LNAT - Large scale IP via Network Address Translation

Van Jacobson
January 20, 1992

the proposal

1 Introduction

The IP protocol suite requires that communicating hosts be identified by globally unique, 32-bit addresses. The Internet protocol suite has been wildly successful (to the tune of several million interconnected hosts at present) and, consequently, the 32-bit IP address space is rapidly being exhausted. There are two ways to remedy this:

1. Convert to a protocol with larger addresses such as ISO NSAPs in CLNP.

2. Make IP addresses 'locally' unique rather than globally unique.

This proposal describes the latter alternative. IP addresses are unique within a region (an *addressing domain*) and are mapped to something else (either another region's addresses or to global identifiers with a different encapsulation) at the boundary of the region.

the rationale

It is important to realize that the two alternatives above are not mutually exclusive. In particular, approach (1) will almost certainly contain approach (2) as a component: Approach (1) requires (eventually) converting all hosts to a new protocol. Several million hosts can communicate today and expect to be able to communicate tomorrow. Converting this many existing hosts to a new protocol will take many years. Any viable transition plan must include some means to let unconverted and converted hosts interoperate, uninterrupted, during the entire period of the conversion. The author contends that the mapping mechanism introduced in this proposal is such a transition plan and could be used as the medium term vehicle to migrate to a new, long term change in address structure.[1] However, it is the intent of this note to point out that address mapping *by itself* is a solution to the long term addressing problem and it may not be necessary to do anything else.

the rationale

It is important to realize that the two alternatives above are not mutually exclusive. In particular, approach (1) will almost certainly contain approach (2) as a component: Approach (1) requires (eventually) converting all hosts to a new protocol. Several million hosts can communicate today and expect to be able to communicate tomorrow. Converting this many existing hosts to a new protocol will take many years. Any viable transition plan must include some means to let unconverted and converted hosts interoperate, uninterrupted, during the entire period of the conversion. The author contends that the mapping mechanism introduced in this proposal is such a transition plan and could be used as the medium term vehicle to migrate to a new, long term change in address structure.[1] However, it is the intent of this note to point out that address mapping *by itself* is a solution to the long term addressing problem and it may not be necessary to do anything else.

the model

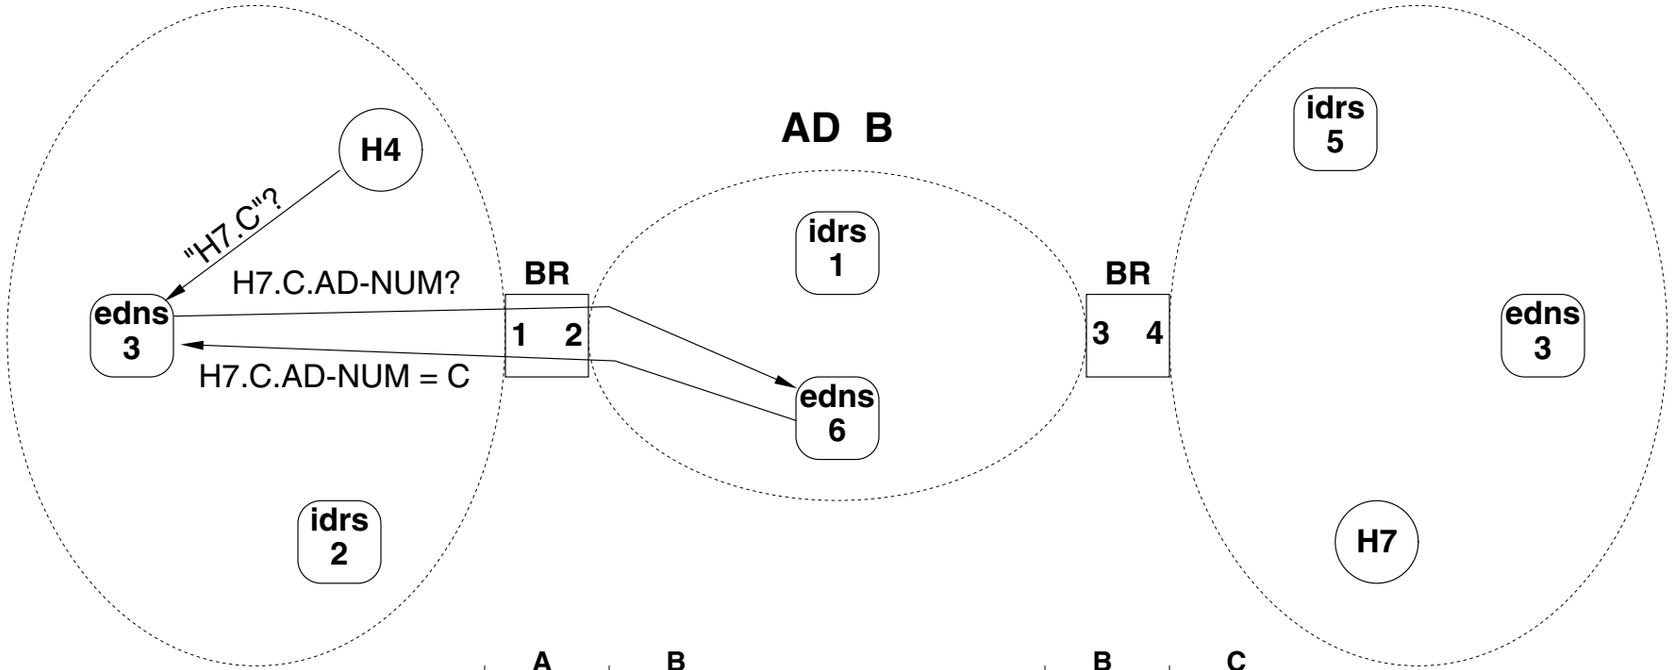

	A	B
"edns6.B"	1.0.53	0.0.6
"edns3.A"	0.0.3	2.0.53

	B	C
"edns6.B"	0.0.6	4.0.53
"edns3.C"	3.0.53	0.0.3

Host H4.A requests IP number for host H7.C. Request percolates through DNS in A and eventually ends up at A's EDNS (**edns3**). First step for EDNS is to find out which AD host H7.C lives in.

what did Jon think?

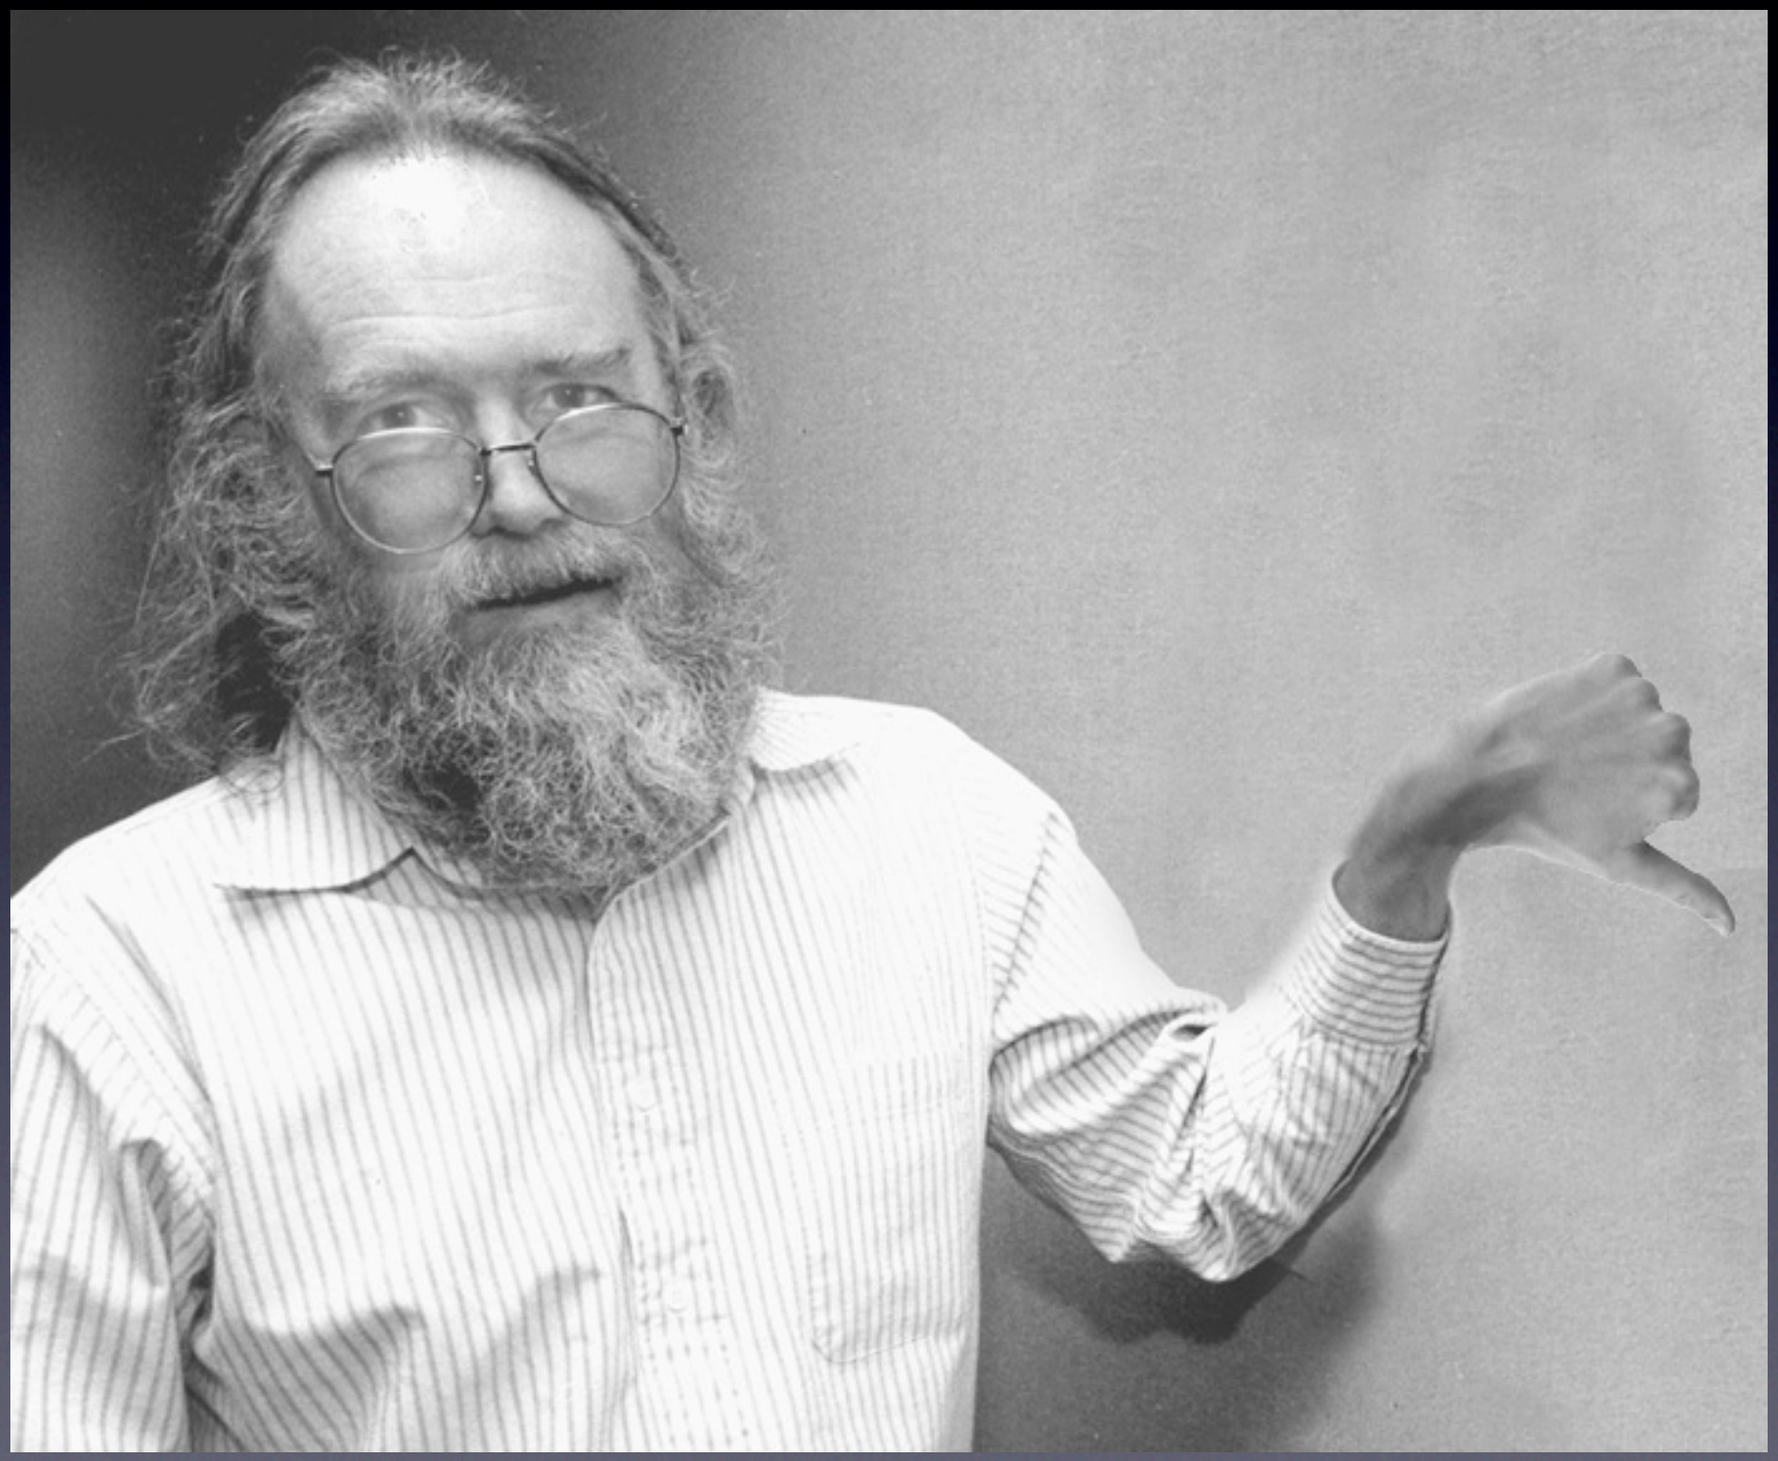

NAT does violence to the core simplicity of the Internet architecture (see RFC-3439)

- It creates new, transient, address mapping state in the core (but the state lifetime can be long if you embed a small number of hosts into a large address space).
- Addresses show up in many places besides packet source & destination fields. It's hard to find and fix all of them.