

Network Working Group
Request for Comments: 3194
Updates: 1715
Category: Informational

A. Durand
SUN Microsystems
C. Huitema
Microsoft
November 2001

The Host-Density Ratio for Address Assignment Efficiency:
An update on the H ratio

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2001). All Rights Reserved.

Abstract

This document provides an update on the "H ratio" defined in RFC 1715. It defines a new ratio which the authors claim is easier to understand.

1. Evaluating the efficiency of address allocation

A naive observer might assume that the number of addressable objects in an addressing plan is a linear function of the size of the address. If this were true, a telephone numbering plan based on 10 digits would be able to number 10 billion telephones, and the IPv4 32 bit addresses would be adequate for numbering 4 billion computers (using the American English definition of a billion, i.e. one thousand millions.) We all know that this is not correct: the 10 digit plan is stressed today, and it handles only a few hundred million telephones in North America; the Internet registries have started to implement increasingly restrictive allocation policies when there were only a few tens of million computers on the Internet.

Addressing plans are typically organized as a hierarchy: in telephony, the first digits will designate a region, the next digits will designate an exchange, and the last digits will designate a subscriber within this exchange; in computer networks, the most significant bits will designate an address range allocated to a network provider, the next bits will designate the network of an organization served by that provider, and then the subnet to which the individual computers are connected. At each level of the

hierarchy, one has to provide some margins: one has to allocate more digits to the region code than the current number of regions would necessitate, and more bits in a subnet than strictly required by the number of computers. The number of elements in any given level of the hierarchy will change over time, due to growth and mobility. If the current allocation is exceeded, one has to engage in renumbering, which is painful and expensive. In short, trying to squeeze too many objects into a hierarchical address space increases the level of pain endured by operators and subscribers.

Back in 1993, when we were debating the revision of the Internet Protocol, we wondered what the acceptable ratio of utilization was of a given addressing plan. Coming out with such a ratio was useful to assess how many computers could be connected to the Internet with the current 32-bit addresses, as well as to decide the size of the next generation addresses. The second point is now decided, with 128-bits addresses for IPv6, but the first question is still relevant: knowing the capacity of the current address plan will help us predict the date at which this capacity will be exceeded.

Participants in the IPNG debates initially measured the efficiency of address allocation by simply dividing the number of allocated addresses by the size of the address space. This is a simple measure, but it is largely dependent on the size of the address space. Loss of efficiency at each level of a hierarchical plan has a multiplicative effect; for example, 50% efficiency at each stage of a three level hierarchy results in a overall efficiency of 12.5%. If we want a "pain level indicator", we have to use a ratio that takes into account these multiplicative effects.

The "H-Ratio" defined in RFC 1715 proposed to measure the efficiency of address allocation as the ratio of the base 10 logarithm of the number of allocated addresses to the size of the address in bits. This provides an address size independent ratio, but the definition of the H ratio results in values in the range of 0.0 to 0.30103, with typical values ranging from 0.20 to 0.28. Experience has shown that these numbers are difficult to explain to others; it would be easier to say that "your address bits are used to 83% of their H-Density", and then explain what the H-Density is, than to say "you are hitting a H ratio of 0.25" and then explain what exactly the range is.

This memo introduces the Host Density ratio or "HD-Ratio", a proposed replacement for the H-Ratio defined in RFC 1715. The HD values range from 0 to 1, and are generally expressed as percentage points; the authors believe that this new formulation is easier to understand and more expressive than the H-Ratio.

2. Definition of the HD-ratio

When considering an addressing plan to allocate objects, the host density ratio HD is defined as follow:

$$HD = \frac{\log(\text{number of allocated objects})}{\log(\text{maximum number of allocatable objects})}$$

This ratio is defined for any number of allocatable objects greater than 1 and any number of allocated objects greater or equal than 1 and less than or equal the maximum number of allocatable objects. The ratio is usually presented as a percentage, e.g. 70%. It varies between 0 (0%), when there is just one allocation, and 1 (100%), when there is one object allocated to each available address. Note that for the calculation of the HD-ratio, one can use any base for the logarithm as long as it is the same for both the numerator and the denominator.

The HD-ratio can, in most cases, be derived from the H ratio by the formula:

$$HD = \frac{H}{\log_{10}(2)}$$

3. Using the HD-ratio as an indicator of the pain level

In order to assess whether the H-Ratio was a good predictor of the "pain level" caused by a specific efficiency, RFC1715 used several examples of networks that had reached their capacity limit. These could be for example telephone networks at the point when they decided to add digits to their numbering plans, or computer networks at the point when their addressing capabilities were perceived as stretched beyond practical limits. The idea behind these examples is that network managers would delay renumbering or changing the network protocol until it became just too painful; the ratio just before the change is thus a good predictor of what can be achieved in practice. The examples were the following:

* Adding one digit to all French telephone numbers, moving from 8 digits to 9, when the number of phones reached a threshold of 1.0 E+7.

$$\text{HD(FrenchTelephone8digit)} = \frac{\log(1.0\text{E}+7)}{\log(1.0\text{E}+8)} = 0.8750 = 87.5\%$$

$$\text{HD(FrenchTelephone9digit)} = \frac{\log(1.0\text{E}+7)}{\log(1.0\text{E}+9)} = 0.7778 = 77.8\%$$

* Expanding the number of areas in the US telephone system, making the phone number effectively 10 digits long instead of "9.2" (the second digit of area codes used to be limited to 0 or 1) for about 1.0 E+8 subscribers.

$$\text{HD(USTelephone9.2digit)} = \frac{\log(1.0\text{E}+8)}{\log(9.5\text{E}+9)} = 0.8696 = 87.0 \%$$

$$\text{HD(USTelephone10digit)} = \frac{\log(1.0\text{E}+8)}{\log(1\text{E}+10)} = 0.8000 = 80.0 \%$$

* The globally-connected physics/space science DECnet (Phase IV) stopped growing at about 15K nodes (i.e. new nodes were hidden) in a 16 bit address space.

$$\text{HD(DecNET IV)} = \frac{\log(15000)}{\log(2^{16})} = 0.8670 = 86.7 \%$$

From those examples, we can note that these addressing systems reached their limits for very close values of the HD-ratio. We can use the same examples to confirm that the definition of the HD-ratio as a quotient of logarithms results in better prediction than the direct quotient of allocated objects over size of the address space. In our three examples, the direct quotients were 10%, 3.2% and 22.8%, three very different numbers that don't lead to any obvious generalization. The examples suggest an HD-ratio value on the order of 85% and above correspond to a high pain level, at which operators are ready to make drastic decisions.

We can also examine our examples and hypothesize that the operators who renumbered their networks tried to reach, after the renumbering, a pain level that was easily supported. The HD-ratio of the French or US network immediately after renumbering was 78% and 80%, respectively. This suggests that values of 80% or less corresponds to comfortable trade-offs between pain and efficiency.

4. Using the HD-ratio to evaluate the capacity of addressing plans

Directly using the HD-ratio makes it easy to evaluate the density of allocated objects. Evaluating how well an addressing plan will scale requires the reverse calculation. We have seen in section 3.1 that an HD-ratio lower than 80% is manageable, and that HD-ratios higher than 87% are hard to sustain. This should enable us to compute the acceptable and "practical maximum" number of objects that can be allocated given a specific address size, using the formula:

```
number allocatable of objects
    = exp( HD x log(maximum number allocatable of objects))
    = (maximum number allocatable of objects)^HD
```

The following table provides example values for a 9-digit telephone plan, a 10-digit telephone plan, and the 32-bit IPv4 Internet:

| | Reasonable HD=80% | Painful HD=85% | Very Painful HD=86% | Practical Maximum HD=87% |
|-------------------|----------------------|-------------------|------------------------|--------------------------------|
| ----- | ----- | ----- | ----- | ----- |
| 9-digits plan | 16 M | 45 M | 55 M | 68 M |
| 10-digits plan | 100 M | 316 M | 400 M | 500 M |
| 32-bits addresses | 51 M | 154 M | 192 M | 240 M |

Note: 1M = 1,000,000

Indeed, the practical maximum depends on the level of pain that the users and providers are willing to accept. We may very well end up with more than 154M allocated IPv4 addresses in the next years, if we are willing to accept the pain.

5. Security considerations

This document has no security implications.

6. IANA Considerations

This memo does not request any IANA action.

7. Author addresses

Alain Durand
SUN Microsystems, Inc
901 San Antonio Road MPK17-202
Palo Alto, CA 94303-4900
USA

EMail: Alain.Durand@sun.com

Christian Huitema
Microsoft Corporation
One Microsoft Way Redmond, WA 98052-6399
USA

EMail: huitema@microsoft.com

8. Acknowledgment

The authors would like to thank Jean Daniau for his kind support during the elaboration of the HD formula.

9. References

- [RFC1715] Huitema, C., "The H Ratio for Address Assignment Efficiency", RFC 1715, November 1994.
- [IANAV4] INTERNET PROTOCOL V4 ADDRESS SPACE, maintained by the IANA, <http://www.iana.org/assignments/ipv4-address-space>
- [DMNSRV] Internet Domain Survey, Internet Software Consortium, <http://www.isc.org/ds/>
- [NETSZR] Netsizer, Telcordia Technologies, <http://www.netsizer.com/>

10. Full Copyright Statement

Copyright (C) The Internet Society (2001). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

