Establish national registries in each nation
- Promote competition between the national registries and the RIRs
- Allow LIRs / ISPs a choice of service entity between RIR and national registry

Some Attributes and Assumptions

- Addresses are a **global resource**
  - should be distributed between countries in a fair manner
- Addresses are a **public resource**
  - allows national public policy processes to set national address distribution policies
- Addresses are a **critical resource**
  - Establishes locally controlled address pools for each nation
- Addresses are a **network resource**
  - Without addresses network services are difficult to support
- Addresses are an **infinite resource**
  - There is enough address space to create 200 new registries with enough space for each such that all countries can agree

Some Issues

Allows for 200 different policy regimes and **policy confusion**

- "Recommendations" to sovereign national entities is ineffectual as a network control mechanism

Does not align to regional and **global business models**

- Does a global enterprise need to deal with up to 200 different address sources?

Has no visible relationship to **known routing capabilities**

- Route fragmentation at an entirely new level

Creates competition regimes based on **policy dilution**

- Creates impetus for rapid consumption, hoarding and address trading markets

Eliminates **common interest in one network**

- Places short term national interest well above common network interest

Compromises any hope to enhance **routing integrity and security**

- Eliminates hope for a robust and resilient trust hierarchy to support a viable secure network routing environment

Creates further churn in perceptions of **stability and viability of V6**

- Increases barriers to business investment in V6 infrastructure and services

Some Options

- Agree

  It’s a really good idea – go for it!
- Disagree

  It’s yet another really bad idea – go away!
- Discuss

  There are some valid assumptions here – but is there a way to do this that does not utterly destroy IPV6 at the same time?