HD Ratio for IPv4

- You get a new allocation when your entire hoard of IPv4 addresses is sufficiently utilized
- You must pass both HD ratio tests
  - last allocated block must be >= .930
  - total hoard must be >= .966
- HD ratio is log(utilized addresses) divided by log(entire hoard)

What does it mean?

Slash	Num Ips	Percentage at HD 0.966
24	256	82.8%
23	512	80.9%
22	1024	79.0%
21	2048	77.2%
20	4096	75.4%
19	8192	73.6%
18	16384	71.9%
17	32768	70.2%
16	65536	68.6%
15	131072	67.0%
14	262144	65.4%
13	524288	63.9%
12	1048576	62.4%
11	2097152	61.0%
10	4194304	59.5%
9	8388608	58.2%
8	16777216	56.8%

Old Threshold of 80%

New policy kicks in for providers with /20 or more

What does it mean?

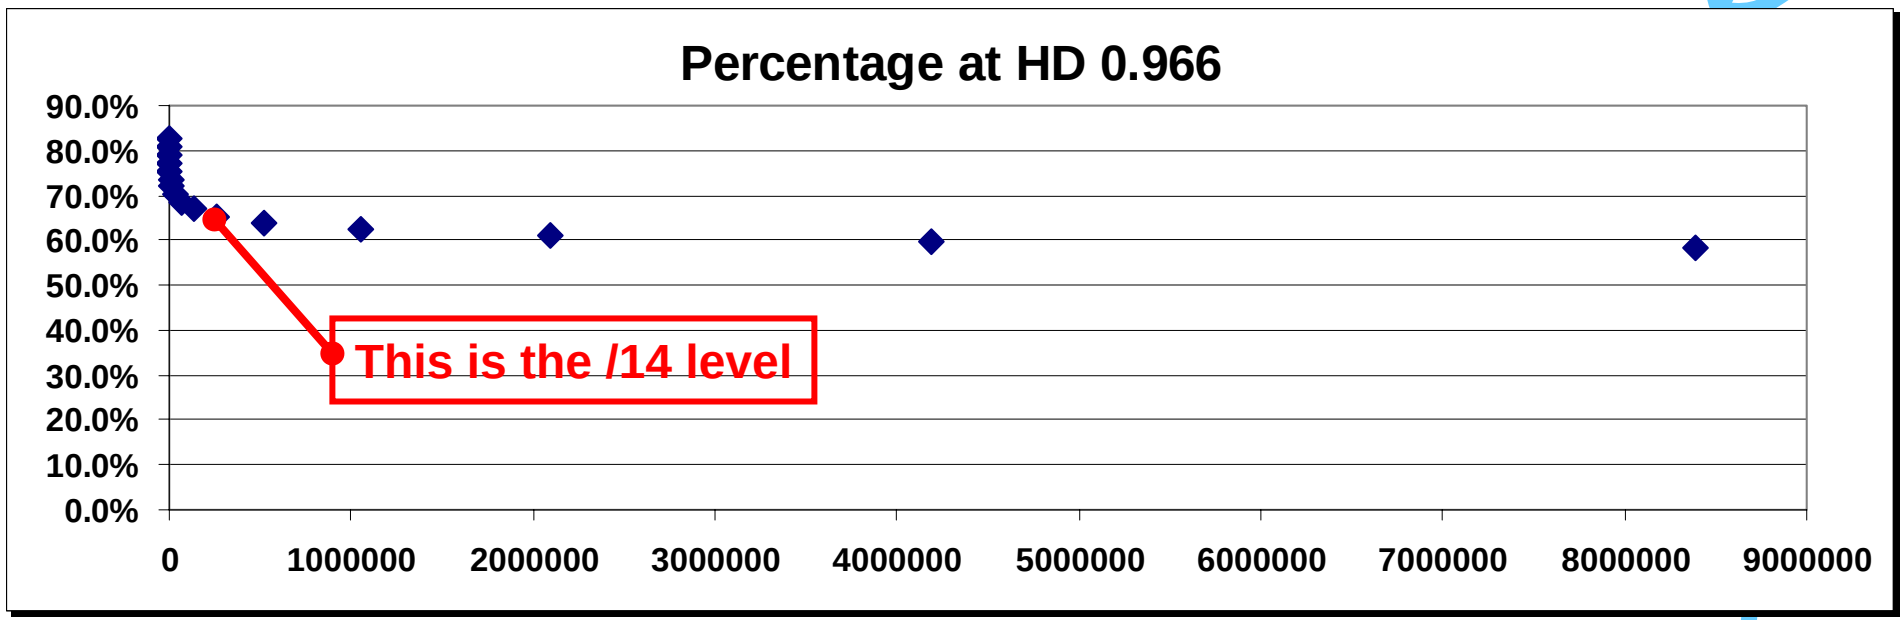

- Benefit increases rapidly as an ISP grows to mid-sized
- Not just for the few big guys

What does it mean?

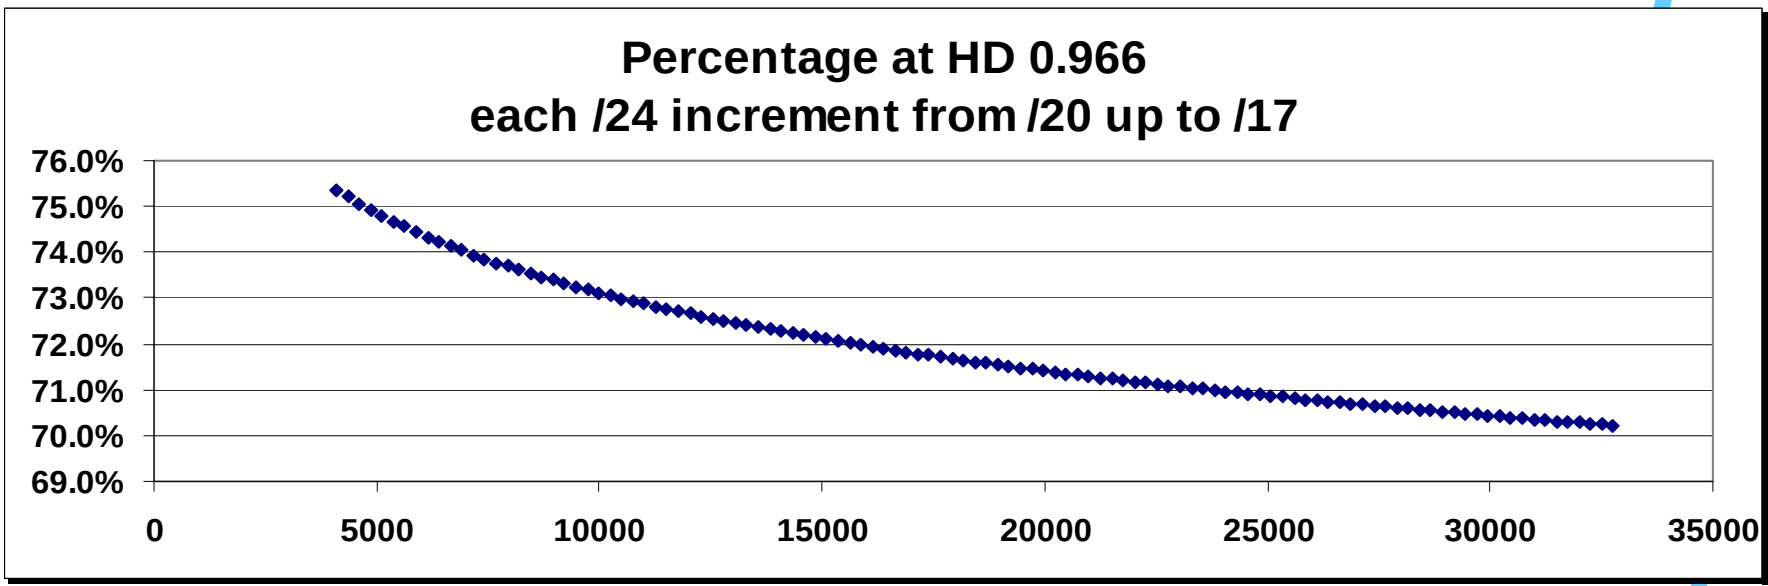

- The smaller ISPs will only need between 70% and 75% utilization rates

What does it mean?

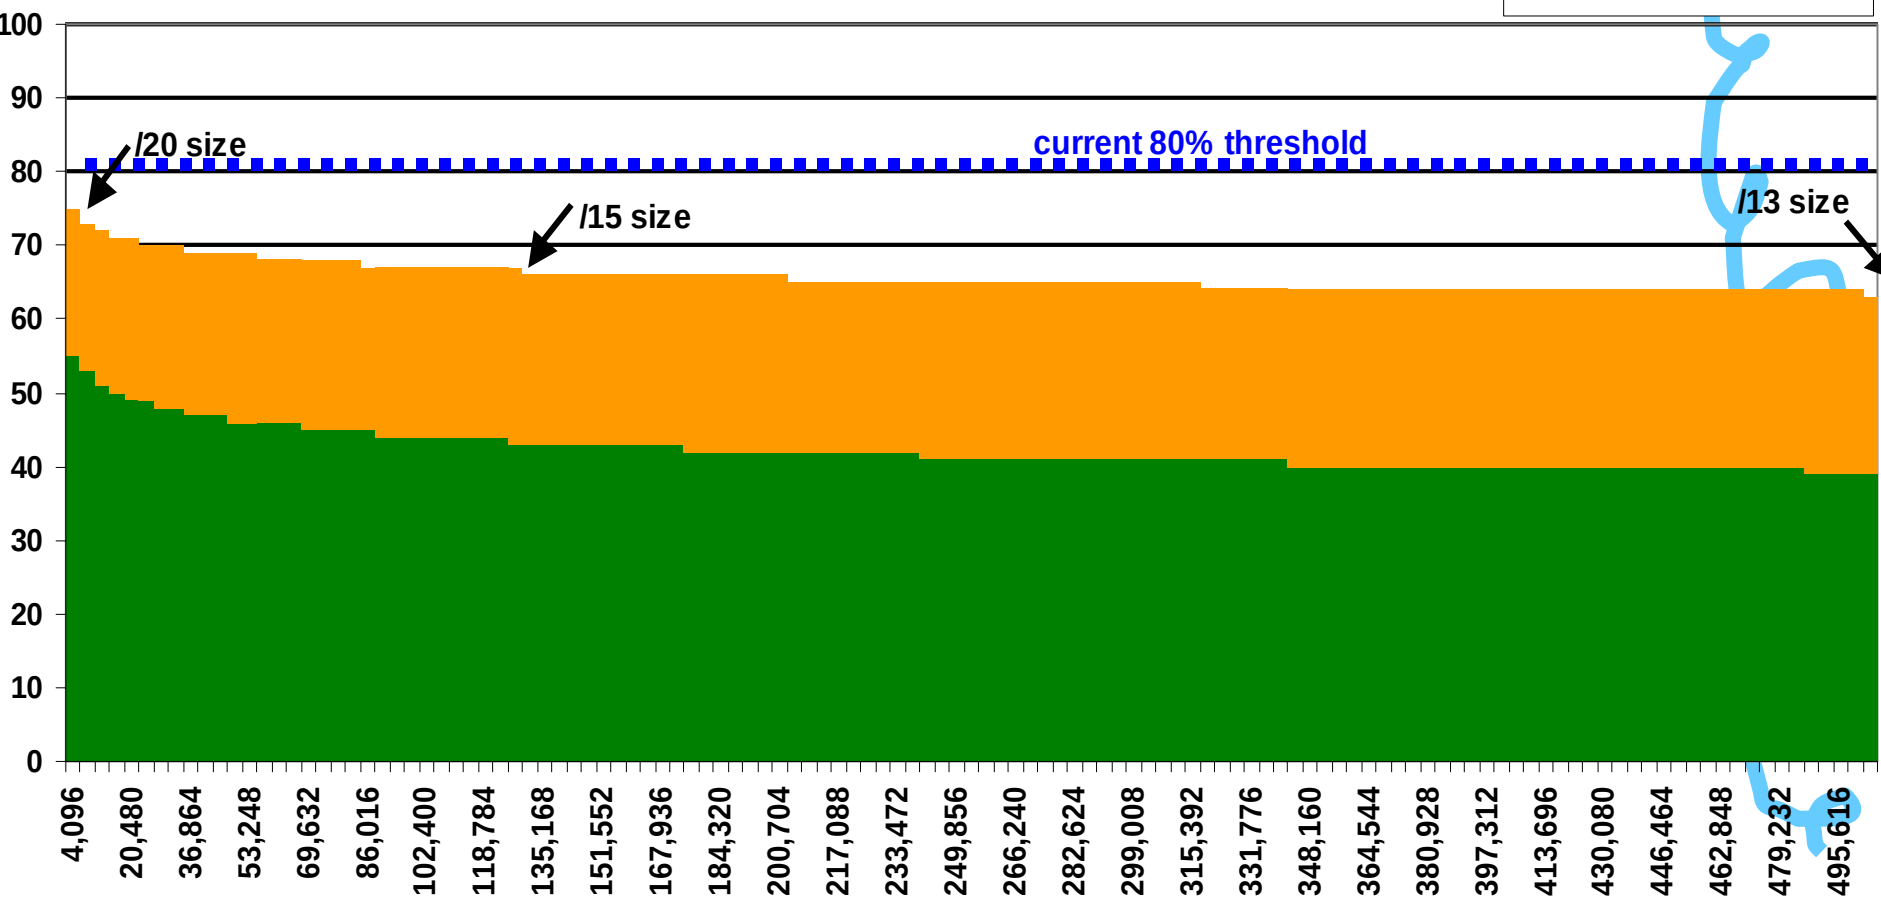

Why, why, why?

- Why do you have to count utilization of all previous blocks?
- Because the HD ratio calculation is meant to recognize that overhead increases as a network grows larger. We need to count the ISP’s entire hoard of IPv4 addresses in order to properly make allowance for this overhead.

Why, why, why?

- Why did you choose .966 for the HD ratio magic number?
- Mainly because I felt that Paul Wilson's argument made sense in his discussion paper for APNIC http://www.apnic.net/mailing-lists/sig-policy/archive/2003/08/msg00000.html

Why, why, why?

- Why set a separate .930 ratio for the previous allocation?
- This is to simplify ISP planning. When they receive a new block, an ISP can calculate the .930 ratio for it and use this as a target for planning. Only ISPs with significant churn of older addresses will need to closely track the overall .966 ratio.

Why, why, why?

- But why be more lenient with .930 ratio?
- We believe that it is prudent for network operators to have change control processes and we recognize that such processes take time. The .930 ratio demands that ISPs show they are indeed deploying the last allocation but it avoids unnecessary brinkmanship with their change control timing.

Why, why, why?

- Why are you using a complicated natural logarithm?
- A logarithmic ratio is used because that creates a sliding scale that recognizes the sliding scale of overhead inherent in hierarchical allocation of addresses. The natural logarithm was chosen because it is supported by both PERL and Excel.

Why, why, why?

- Why do you care so much about PERL?
- While I wish the world would throw away all PERL code and replace it with Python or Ruby, I recognize that it is a major tool for ISPs and many IP address management packages are written in PERL. Similarly, many small ISPs manage IP addresses using spreadsheets like Excel.

Understanding the formula

- Given a range of IPv4 addresses
- Let utiladdr = the number of IPv4 addresses utilized out of the range
- Let totaddr = the total number of IPv4 addresses in the range
- In Excel: HD = ln(utiladdr)/ln(totaddr)
- In PERL: HD = log(utiladdr)/log(totaddr)

Understanding the Formula

- Basic ratio calculation.

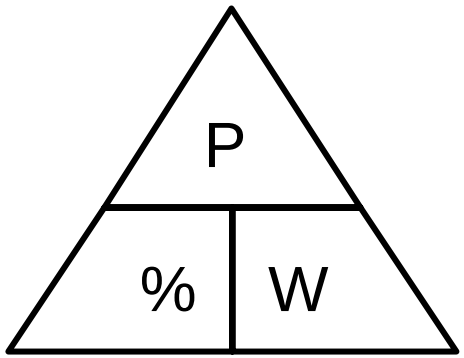

% - Percentage
P - Part
W - Whole

★Percentage = Part / Whole

★Part = Percentage * Whole

★Whole = Part / Percentage

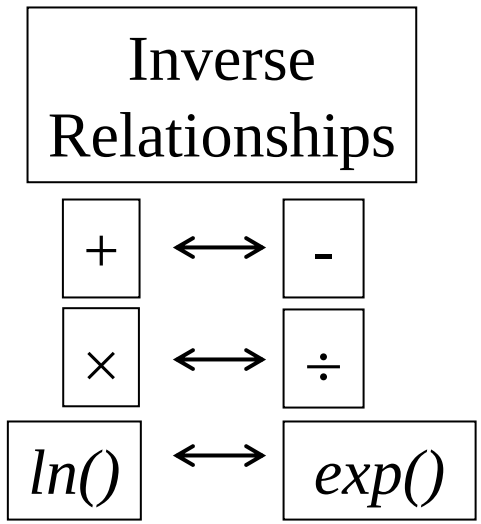

Key formulas are marked with★

Understanding the Formula

- HD Ratio is a ratio just like percentage.

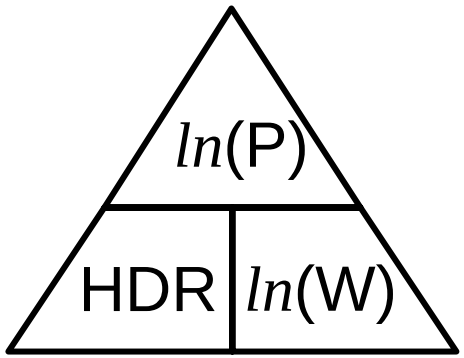

★HDRatio = *ln*(Part) / *ln*(Whole)

ln(Part) = HDRatio * *ln*(Whole)
ln(Whole) = *ln*(Part) / HDRatio

exp() inverts *ln*() → Part = *exp*(*ln*(Part))
Whole = *exp*(*ln*(Whole))

therefore:

★Part = *exp*(HDRatio * *ln*(Whole))
★Whole = *exp*(*ln*(Part) / HDRatio)

Key formulas are marked with★